A RAND NOTE

Policy Analysis of Water Management for the Netherlands

Vol. V, Design of Managerial Strategies

J. H. Bigelow

May 1982

N-1500/5-NETH

Prepared for

The Netherlands Rijkswaterstaat

Rand
SANTA MONICA, CA. 90406

PREFACE

For some time the Netherlands has had a problem with water quality, particularly salinity, eutrophication, and thermal pollution. Moreover, the future demand for fresh water is expected to exceed the supply. The growing demand for the limited supply of groundwater is leading to increased competition among its users: agriculture, industry, nature preserves, and companies that supply drinking water. The supply of surface water is sufficient except in dry years, when there is competition not only among such users as agriculture, power plants, and shipping, but also among different regions.

Facing such water management problems, the Dutch government wanted an analysis to help draft the first national water management law and to select the overall water management policy for the Netherlands. It established the Policy Analysis for the Water Management of the Netherlands (PAWN) Project in August 1976 as a joint research project of Rand (a nonprofit corporation),[1] the Rijkswaterstaat (the government agency responsible for water control and public works),[2] and the Delft Hydraulics Laboratory (a leading Dutch research organization).[3]

The primary tasks of the PAWN project were to:

1. Develop a methodology for assessing the multiple consequences of water management policies.
2. Apply it to develop alternative water management policies[4] for the Netherlands and to assess and compare their consequences.
3. Create a Dutch capability for further such analyses by training Dutch analysts and by documenting and transferring methodology developed at Rand to the Netherlands.

The methodology and results of the PAWN project are described in a series of publications entitled Policy Analysis of Water Management for the Netherlands. The series contains the following volumes:

- Volume I, Summary Report (Rand R-2500/1)
- Volume II, Screening of Technical and Managerial Tactics (Rand N-1500/2)
- Volume III, Screening of Eutrophication Control Tactics (Rand N-1500/3)
- Volume IV, Design of Long-Run Pricing and Regulation Strategies (Rand N-1500/4)
- Volume V, Design of Managerial Strategies (Rand N-1500/5)
- Volume VA, Methodological Appendixes to Vol. V (Rand N-1500/5A)
- Volume VI, Design of Eutrophication Control Strategies (Rand N-1500/6)

- Volume VII, Assessment of Impacts on Drinking-Water Companies and Their Customers (Rand N-1500/7)
- Volume VIII, Assessment of Impacts on Industrial Firms (Rand N-1500/8)
- Volume IX, Assessment of Impacts on Shipping and Lock Operation (Rand N-1500/9)
- Volume X, Distribution of Monetary Benefits and Costs (Rand N-1500/10)
- Volume XI, Water Distribution Model (Rand N-1500/11)
- Volume XII, Model for Regional Hydrology, Agricultural Water Demands and Damages from Drought and Salinity (Rand N-1500/12)
- Volume XIII, Models for Sprinkler Irrigation System Design, Cost, and Operation (Rand N-1500/13)
- Volume XIV, Optimal Distribution of Agricultural Irrigation Systems (Rand N-1500/14)
- Volume XV, Electric Power Reallocation and Cost Model (Rand N-1500/15)
- Volume XVI, Costs for Infrastructure Tactics (Rand N-1500/16)
- Volume XVII, Flood Safety Model for the IJssel Lakes (Rand N-1500/17)
- Volume XVIII, Sedimentation and Dredging Cost Models (Rand N-1500/18)
- Volume XIX, Models for Salt Intrusion in the Rhine Delta (Rand N-1500/19)
- Volume XX, Industry Response Simulation Model (Rand N-1500/20)

Four comments about this series of publications seem appropriate. First, the series represents a joint Rand/Rijkswaterstaat/Delft Hydraulics Laboratory research effort. Whereas only some of the volumes list Dutch coauthors, all have Dutch contributors, as can be seen from the acknowledgments pages.

Second, except where noted, these publications describe the methodology and results presented at the final PAWN briefing at Delft on December 11 and 12, 1979. For Rand, this briefing marked the beginning of the documentation phase of the project and the end of the analysis phase. Rand and the Rijkswaterstaat (RWS) considered the results to be tentative because (1) some of the methodology had not become available until late in the analysis phase, and (2) the RWS planned to do additional analysis.

Third, the RWS is preparing its Nota Waterhuishouding, the new policy document on water management scheduled for publication in 1982, by combining some of the PAWN results from December 1979 with the results of considerable additional analysis done in the Netherlands with the PAWN methodology. Because the understanding gained in the original analysis led to improvements in the data--and, in some instances, the models--used to represent the water management system in the additional analysis, the reader is hereby cautioned that the numerical results and conclusions presented in the PAWN volumes will not always agree with those presented in the Nota Waterhuishouding or its companion reports. (It has not been possible to indicate such differences in the volumes since they are being written before the Nota is published.) Thus, the present series of publications puts primary emphasis on documenting the methodology rather than on describing the policy results.

Fourth, Vols. II through XX are not intended to stand alone, and should be read in conjunction with the Summary Report (Vol. I), which contains most of the contextual and evaluative material.

The present volume describes the Managerial Strategy Design Model (MSDM), and the results we have obtained with it. Given a fixed set of facilities, MSDM determines a managerial strategy[4] that yields the lowest possible total cost to all users of water. MSDM relies heavily on many of the other volumes in this series to estimate the costs that a managerial strategy imposes on each user, particularly Vol. IX (for shipping), Vols. XII and XIII (for agriculture), Vol. XV (for electric power generation), and Vol. XVI (for the cost of pumping energy). In addition, we have used MSDM to develop a managerial strategy used in the Water Distributionn Model, described in Vol. XI. Finally, we have relegated much technical detail concerning MSDM to Vol. VA.

Those interested in one or another aspect of the day-to-day management of a fixed set of facilities can read all or part of this volume. For a discussion of what managerial tactics are possible, read Chap. 2. For a description of the minimum-cost strategy, plus other results and our conclusions, read Chaps. 9, 10, and 11. And for a deeper discussion of the effects of managerial tactics on a particular water use or user, read the chapter dealing with that use or user--e.g., improving water quality (Chap. 3); reducing salinity (Chap. 4); cooling electric generating plants (Chap. 5). We do hope, however, that some readers will be sufficiently interested in the entire problem to read the whole volume.

NOTES

1. Rand had had extensive experience with similar kinds of analysis and had been working with the Rijkswaterstaat for several years on other problems.
2. The Rand contract was officially with the Rijkswaterstaat, Directie Waterhuishouding en Waterbeweging (Directorate for Water Management and Water Movement), but numerous other parts of the Rijkswaterstaat contributed to the analysis.
3. Delft Hydraulics Laboratory research was performed under project number R1230, sponsored by the Netherlands Rijkswaterstaat.
4. Each water management policy involved a mix of tactics, each a particular action to affect water management, such as building a particular canal or taxing a particular use. Four kinds of tactics were considered: building new water management facilities (infrastructure) or applying various treatments to the water (called technical tactics); using managerial measures (called

managerial tactics) to change the distribution of water among competing regions and users; and imposing taxes or quotas to affect the quantity or quality of water extracted or discharged by different users (called *price and regulation tactics*, respectively). A mix of tactics of the same kind is called a *strategy*. Thus, the overall policy could be conceived as a combination of technical, managerial, pricing, and regulation strategies.

SUMMARY

S.1. THE MANAGERIAL STRATEGY DESIGN PROBLEM

This volume investigates the problem of designing managerial strategies, which we can state as: Determine what day-to-day actions will bring about the distribution of surface water most beneficial to all water users and uses. We address this problem in a multitude of contexts, which vary according to the amount of water available to be distributed, the day-to-day actions for distributing it made possible by whatever facilities are available, and the demands for water by the various users.

But why can we not assume that the Dutch know how to manage their own surface water system? First, even if the Dutch know how to manage their present system efficiently and effectively, it does not follow that they will know how to manage a system whose facilities have been modified (much of PAWN's effort has been devoted to analyzing changes in facilities). Second, water management goals are changing. Only during the past handful of years have Dutch water managers had to concern themselves with any water quality issue besides salt, and that in only a few locations. Now, they must take into account thermal pollution, heavy metals, BOD (biochemical oxygen demand), and other pollutants as well.

Finally, it is by no means clear that the Dutch manage their present system to best advantage. The Netherlands is a very wet country. Most of the time, almost any managerial strategy will yield satisfactory results. So, most of the time the Dutch have gotten by comfortably with a highly fragmented and uncoordinated managerial strategy, one that evolved piecemeal as major and minor constructions slowly changed the facilities. But the Dutch found the problems posed by the very dry conditions of the summer of 1976 sufficiently unsettling that they initiated the PAWN study[1]. This volume reports a method of developing more rational managerial strategies, and applies the method to both the present facilities and to a surface water system that incorporates some possible future changes.

S.2. WATER SUPPLY

The Netherlands receives its water as rainfall and river flows. In an average year, approximately 750 mm of rain falls on each square meter of the Netherlands, which is equivalent to a flow of 840 m^3/s. The rain is not markedly seasonal, but decades (ten-day periods) with no rain are not uncommon, and even whole months without rain have occurred.

We have considered six rivers that supply water to the Netherlands, and which are shown in Fig. S.1. The largest is the Rijn, which enters the country at Lobith carrying an average of 2200 m^3/s of water. Peak flows of 6000 m^3/s or higher (when averaged over a

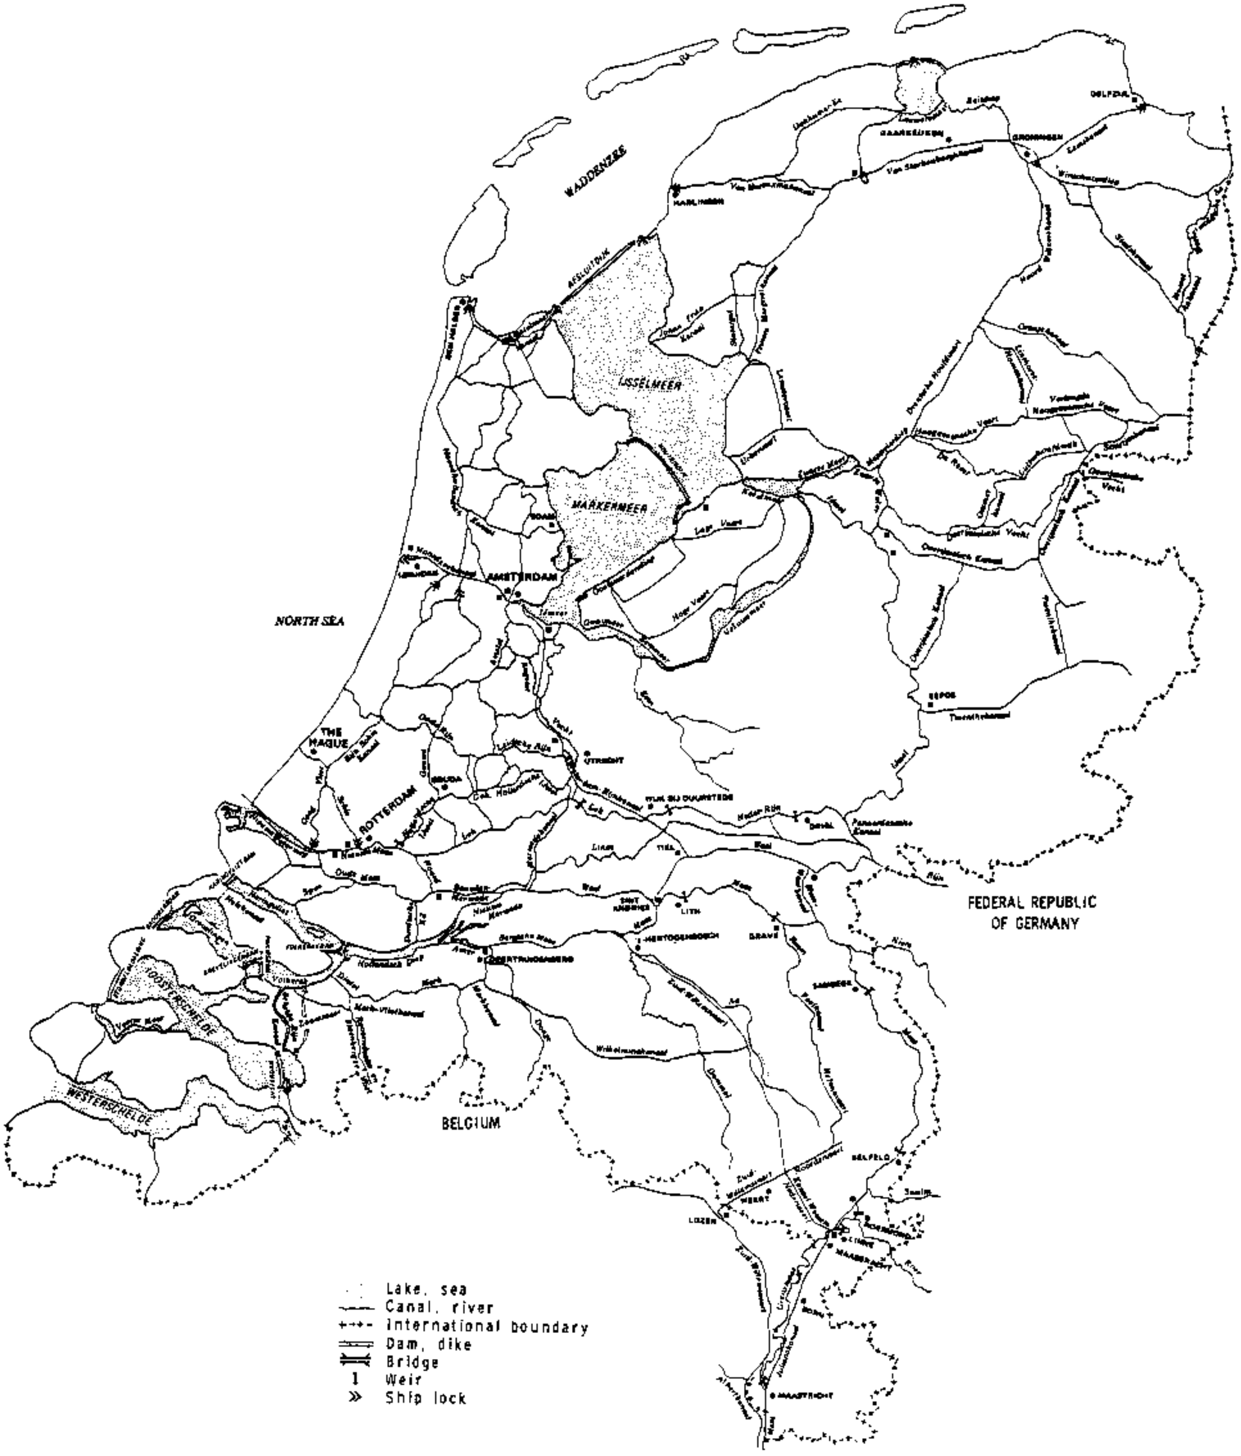

Fig. S.1--The present surface water management system decade) can occur in spring, while flows may fall below 700 m^3/s in late summer or fall.

The second largest river is the Maas, which enters the Netherlands at its southernmost extremity. The average Maas flow is several hundred cubic meters per second, but in summer and fall it may drop almost to zero. We have considered three small rivers that join the Maas as it flows north, the Roer, the Swalm, and the Niers. The Roer flow is always maintained above 10 m^3/s by releases from German reservoirs. The Swalm and the Niers have much smaller flows in summer, but carry heavy pollutant loads.

Finally, the Overijsselsche Vecht enters the Netherlands at De Haandrik on the northeastern border. Its flow in summer is only a few cubic meters per second, but it supplies an area with limited alternative supplies.

S.3. FACILITIES FOR CONTROLLING THE WATER DISTRIBUTION

The present Dutch surface water management system (see Fig. S.1) can be considered as a network of rivers, canals, lakes and reservoirs, called an _infrastructure_. Day-to-day control can be exercised over the flows in some rivers and canals by means of pumping stations and weirs. Some of the withdrawals of water from the infrastructure can also be controlled. We call the individual day-to-day control measures _managerial tactics_. Combinations of managerial tactics form a _managerial strategy_.

We distinguish managerial tactics from other, longer-term kinds of control that might be exercised over the distribution of surface water. Clearly, the digging of a new canal, the enlargement of an old one, or the construction of a new sluice or pumping station will influence water distribution. We term such tactics as these _major technical tactics_. They differ from managerial tactics in that they require a considerable time to implement, and they are more or less permanent once implemented. Managerial tactics, by contrast, can be implemented almost at a moment's notice, and as quickly canceled.

S.3.1. The Weir at Driel

From Lobith, where it enters the Netherlands, the Rijn River flows west in the Waal and the Neder-Rijn, and north in the IJssel. The division of the Rijn flow among these three river branches is determined in part by hydrological factors,[2] and in part by the weir at Driel, on the Neder-Rijn. Altering this division by adjusting the weir at Driel is one of the most important managerial tactics.

When the weirs on the Neder-Rijn were completed, in the early 1970s, it was decided to operate them so as to send a maximum amount of water north to the IJssel lakes via the IJssel River, which would also maximize the IJssel depth for the benefit of shipping. Thus, only at relatively high Rijn flows is the weir at Driel set to allow more than a minimum flow in the Neder-Rijn.

S.3.2. The Amsterdam-Rijnkanaal

At Tiel, farther west on the Waal, water can be withdrawn into the Amsterdam-Rijnkanaal and sent north to the Neder-Rijn at Wijk bij Duurstede, where it will augment the Neder-Rijn flow. Or it may be sent across the Neder-Rijn (where it can be augmented by Neder-Rijn water), and along the Amsterdam-Rijnkanaal to Amsterdam, and then west in the Noordzeekanaal into the North Sea. The adjustments of the flows on the Waal, the Neder-Rijn, and the Amsterdam-Rijnkanaal constitute another group of important managerial tactics.

The Amsterdam-Rijnkanaal was completed in 1952. It was decided to maintain the flow high enough to permit desired extractions from the canal, and to keep both the Amsterdam-Rijnkanaal and the Noordzeekanaal flushed adequately for water quality maintenance. Originally, this flow could be supplied from the Neder-Rijn, but after the weirs on the Neder-Rijn had been completed the water had to be taken from the Waal at Tiel.

S.3.3. The IJssel Lakes

The water sent north on the IJssel River flows into the only major water storage basin in the Netherlands, the IJssel lakes. There are two large lakes, the IJsselmeer and the Markermeer, plus several smaller randmeren (border lakes, such as the Gooimeer, Eemmeer, Veluwemeer, etc.).

Once the Afsluitdijk (barrier dam) was completed in 1932, and the freshwater IJssel lakes were formed, it was discovered that the IJssel flow is almost always large enough to maintain the lakes at their maximum safe levels. Thus, controlling the lake levels involved only getting rid of excess water. The preferred method extracts water from the IJmeer (at the southwest corner of the Markermeer) to augment the flow in the Noordzeekanaal; only when there is a large excess of water is any discharged through the sluices in the Afsluitdijk. So seldom have the IJssel lakes levels dropped significantly below their maxima that no doctrine exists for dealing with that circumstance; in fact, a major reason that the events of 1976 unsettled Dutch water managers was that the ad hoc responses to the shortage of water allowed the lakes to drop near their minimum levels.

S.3.4. The Maas

The Maas River is fully canalized by a system of weirs, so that even during periods of very low Maas flow (which happens frequently), the water can be maintained deep enough to permit shipping to use the river. Water can be diverted from the upper reaches of the Maas into a system of canals (the Zuidwillemsvaart, the Wessem-Nederweert, and the Wilhelminakanaal) built to make inland shipping possible. The canals also deliver water for irrigation and other purposes to areas throughout the southern provinces of the Netherlands.

Most of this area is at a high elevation (as such things go in the Netherlands), and at present can only be supplied from the upper reaches of the Maas. At times of low Maas flow, there is little freedom to choose among managerial strategies. The water manager must conserve water as best he can, mostly by reducing the losses at shipping locks. He must release water impounded behind the weirs to make up any deficits, thus reducing the depth of the Maas to the detriment of shipping. And if there is still too little water, users (mostly farmers) must go without, and shipping will cease to move.

S.3.5. Regional Waterways

In addition to the major parts of the surface water network discussed above, which have national importance, there are many small rivers and canals with only regional significance. North Holland, which is just to the west of the IJsselmeer and Markermeer, extracts its surface water from those lakes. The northern provinces to the east of the IJsselmeer also obtain their surface water from the IJssel lakes, by way of the Prinses Margrietkanaal. The provinces in the northeast highlands pump water from the IJssel River via the Twenthekanaal, and from the Zwartemeer via the Mepplerdiep. The northeast highlands also obtain some water from the Overijsselsche Vecht, that enters the Netherlands at De Haandrik. The midwest region is bounded on the south by the Lek and Nieuwe Maas rivers, on the east by the Amsterdam-Rijnkanaal, and on the north by the Noordzeekanaal. Its surface water is supplied partly from the Amsterdam-Rijnkanaal and partly from the Nieuwe Maas via the Hollandsche IJssel River. Every diversion of water from the national system into a regional waterway can be controlled, at least in part, by managerial tactics. In general, present practice supplies all the water demanded to each region, as long as the water is available.

S.3.6. Possible Future Changes

Changes to the infrastructure are likely to occur in the future, making it necessary to extend and modify the managerial strategy. For example, pumping stations have been proposed on the Zuid-Willemsvaart to permit water from the lower reaches of the Maas to supply more of the southern highlands. If a canal were also built to connect the Waal and the Maas at St. Andries, Waal water could be used in place of Maas water to supply the southern highlands.

S.4. WATER USERS AND USES, AND THEIR BENEFITS AND COSTS

We have considered the following broad categories of water users and uses:

1. Water management, including flushing and level control
2. Water Quality, Environment, and Public Health
3. Shipping
4. Agriculture
5. Power Plants (for cooling)
6. Households, Commerce, and Industry (except power plants)

S.4.1. Monetary Value of Some Water Uses

For some water users or uses, the benefit from receiving a certain amount of water can be measured in monetary terms. For example, if water is diverted from the Waal and IJssel rivers, their depths will decrease and some shippers will be unable to carry full loads. We have estimated the additional cost for transporting goods in part-loaded ships, and we infer that the value to shipping of leaving the water in the major rivers may be as high as 0.1 Dfl/m^3 when the river flows are unusually small. But to leave the water in the major rivers might mean depriving some farmers of irrigation water. We have estimated that the value of water for irrigation can range as high as 0.3 or 0.4 or even 1.0 Dfl/m^3, in a period with no rain and dry, hot conditions favoring a high rate of evaporation.

The diversion of water from the Waal could also have the purpose of providing water to cool the electric generating plants on the Amsterdam-Rijnkanaal and the Noordzeekanaal. At the present time, water temperatures in these canals are permitted to rise as much as seven degrees Celsius above their natural or background levels, but we have also considered the effect of imposing a three-degree limit on the excess temperature. Whatever the limit, if the flow of cooling water is insufficient these plants must reduce their heat discharges. This is done by operating the plants below capacity, and generating the missing power at other, less efficient plants. We have estimated the additional cost of generating the power at less efficient plants, and have inferred that, for a three-degree standard, the value of water for cooling on these canals is between 0.01 and 0.03 Dfl/m^3.

Diverting water into the Amsterdam-Rijnkanaal from the Waal or Neder-Rijn also reduces the flow of water in the Nieuwe Maas. Water from the North Sea, being denser than the fresh Rijn water, flows upstream in the Nieuwe Waterweg and Nieuwe Maas in a layer at the bottom of the channel. When the flow in the Nieuwe Maas is too low, this salt wedge can penetrate as far inland as the mouth of the Hollansche IJssel. In dry periods, water flows north in the Hollansche IJssel from its mouth to Gouda, where it is extracted to supply most of the water needs of the midwest. The salt admitted along with the water can cause significant damage to the crops grown in glasshouses in the midwest, in large part because the water is not soon flushed out of the region. Thus, diverting water from the Waal or Neder-Rijn into the Amsterdam-Rijnkanaal indirectly damages glasshouse crops in the midwest. We have estimated the damage caused by the diversion of one cubic meter of water to be as high as 0.35 Dfl, when the Rijn flow is extremely low.

S.4.2. Water Uses With Nonmonetizable Values

For other uses of water, it is impossible to estimate a monetary value. For example, it is considered essential by Dutch water managers that the water in regional canals and ditches in the lowlands be maintained at constant levels. To allow the levels to rise would cause flooding, while to allow them to drop would expose uncounted industrial water intakes. Letting the water levels drop would also reduce the hydraulic pressure that limits rate of seepage of sea water into these sub-sea level areas. (Actually, these levels do vary by as much as ten centimeters, probably because they are difficult to control more closely. Ten centimeters sounds small, but the total area of water involved is so large that a 10-cm change in level accommodates a large store of water.) Two other effects of a drop in water level are (i) the drying out and subsequent subsidence of the land, and (ii) the exposure to air of usually submerged wooden foundation piles. The first of these effects would make permanent an increase in the seepage rate of sea water, since the water level could not thereafter be raised as high as previously without flooding. The second effect would accelerate the decay of the foundation piles.

Another use of water for which we cannot estimate a monetary value is water quality improvement. How much is it worth to lower the concentration of phosphate or BOD (biochemical oxygen demand) by 1 mg/l? This might reduce the amount of algae growing in the water, improving its appearance. It might change somewhat the species of fish and water plants found there. There could even be some public health benefits. But we know of no method for turning these benefits into money equivalents.

In this volume we consider five measures of water quality. They are phosphate, which is a nutrient contributing to algae growth (heavy growths of algae are a major nuisance in many Dutch lakes); BOD, which can deplete water of oxygen, causing fish kills and noxious odors; chromium, which is one of a class of toxic substances; salt, which degrades the taste of drinking water and may cause high blood pressure; and thermal pollution, which may drive some species out of their traditional habitats. The damage done to agriculture by salt can be monetized, and we have done so, but the taste and public health aspects cannot. Similarly, the cost of reducing heat discharges from power plants can be estimated, but not the "cost" of living with a high increment of temperature above the background (the excess temperature). For the other pollutants, we have considered no monetizable aspects.

We have treated nonmonetizable water uses by means of constraints. Instead of estimating the damage caused by a change in the water level in a ditch from some preferred value, we simply impose the requirement that the ditch water level remain constant, and we always supply (or extract) enough water to keep it so. Similarly, instead of trying to value a reduction in pollutant concentrations, we impose a water quality standard that prohibits the concentration from rising above a specified value. The values chosen for the water quality standards and the preferred ditch levels are based on prevailing Dutch opinion and practices.

S.5. THE MANAGERIAL STRATEGY DESIGN MODEL

We constructed a Managerial Strategy Design Model (MSDM) to solve the problem stated at the beginning of this summary. The model is formulated as a mathematical programming problem, which is a class of problems having variables whose values are restricted by constraints, and whose purpose is to find, among all variable values that satisfy the constraints, those values that minimize an objective function.

In MSDM the variables are flows in rivers and canals and water levels in lakes and reservoirs, as well as concentrations of pollutants at numerous locations. Some constraints on flows and water levels define the managerial tactics, while others impose requirements for water uses whose values we could not monetize. Many constraints can be changed (usually relaxed) by changing the infrastructure, e.g., building new canals, pumping stations, or weirs. Such changes have the effect of increasing or decreasing the range of managerial tactics available to the water manager. Water quality standards are expressed as contraints on pollutant concentrations. The objective function is the sum of all monetized costs imposed by any set of flows, water levels, and pollutant concentrations.

The model can be decomposed naturally into a subproblem dealing only with the water distribution (i.e., flows and water levels), and five water quality subproblems, each dealing with a single pollutant. The model finds a managerial strategy for any specified water supply and infrastructure by cycling through the six subproblems repeatedly. In the first cycle, it solves the water distribution subproblem, considering only the user costs that are expressed directly in terms of water flows and levels. That is, it does consider pumping costs, costs of low water to shipping, and costs to agriculture of failing to supply water for irrigation; but it does not consider the effect of changing the water distribution on the cost of meeting the demand for electric power under a thermal limitation, or the cost to agriculture of excessive salt concentrations, even though a different water distribution might reduce those concentrations. Continuing the first cycle, it uses the flows and levels thus calculated as inputs to each of the water quality subproblems in turn, and computes the pollutant concentrations.

In the next cycle, MSDM once again solves the water distribution subproblem. This time, however, it uses the solutions of the water quality subproblems to estimate the effect of changing the water distribution on quality-related costs, and adds these estimates to the user costs already expressed directly in terms of water flows and levels. It is not possible to calculate the exact effect, so if the solution to the new water distribution subproblem differs from the previous solution, it is necessary to solve each water quality subproblem again, form another estimate of the effect, and start another cycle.

MSDM does not invariably converge upon a single water distribution, even after arbitrarily many cycles. It may oscillate, or it may drift seemingly at random. When this happens, we can usually obtain convergence by manually imposing some artificial constraints, thus reducing the freedom available to MSDM. Because of this, we cannot be certain that the managerial strategy selected actually minimizes the objective function (i.e., is optimal). But we have experimented enough with MSDM to assure ourselves that the managerial strategies selected are nearly optimal.

S.6. RESULTS

S.6.1. "Dilution Is No Solution To Pollution"

Our first experiments with MSDM demonstrated that, with a few exceptions, managerial tactics have little effect on water quality. In fact, for most pollutants managerial tactics cannot reduce pollutant concentration enough to meet the water quality standards, and MSDM found its problem to be infeasible. Thus, for later MSDM cases we relaxed the water quality standards for phosphate, BOD, chromium, and salt. We retained the thermal standards, at least for some cases, because we had included in MSDM the means to shift power generation from one plant to another in order to meet them. We also retained the costs associated with high salt concentrations.

We must include a cautionary footnote. While MSDM represents the nationally important part of the surface water network in reasonable detail, it represents the regional waterways in a highly aggregated way. The pollutant concentrations it calculates in regional waters, therefore, are averages over relatively large areas. It is quite possible that a more detailed examination of regional waterways would uncover instances where the diversion of water from one small canal to another might dilute a highly polluted, but rather small area.

We mention above that a few instances exist of managerial tactics affecting water quality to a significant degree. One such instance is the salt concentration in the Hollandsche IJssel, which is elevated by salt intrusion from the North Sea when the Nieuwe Maas flow is low. The Nieuwe Maas flow can be increased by opening the weir at Driel or by withdrawing water from the Waal at Tiel to augment the Neder-Rijn flow. Even though the salt standards are ignored, these tactics will be employed under some circumstances because salt in the Hollandsche IJssel will be taken into the midwest and cause damage to crops under glass. This damage is included in the MSDM objective function.

Another instance is the use of water from the IJssel lakes cool the discharges from power plants on the Noordzeekanaal and on the Bergumermeer (a small lake northeast of the IJssel lakes). These tactics have been employed in all of the MSDM cases in which a thermal standard has been imposed.

S.6.2. The Value of Storing Water for Future Uses

Using MSDM we have developed superior managerial strategies for decades with widely varying water supply conditions. As expected, the costs to all users rise as the water supply falls, whether the water be rain or river water, or water stored up in the IJssel lakes from previous decades. But if having water in storage at the start of a decade may reduce the costs to users during the decade, then there must have been some value to putting water in storage in previous decades. And by shifting one's time perspective, one can see that there must be some value to putting water in storage during the present decade for possible use in later decades.

Using cost differentials from MSDM cases, and probability estimates from rain and river flow data, and from Distribution Model simulations, we have estimated the average value of storing water in the IJssel lakes to be between 0.0011 and 0.0045 Dfl/m^3. This value is smaller that the value of using water to cool the power plants on the Noordzeekanaal and the Bergumermeer to meet a three-degree thermal standard (if so strict a standard is imposed). Thus, it would seem preferable to use the water for cooling the power plants rather than saving it for possible future needs.

We caution, however, that the average value of stored water hides a large variation. Most of the time, the stored water is not needed at all, and has zero value. But on rare occasions, when dry conditions persist long enough to draw down the IJssel lakes to their minimum tolerable level, water must be diverted from all possible users to avoid drawing down the lakes further. On these occasions, the value of having more water in the lakes is as large the highest value of water found among the users who have been deprived. Many people might therefore prefer to store up water even when average value calculations indicate it is not economical to do so, in order to avoid occasional, but catastrophic, losses.

S.6.3. A Superior Managerial Strategy

S.6.3.1. The MSDM Strategy. Different infrastructures, (e.g., new canals or pumping stations) and different water supply conditions demand managerial strategies that are different in detail. However, regardless of the water supply, we can describe all the strategies generally in terms of the same priority list. The priority ordering of water uses corresponds roughly to the relative economic values that water has in the various uses, but cannot precisely reflect the different values of water because some of the values vary strongly with the amount of water devoted to the corresponding use. The order, from high priority to low, is the following:

Priority 1: Supply level control requirements for boezems and lakes, and meet all other constraints on MSDM.

Priority 2: Supply water to farmers for irrigating their crops. Also, establish certain nominal flushing rates for locks at which salt intrusion causes damage to crops grown locally.

Priority 3: Trade off shipping losses due to low water on the Waal and the IJssel, and salt damage to agriculture due to the Rotterdam salt wedge, by simultaneously adjusting the Neder-Rijn flow (by adjusting the weir at Driel) and withdrawals at Tiel.

Priority 4: Use water from the IJssel lakes for cooling the power plants on the Noordzeekanaal and the Bergumermeer to whatever thermal standard has been set for them.

Priority 5: Raise the IJssel lakes to their maximum levels to meet the possible future needs for water.

Priority 6: Use water for flushing boezems and ditches, and for raising the flushing rates of locks with local salt intrusion above the nominal rates established in Priority 2.

Frequently, a particular cubic meter of water is in a location to from which it can be used to satisfy only some demands. Under such circumstances, the priority scheme requires that each cubic meter of water be used for the highest priority use for which it can be used. Thus, water stored in the IJssel lakes might be used to cool the power plants on the Noordzeekanaal or the Bergumermeer (priority 4), since no means exist to transport it to the southern highlands to meet demands for irrigation water (Priority 2). Note, however, that if the infrastructure were changed to make IJssel lakes water accessible to the southern highlands (a highly uneconomical but nevertheless technically feasible undertaking), the water would then be diverted to the higher priority use (irrigation).

S.6.3.2. Comparison of Strategies. The major difference between the MSDM strategy described above and present Dutch practice is concerned with the trade-off called for in Priority 3. For Rijn flows below average, present practice closes the weir at Driel almost completely, leaving only a minimum flow on the Neder-Rijn. This maximizes the depth of the IJssel, and hence benefits IJssel shipping. However, it minimizes the total flow to the west (Neder-Rijn plus Waal), which allows a maximum of salt damage to agriculture due to the Rotterdam salt wedge. Finally, by minimizing the Neder-Rijn flow it requires a maximum withdrawal from the Waal at Tiel, which is unfavorable to Waal shipping. The trade-off performed in the MSDM strategy can yield large savings over the present practice during decades with very low Rijn flows.

Another difference between the two strategies involves the flow in the Amsterdam-Rijnkanaal. Present practice supplies a considerable amount of water to cool the Noordzeekanaal power plants by withdrawing at Tiel. According to Priority 4, the MSDM strategy provides the water from the IJssel lakes instead. This reduces the cost to Waal shipping, and leaves more water in the Waal to help combat the salt wedge, but it does tend to draw down the lakes to lower levels.

Finally, present practice carries out some flushing of boezems and ditches using water from the IJssel lakes, even when the lakes are not at their maximum levels. This is a reversal of Priorities 5 and 6, and is not done by the MSDM strategy.

We have succeeded in implementing the trade-off of Priority 3, and the change in the source of Noordzeekanaal cooling water from the Waal to the IJssel lakes in the Distribution Model, and have compared simulations using this partial MSDM strategy with simulations using strategies that approximate present Dutch practice. All the strategies are designed to cool the Noordzeekanaal power plants only to meet a seven-degree standard, which is the present requirement. Even though only part of the MSDM strategy was implemented, these comparisons show a savings of 0.5 Dflm[3] to shipping and agriculture in a year with an average water supply from rain and rivers, and up to 10 Dflm savings in an extremely dry year.

The reason that the estimated average savings are so small, and that the extreme savings are so much larger, is that most of the time there is ample water available to provide for all users, of all priority classes. The salt wedge rarely intrudes as far as the mouth of the Hollandsche IJssel. The Rijn flow is usually large enough to minimize shipping costs. The lakes are ordinarily at their maximum levels. In fact, there is usually a need to discharge water from the lakes to prevent them from rising above their maximum levels. This excess water might as well be used to cool power plants or flush the the boezems of North Holland. Thus, the model estimates the managerial strategy discussed here to be an improvement over present practice only a small fraction of the time.[4]

NOTES

1. Our Dutch reviewers inform us that preparations for the PAWN study had actually started prior to 1976, but that the dry conditions of 1976 did shape the study significantly.
2. According to Webster's New Collegiate Dictionary, 1975, hydraulics is "a branch of science that deals with practical applications (as the transmission of energy or the effects of flow) of liquid (as water) in motion," whereas hydrology is "a science dealing with the properties, distribution, and circulation of water on the surface of the land, in the soil and underlying rocks, and in the atmosphere." The main distinction seems to be that the term "hydraulics" and related terms (e.g., hydraulic, hydraulically) refer to man-caused effects on water, while terms such as hydrology, hydrological, and the like, refer to natural effects. The hydrological factors meant here are the shape, resistance, and gradient of the river channels. Although man has certainly influenced these factors, for example by dredging, they were largely shaped by nature. By contrast, the weir at Driel, being a man-made object intended to influence the flow of water, would be referred to here as a hydraulic factor.
3. Dflm is our abbreviation for "millions of Dutch florins." At the time of the study, a Dutch florin, or guilder, was valued at between U.S. $0.35 and $0.40.
4. We also caution the reader that, as is usual in studies of this nature, and especially in studies as broadly scoped as PAWN, estimates of costs and benefits are likely to be uncertain. The monetary differences between strategies that we quote here rely on numerous assumptions and approximations made throughout the study. In addition, PAWN was a policy study, and thus necessarily ignored many factors that could be important in implementation questions. Accordingly, one should be cautious in applying the results of this study. In the present instance, the fact that MSDM has found a "better" managerial strategy does not mean that the Dutch should rush to implement it. But it does suggest that it may be worthwhile to investigate the question further, to determine the effect of the approximations and assumptions used in PAWN, or the factors excluded from consideration, would have on this conclusion.

ACKNOWLEDGMENTS

Although this volume and its companion Vol. VA have only one author, it should be recognized that many people have contributed to the work reported herein. Some of these contributions are documented in other PAWN volumes, and can be referred to there. But other contributions appear only here.

E. van Beek of the Delft Hydraulics Laboratory formulated and tested an aggregated and simplified representation of agriculture for MSDM. J. W. Pulles of the Rijkswaterstaat scrutinized early outputs of the model, and suggested corrections and improvements.

R. L. Petruschell of Rand assisted in analyzing the results of MSDM runs, and suggested many lines of inquiry. C. Dzitzer searched and synopsized the literature on water quality standards and the effects of water pollution; she also prepared the pollutant discharge tables.

B. F. Goeller, Project Leader of PAWN at Rand, provided broad guidance throughout the project, as well as very helpful administrative support.

Many of our Dutch colleagues reviewed these pages, making numerous helpful comments. Among those serving in this capacity were E. van Beek and G. Baarse of the Delft Hydraulics Laboratory, and J. W. Pulles of the Rijkswaterstaat.

L. W. Miller, a Rand consultant, made a variety of constructive suggestions, whose incorporation has much improved the text.

Special thanks must go to E. T. Gernert, Managing Editor of PAWN, and her colleagues in the Publications Department for their care in guiding this volume to press; to N. Katz for preparing various auxilliary computer programs to aid in preparing inputs and displaying outputs of MSDM; and to M. P. Dobson amd M. LaPrell for painstakingly incorporating changes in the draft.

CONTENTS

FIGURES

TABLES

Chapter 1

INTRODUCTION

1.1. THE MANAGERIAL STRATEGY DESIGN PROBLEM

This volume investigates what day-to-day actions will bring about the distribution of surface water most beneficial to all water users and uses.

1.1.1. Water Users and Uses, and Their Benefits

In this investigation, we have considered the following broad categories of water users or uses:

1. Water management, including flushing and level control
2. Water quality, environment, and public health
3. Shipping
4. Agriculture
5. Power plants (for cooling)
6. Households, commerce, and industry (except power plants)

For some water users or uses, the benefit from receiving a certain amount of water can be measured in monetary terms. This is true for shipping, where changes in water levels affect costs; for agriculture, where a lack of water will reduce crop yields; and for power plants, where a lack of water for cooling will require that heat discharges, and hence power generation, be reduced at critical locations. A reduction in power generation at one plant requires an increase at another, less efficient plant located elsewhere, thus raising the cost of meeting the demand for electric power.

For other users and uses, a monetary value of water is impossible to estimate. Water management uses, such as flushing to reduce local salinity, and the control of water levels in ditches and canals, have no direct economic value, although the indirect effects could hardly be more widespread. A vast and controversial literature exists on ways to estimate the monetary value of maintaining good water quality, or preserving the environment, or protecting the public health. For these uses and users, we have introduced constraints on the day-to-day actions taken. These can be limits on flows or water levels, or water quality standards.

1.1.2. Day-to-Day Actions Affecting the Water Distribution

The Dutch surface water management system is a network of rivers, canals, lakes and reservoirs, called an infrastructure. Day-to-day control can be exercised over the movement and storage of water in the infrastructure by means of pumping stations, weirs, discharge sluices, etc. Control can also be exerted over the withdrawal of water from the infrastructure. We call individual day-to-day control measures managerial tactics.[1] Combinations of managerial tactics form a managerial strategy.

We distinguish managerial tactics from other, longer-term kinds of control that might be exercised over the movement and uses of surface water. Clearly, the digging of a new canal, the enlargement of an old one, or the construction of a new sluice or pumping station will influence water distribution. We term such tactics as these major technical tactics. They differ from managerial tactics in that they require a considerable time to implement, and they are more or less permanent once implemented. Managerial tactics, by contrast, are those which can be implemented almost at a moment's notice, and which can be as quickly canceled.

1.1.3. The Need for Managerial Strategy Design

The question arises, why should any effort be spent designing new managerial strategies for the Netherlands? Can we not reasonably assume that the Dutch know how to manage their own surface water system? There are several replies.

First, even if the Dutch know how to manage their present system efficiently and effectively, it does not follow that they will know how to manage a system whose infrastructure has been modified. Much of PAWN's effort has been devoted to analyzing changes in infrastructure, so this is no idle consideration. Second, the goals for which the system is to be managed are changing. Only during the past handful of years have Dutch water managers had to concern themselves with any water quality issue besides salt, and that in only a few locations. Now, they must take into account thermal pollution, heavy metals, BOD, and other pollutants as well.

Finally, it is by no means clear that the Dutch manage their present system to best advantage. The Netherlands is a very wet country. Most of the time, almost any managerial strategy will yield satisfactory results. So, most of the time the Dutch have gotten by comfortably with a relatively uncoordinated managerial strategy. But the Dutch found the problems posed by the very dry conditions of the summer of 1976 sufficiently unsettling that they initiated the PAWN study.[2]

1.2. APPROACH TO THE PROBLEM

We have followed two parallel paths to investigate new managerial strategies. The more straightforward path involved the use of the Distribution Model (see Vol. XI). Simulations covering several years were carried out with each infrastructure of interest. If the simulation for any year gave rise to a catastrophic cost to any user of water (usually agriculture), the strategy was adjusted to avoid that event. (The usual cause of a catastrophic cost was that the water level of the IJsselmeer--the largest lake in the Netherlands, and the most important for water storage--dropped to its minimum allowed value, forcing deliveries of water for sprinkling in the North and North Holland to be curtailed.)

The other path we followed toward managerial strategy design aimed at understanding the trade-offs among the different users of water when no user's costs were catastrophically large. It aimed at designing strategies that would improve on those that merely avoided instances of very large costs. It aimed, in a word, at "superior" managerial strategies. Our vehicle for traveling the second path was the Management Strategy Design Model (MSDM), which is described in this volume.

1.2.1. Formulation of the Problem

We formulated the problem of designing managerial strategies as a nonlinear program. Problems of this form have three different elements, called <u>variables</u>, <u>constraints</u>, and an <u>objective function</u>. A <u>feasible solution</u> to the problem is a set of values for the variables that satisfies all the constraints. An <u>optimal solution</u> is a feasible solution that yields the smallest possible value of the objective function. Recall the opening sentence of this chapter, that defined the problem of managerial strategy design as <u>finding what day-to-day actions would bring about the distribution of surface water most beneficial to all water users and uses</u>. We can identify the variables of the problem as describing the water distribution, the constraints as determining what day-to-day actions are possible, and the objective function as defining how beneficial a particular water distribution is to water users and uses.

1.2.2. Organization into Subproblems

We organized the managerial strategy design problem as a series of subproblems. The most central subproblem, the <u>water distribution subproblem</u>, concerns water quantity variables such as flows and water levels in the infrastructure. These are the quantities over which water managers have the greatest and most direct control. The other subproblems are concerned with water quality, with one subproblem devoted to each pollutant. The pollutants considered in the present version of MSDM are salt (expressed as chloride), thermal pollution, heavy metal (we take chromium as our example), BOD (biochemical oxygen demand), and phosphate.

The water distribution subproblem is formulated as a linear program, the solution to which is the water distribution, expressed in terms of water flows in rivers and canals, and water levels in reservoirs and lakes. Each water quality subproblem is expressed as a linear program whose coefficients depend on the flows and levels that describe the water distribution. The results of solving the quality subproblems influence the quantity subproblem as follows. An estimate is made from the results of solving each quality subproblem of how sensitive the quality-related costs are to changes in the water flows and levels. Then these sensitivities are used to modify the quantity-related cost components that form the objective function of the water distribution subproblem.

1.3. SOME CONCLUSIONS AND RECOMMENDATIONS

Our first conclusion is that managerial tactics generally have little effect on water quality. A great deal of water can be made to flow into the North Sea instead of into the IJsselmeer, or along the Neder-Rijn River instead of the Waal or IJssel rivers, without much affecting pollutant concentrations. But such diversions of water can cause great harm to shipping, agriculture, and other users. We find, therefore, that improving water quality should not be given much weight in the choice of a managerial strategy.

We hasten to add that a concern for water quality may well influence the choice of longer-term tactics, either within or related to the field of water management. The same concern may also affect what industrial or agricultural developments are allowed in the future. We also must add that special, local circumstances can justify exceptions to this rule. For example, it may be justified to maintain a flow of 10 m^3/s in the city canals of Amsterdam, to prevent them from becoming open sewers. But in general, day-to-day managerial decisions--especially those of larger scope--should not much depend on water quality.

Our second conclusion is that, although in most instances the present Dutch managerial practice is satisfactory, there are some circumstances under which high costs can result. Implementing a more coordinated managerial strategy would result in only small savings in an average year (under one million guilders), but it could well save ten million guilders or more in a year with little rain and low river flows.[3] The year 1976 was such a year.

Third, some of the many assumptions regarding water management should be questioned. On advice from the Dutch, PAWN has considered it to be absolutely necessary to maintain the water levels in most ditches and canals at constant levels. At times, the MSDM strategy will divert water (e.g., from irrigation) to maintain these levels, at considerable cost to agriculture and other users. If the water levels could be allowed to decline by even 10 cm, these costs could be at least delayed, and very likely avoided. (Of course, the fact that an assumption is questioned does not imply that it will necessarily be changed. Upon investigation, one might discover that the costs of allowing the water level to drop as suggested here exceed the benefits. But we are not aware of any presently existing estimate of such costs.)

1.4. ORGANIZATION OF THIS VOLUME

The remainder of this volume discusses the formulation of the managerial strategy design model (MSDM) in Chaps. 2 through 7, the computational features and mechanics of using the model in Chap. 8, and some results and conclusions obtained from exercising MSDM in Chaps. 9 through 11.

1.4.1. The Formulation Chapters

Chapter 2 discusses the water distribution subproblem. It describes the water management infrastructure and the managerial tactics that the infrastructure makes possible. Some of these managerial tactics have associated costs, such as energy costs for pumping. These costs are included in the objective function of the water distribution subproblem. The chapter also discusses the sources of water for the Netherlands, and the relation between surface water (which is our direct concern in MSDM) and groundwater.

Chapter 3 formulates the water quality subproblems for the pollutants heavy metals, BOD, and phosphate. It discusses the sources of these pollutants, and what standards should be applied to pollutant concentrations. It also discusses the process of calibrating the model to observed pollutant concentrations.

Chapter 4 deals with salt. It formulates the salt subproblem, and discusses sources, means of control, and water quality standards for salt. Unlike the pollutants discussed in Chap. 3, high salt concentrations can result in monetary costs. Some of these costs are included in the objective function of the salt subproblem, while others appear as terms in the objective function of the water distribution subproblem.

Chapter 5 formulates the thermal subproblem. The major source of heat we consider is waste heat discharged from power plants. In this chapter, we discuss the relation between waste heat discharged and power generated. Standards limit heat discharges at various points in the infrastructure, and demands for power require that some waste heat be produced, and discharged somewhere. The objective function of the thermal subproblem consists of the marginal cost of generating and transmitting enough power to meet regional demands. The variables are the amounts of power generated by each of the power plants.

Chapter 6 estimates the value of water to agriculture. There are many different crops grown on different soils with differing degrees of access to water. Therefore, there are many different values that water may have to different examples of agriculture. These different values enter the water distribution subproblem as part of the objective function.

Chapter 7 estimates the value of water to shipping. Restricted water depths in the major rivers result in shipping costs. Major river depths can become restricted either due to low water flows or to the deposition of sediment, which may in turn be caused by unwise withdrawals of water from the rivers. Another source of costs to shipping is delays at locks, which may be lengthened if certain water conservation measures are employed. These values are incorporated into the objective function of the water distribution subproblem.

1.4.2. The Use of MSDM

Chapter 8 assembles the six subproblems formulated in the earlier chapters, and discusses how they depend upon one another. Although it is easy to solve any subproblem in isolation, the interactions between them makes it horribly difficult to obtain an overall solution. We have not devised a completely satisfactory way through the difficulties, but we have developed a method which, with judicious human intervention, can produce good solutions.

1.4.3. Results and Conclusions

Chapter 9 presents results of using MSDM to investigate managerial strategies for the present infrastructure and present patterns of water demands (especially by agriculture). We compare the optimal strategy with present practice, and investigate managerial strategies to improve water quality. We also develop a description of the strategy found by MSDM in the form of a list of priorities for water users and uses. Using this list, one can approximate the MSDM strategy under a wide variety of conditions (river flows, rainfall, etc.) without relying on calculations by the rather cumbersome MSDM model.

Chapter 10 repeats most of Chap. 9, but for an infrastructure to which several major technical tactics have been added, and for patterns of water demands that include a four-fold increase in the agricultural use of irrigation water. We find that the priority list of water users developed in Chap. 9 continues to serve as a good description even for the changed infrastructure and water demands.

Chapter 11 presents our main conclusions and recommendations.

1.4.4. The Technical Appendixes

Volume VA of this series is a companion to this volume. Looking at the size of this volume, the reader might suspect that no detail, however small, could have been omitted. In fact, frequent references in this volume will show that many details were omitted, only to appear in one of seven technical appendixes contained in Vol. VA.

NOTES

1. There exists at the present time a doctrine specifying how the infrastructure should be managed from day to day. We have described this body of doctrine as best we can as a set of managerial rules, each covering a particular action--e.g., setting a pumping rate or a weir position. The rules have been incorporated in the Distribution Model (see Vol. XI). Technically, what we mean by a managerial tactic is a change of a managerial rule.
2. Our Dutch reviewers inform us that preparations for the PAWN study had actually started prior to 1976, but that the dry conditions of 1976 did shape the study significantly.
3. We also caution the reader that, as is usual in studies of this nature, and especially in studies as broadly scoped as PAWN, estimates of costs and benefits are likely to be uncertain. The monetary differences between strategies that we quote here rely on numerous assumptions and approximations made throughout the study. In addition, PAWN was a policy study, and thus necessarily ignored many factors that could be important in implementation questions. Accordingly, one should be cautious in applying the results of this study. In the present instance, the fact that MSDM has found a "better" managerial strategy does not mean that the Dutch should rush to implement it. But it does suggest that it may be worthwhile to investigate the question further, to determine the effect of the approximations and assumptions used in PAWN, or the factors excluded from consideration, would have on this conclusion.

Chapter 2

THE WATER MANAGEMENT INFRASTRUCTURE

2.1. NODES AND LINKS

In both the Distribution Model (DM) (see Vol. XI) and MSDM, the water management infrastructure is expressed as a network consisting of nodes and links. Links represent waterways, such as sections of rivers and canals. We have included a link in each network wherever we wish to allow water to flow. Links are also used to represent extractions from the surface water system that lead to points outside, and discharges into the system that originate outside. An example of an extraction is the taking of water from a lake to irrigate cultivated land. An example of a discharge is water in a river (e.g., the Rijn) entering the country.

Nodes represent junctions of two or more links (waterways, extractions, discharges, or any combinations), or locations where water can be stored. Thus, if a river divides into two branches, the branch point is represented by a node. The same is true wherever two rivers or two canals join, or where an extraction or a discharge occurs. In addition, each lake (e.g., the IJsselmeer) is represented by a node. In our network representation of the infrastructure, a link may connect two nodes (a waterway), or it may lead into a node from outside the system (a discharge), or from a node to points outside (an extraction).

We have used two network representations of the infrastructure in PAWN, one for the distribution model, and one for MSDM. The reason is that computationally, MSDM is much more complex, and must operate with a simpler network if it is to be economically feasible to run. There is, however, a definite relation between the two networks, which is described in App. A of Vol. VA.

Figure 2.1 below shows the network used in MSDM. The reader can find the names of the nodes and links in Tables 2.1 and 2.2, respectively. Both nodes and links are numbered in the figure in the same order as in the tables. Links in Fig. 2.1 are designated as either major waterways (solid lines) or minor waterways (dashed lines), on the basis of the amount of water typically carried by the different links. In Table 2.2, links shown with no "To Node" are major extractions, while those with no "From Node" are discharges. Not shown in either Fig. 2.1 or Table 2.2 are links that represent minor extractions for irrigation, industry, or leakage from surface water into the groundwater. Such minor extractions occur at every node.

In both DM and MSDM, each link is represented by a variable, the value of which is calculated for each decade considered in a run of the model. The variable denotes the flow of water in cubic meters per second (m^3/s) through the associated link. In MSDM, these

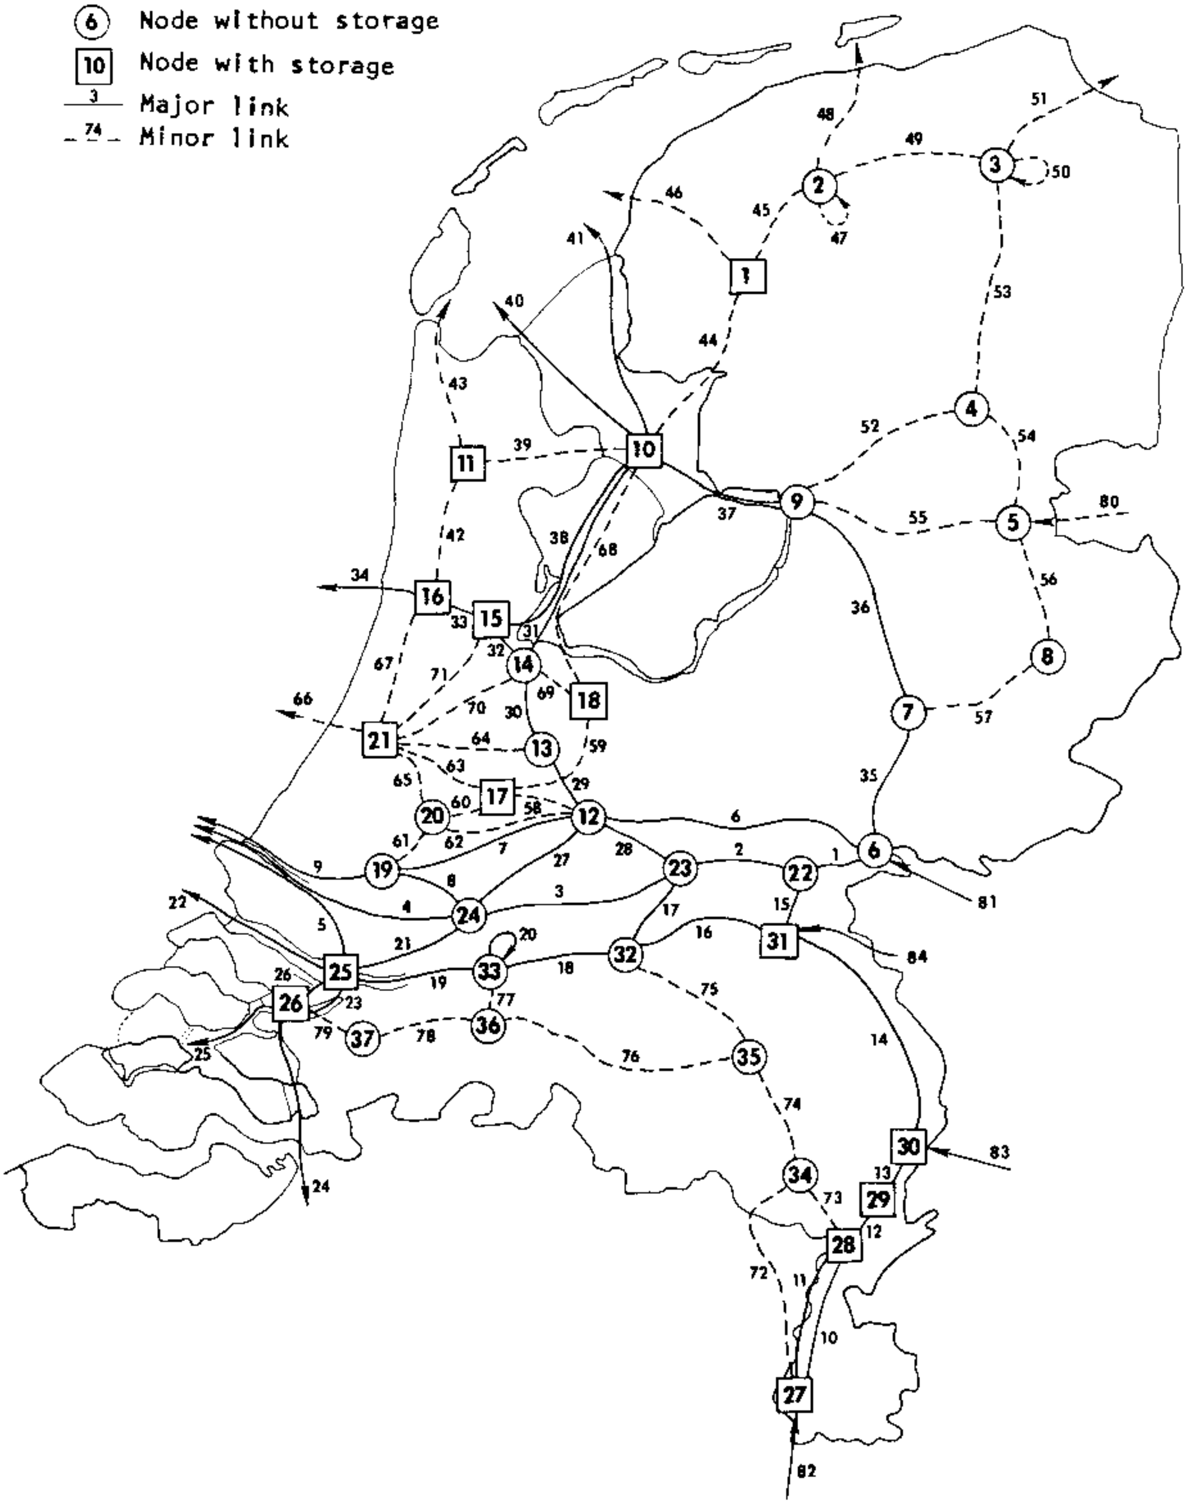

Fig. 2.1--The MSDM network

Table 2.1

NODES IN THE MSDM NETWORK

1. FRIELAND(a)
2. B'GUMMER
3. GRONETAL
4. NEHIGH
5. OVIJVECH
6. UPRIVER
7. TWENMOND
8. TWENTEND
9. IJSLEND
10. IJSLAKES(a)
11. NORHOLL(a)
12. A-R.MOND
13. UTR+MAAR
14. DIEMEN
15. AMSTEDAM(a)
16. HAL+IJMU(a)
17. LOPIKWAR(a)
18. VECHT(a)
19. IJSLMOND
20. GOUDA
21. MIDWEST(a)
22. NIJMEGEN
23. TIEL
24. GOR+DOR
25. LOWRIVER(a)
26. DLTALAKE(a)
27. MAASLOO(a)
28. BORN+PAN(a)
29. LINNE(a)
30. ROER+BEL(a)
31. SAMGRALI(a)
32. DBOS+BOX
33. GERTRUID
34. WEER+MEY
35. HELMOND
36. OSTRHOUT
37. FIJNAART

(a) Water storage associated with this node.

Table 2.2

LINKS IN THE MSDM NETWORK

	Link Name		From Node		To Node
1.	WAAL1	6.	UPRIVER	22.	NIJMEGEN
2.	WAAL2	22.	NIJMEGEN	23.	TIEL
3.	WAAL3	23.	TIEL	24.	GOR+DOR
4.	OUDEMAS1	24.	GOR+DOR		
5.	KIL+SPUI	25.	LOWRIVER		
6.	NEDRIJN2	6.	UPRIVER	12.	A-R.MOND
7.	LEK3	12.	A-R.MOND	19.	IJSLMOND
8.	NOORD	24.	GOR+DOR	19.	IJSLMOND
9.	NIEWMAAS	19.	IJSLMOND		
10.	JULCANL1	27.	MAASLOO	28.	BORN+PAN
11.	MAAS1	27.	MAASLOO	28.	BORN+PAN
12.	MAAS2	28.	BORN+PAN	29.	LINNE
13.	MAAS3	29.	LINNE	30.	ROER+BEL
14.	MAAS5	30.	ROER+BEL	31.	SAMGRALI
15.	MAWAKAN	22.	NIJMEGEN	31.	SAMGRALI
16.	MAAS8	31.	SAMGRALI	32.	DBOS+BOX
17.	SANDRIES	23.	TIEL	32.	DBOS+BOX(a)
18.	BERGMAAS	32.	DBOS+BOX	33.	GERTRUID
19.	AMER1	33.	GERTRUID	25.	LOWRIVER
20.	AMER REC	33.	GERTRUID	33.	GERTRUID(b)
21.	NIEWMERW	24.	GOR+DOR	25.	LOWRIVER
22.	HARINGSL	25.	LOWRIVER		
23.	VOLKERAK	25.	LOWRIVER	26.	DLTALAKE
24.	KREEKRSL	26.	DLTALAKE		
25.	PHILDAM	26.	DLTALAKE		
26.	HALSKAN	25.	LOWRIVER	26.	DLTALAKE(a)
27.	MERWKAN1	24.	GOR+DOR	12.	A-R.MOND(a)
28.	ARKANAL1	23.	TIEL	12.	A-R.MOND
29.	ARKANAL3	12.	A-R.MOND	13.	UTR+MAAR
30.	ARKANL4B	13.	UTR+MAAR	14.	DIEMEN
31.	IJSYPHON	14.	DIEMEN	10.	IJSLAKES(a)
32.	ARKANAL5	14.	DIEMEN	15.	AMSTEDAM
33.	NZKANAL1	15.	AMSTEDAM	16.	HAL+IJMU
34.	NZKANLSL	16.	HAL+IJMU		
35.	IJSSEL1	6.	UPRIVER	7.	TWENMOND
36.	IJSSEL4	7.	TWENMOND	9.	IJSLEND
37.	KETLIJSL	9.	IJSLEND	10.	IJSLAKES
38.	ORANJESL	10.	IJSLAKES	15.	AMSTEDAM
39.	SCHERMIN	10.	IJSLAKES	11.	NORHOLL
40.	D.OEVER	10.	IJSLAKES		
41.	KORNWEDR	10.	IJSLAKES		
42.	ZAAN	11.	NORHOLL	16.	HAL+IJMU
43.	NOHOLKAN	11.	NORHOLL		
44.	MARGKAN1	10.	IJSLAKES	1.	FRIELAND
45.	MARGKAN2	1.	FRIELAND	2.	B'GUMMER

Table 2.2 (continued)

Link Name	From Node	To Node
46. FRIEHARL	1. FRIELAND	
47. B'GM REC	2. B'GUMMER	2. B'GUMMER(b)
48. B'GMSINK	2. B'GUMMER	
49. STABOKAN	2. B'GUMMER	3. GRONETAL
50. H-H RECY	3. GRONETAL	3. GRONETAL(b)
51. EEMSKAN	3. GRONETAL	
52. MEPLDIEP	4. NEHIGH	9. IJSLEND
53. NOWILKAN	4. NEHIGH	3. GRONETAL
54. OMMERKAN	4. NEHIGH	5. OVIJVECH
55. OVIJVEC2	5. OVIJVECH	9. IJSLEND
56. OVIJKAN1	8. TWENTEND	5. OVIJVECH
57. TWENKAN1	8. TWENTEND	7. TWENMOND
58. MERWKAN2	12. A-R.MOND	17. LOPIKWAR
59. MERWKAN3	17. LOPIKWAR	18. VECHT
60. KANHOLIJ	17. LOPIKWAR	20. GOUDA
61. HOLIJSEL	20. GOUDA	19. IJSLMOND
62. KRIMPKAN	12. A-R.MOND	20. GOUDA(a)
63. WIERICKE	17. LOPIKWAR	21. MIDWEST
64. LEIDRIJN	13. UTR+MAAR	21. MIDWEST
65. GOUWE	20. GOUDA	21. MIDWEST
66. KATWIJK	21. MIDWEST	
67. HARLMEER	21. MIDWEST	16. HAL+IJMU
68. VECHT2	10. IJSLAKES	18. VECHT
69. ARKVECHT	18. VECHT	14. DIEMEN
70. ANGSTEL	14. DIEMEN	21. MIDWEST
71. AMSTEL	21. MIDWEST	15. AMSTEDAM
72. ZUIDWLM2	27. MAASLOO	34. WEER+MEY
73. WESNVERT	34. WEER+MEY	28. BORN+PAN
74. ZUIDWLM3	34. WEER+MEY	35. HELMOND
75. ZUIDWLM4	35. HELMOND	32. DBOS+BOX
76. WILHEKAN	35. HELMOND	36. OSTRHOUT
77. DONGE	36. OSTRHOUT	33. GERTRUID
78. MARK1	36. OSTRHOUT	37. FIJNAART
79. MARK2	37. FIJNAART	26. DLTALAKE
80. OVIJVECH		5. OVIJVECH
81. RIJN		6. UPRIVER
82. MAAS		27. MAASLOO
83. RUR/SWLM		30. ROER+BEL
84. NIERS		31. SAMGRALI

(a) Link capacity is zero in present infrastructure. Link was included for possible future expansion.

(b) Link simulates cooling of a power plant by recirculation (or in the case of Amer, by tidal action).

variables are classified according to the purpose or the origin of the flow, as we describe below.

For each node, there is a requirement that inflows during a decade equal outflows, plus any increase in the amount of water stored at the node. We call each such requirement a water balance constraint. Water balance constraints are further discussed in App. A of Vol. VA.

2.2. WATER INPUTS AND OUTPUTS OF THE INFRASTRUCTURE

To understand the water inputs and outputs of the water management infrastructure, it is necessary to first broaden our view to include all of the water in the Netherlands. The infrastructure defined above consists only of the structures containing surface water, but much water is also stored in, and passes through, the soil. It is convenient to describe the water in the Netherlands as residing in three pools, the moisture in the root zone, the moisture in the subsoil, and the surface water. These three pools are shown schematically in Fig. 2.2.

It is clear what we mean by surface water; it consists of all water in rivers, canals, streams, lakes, and reservoirs--i.e., the water found in the water management infrastructure. Conversely, the sum of root zone and subsoil moisture is all water stored in the ground. Following the convention established in PAWN's study of agriculture (Vol. XII), we have divided the moisture in the ground into two pools, namely the moisture directly accessible to the roots of crops, called the root zone moisture, and the moisture in the ground that is too deep for roots to access directly, called the subsoil moisture. Subsoil moisture could be further divided into the moisture in a saturated layer (called the groundwater; the depth at which the saturated layer is first encountered is the groundwater level), and an unsaturated layer. This division is not made in PAWN, although in Vol. XII a relation is developed between subsoil moisture and the groundwater level, in order that measurements of the groundwater level can be compared with model outputs.

The movements of water into and out of each pool, and between pools, are shown as arrows in Fig. 2.2. We will devote the remainder of this section to discussing these flows. Since MSDM is directly concerned only with the surface water, we limit our discussion largely to the flows involving the surface water pool, mentioning the others only to point out their indirect effects on surface water.

A point that deserves some discussion is why we have restricted our attention in MSDM to surface water alone, and have chosen to exclude groundwater. There are two classes of tactics that will affect groundwater. One is the day-to-day control of the amount of groundwater extracted. This is the only class of groundwater tactics that would be of interest in MSDM, because MSDM deals with only the very short term. The other class of tactics, of course, are those

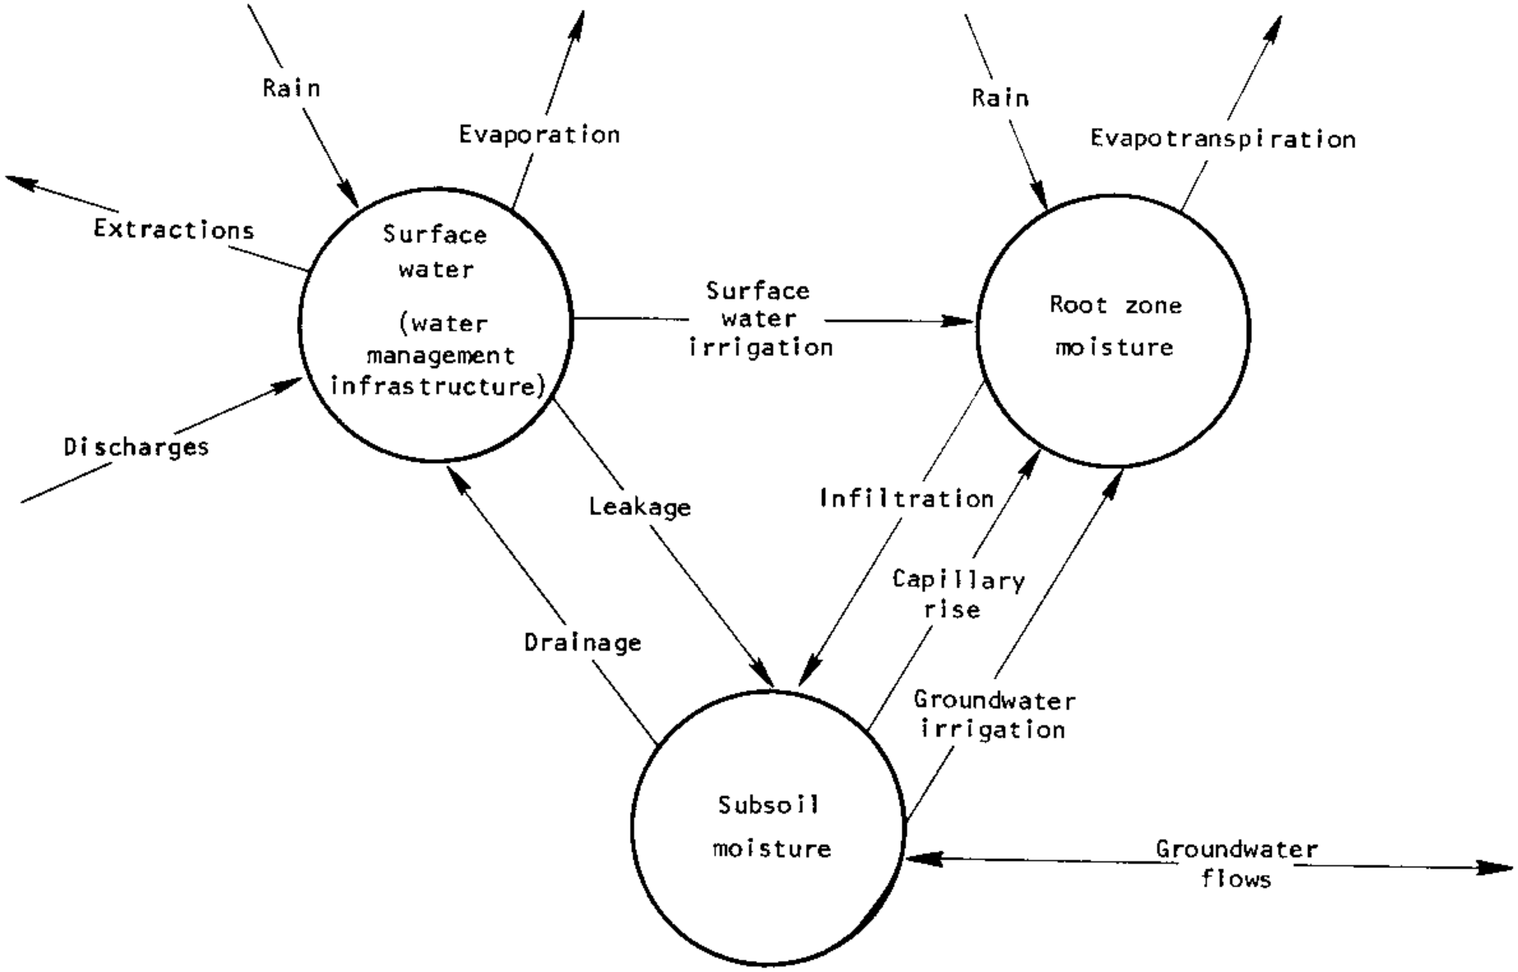

Fig. 2.2--Moisture pools and water flows in the Netherlands

which affect the amount of groundwater extracted in the longer term, for example by restricting the number and total capacity of the wells and pumps used for groundwater extraction.

Day-to-day groundwater tactics have three significant drawbacks. First, the major extractors of groundwater are industries and drinking water companies. Sudden reductions in the amounts they are permitted to extract would cause great hardship (see Ref. 2.1 and Vol. IV). Second, agriculture is responsible for the remainder of the groundwater extractions. Agriculture extracts groundwater from a large number of very small wells (less than 10 m^3/hr). It would be difficult to monitor or control the daily or weekly extractions from such a large number of wells (see Ref. 2.2 and Vol. IV).

Third, the effect of a reduction in groundwater extractions cannot be felt immediately, except by him who reduced his extractions. For example, suppose an extraction of groundwater were taking place a distance of one kilometer from a stream. Suppose further that the subterranean flow from the stream is the ultimate source of the water being extracted. If the extraction were suddenly reduced by 1 m^3/s, it is clear that the subterranean loss from the stream would eventually decline by that same 1 m^3/s. But it is possible to calculate, using data and methods from Ref. 2.3, that only 20-40 percent of that reduction will have occurred by the time one decade had passed. If the distance were as much as five kilometers, less than two percent of the reduction would have occurred in a decade.

Accordingly, we have rejected as unpromising tactics that would control day-to-day extractions of groundwater, and have instead concentrated on tactics that control long-term average extractions. Discussion of these results can be found in Vol. VII. In MSDM, we will hereafter consider groundwater extractions to be exogenously specified, and to no degree influenced by the managerial strategy.

2.2.1. River Discharges

The major input of fresh water to the surface-water pool, and indeed to the Netherlands as a whole, is from rivers. The Rijn contributes almost two-thirds of the fresh water of the Netherlands. The Maas brings only about five percent of the Netherlands' fresh water, but it serves a region not reached by Rijn water, and hence cannot be ignored. Small rivers, such as the Roer, the Swalm, and the Niers (which flow into the Maas) and the Overijsselsche Vecht (which serves part of the Northeast Highlands) also provide locally important supplies of fresh water.

Since river discharges are flows of water, they are represented as links in the MSDM network, with attached variables. The variables are called discharge variables, and are associated with links that originate outside the network, and terminate at a network node. Examples of such links in Fig. 2.1 are link 81 RIJN, and link 82 MAAS. The values of these discharge variables are specified as inputs to

MSDM, which are allowed to be respecified (i.e., which can vary) from decade to decade.

2.2.2. Sinks and Extractions

Most of the water that enters the Netherlands flows out again into the North Sea. The largest part enters the North Sea through the Rotterdam Waterway (links 4, 5, and 9 in Fig. 2.1) and the Haringvlietsluizen (link 22), with the Rotterdam Waterway carrying the larger amount when river flows are low, and the Haringvlietsluizen when flows are high. The Noordzeekanaal (link 34) also conducts important amounts of water into the North Sea even when river flows are low. Other losses to the North Sea are small when the rivers are low, but can be made as large as desired when the rivers are high, to get rid of unwanted water. A flow of water into the North Sea from a node of the network is represented by a <u>sink</u> variable.

Sinks are not the only flows leaving the network. There are extractions for such purposes as industial processes or drinking water (we defer extractions for agricultural purposes--irrigation--to a later section). Generally they are very small. We represent them by <u>extraction</u> variables, which are associated with links in the network that we have elected not to show in Fig. 2.1. There is such a variable (and hence such a link) leaving each node of the MSDM network.

The links corresponding to sink and extraction variables all originate at a node of the network, and terminate outside the water management infrastructure.

2.2.3. Rain and Evaporation

Rain contributes about 30 percent of the total freshwater supply of the Netherlands, but only a small fraction falls directly on surface water. Most of the rain (about 90 percent) falls on soil. Some of this rain will eventually reach the surface water pool indirectly, first infiltrating from the root zone to the subsoil, and then draining into the surface water, but this can happen only after a considerable delay.

Evaporation is responsible for large losses of water from the Netherlands. Water evaporates directly from bodies of open water such as lakes, rivers and canals, almost two-thirds of open water evaporation occurring from the IJssel lakes (node 10 in Fig. 2.1). Water also evaporates from the soil at a rate that depends on how moist the soil is. Unless the soil is very dry, however, the evaporation rate from soil is about 80 percent of the open water rate. Because the land area of the Netherlands exceeds the open water area by a factor of nine, a great deal more moisture evaporated from the soil than from open water.

Since MSDM is directly concerned only with surface water, its direct inputs include only that part of the rain that falls on surface water, and that part of evaporation that occurs from surface water. Neither of these quantities is represented by a variable. Instead, they are specified as inputs or outputs of water at each node in each decade. The amounts specified are calculated from the intensity of the rainfall and evaporation, measured in millimeters per decade for the decade being considered, times the surface water area represented by the node. All the surface area of the larger canals, rivers, lakes, reservoirs, and the like, is assigned to one node or another, so that essentially all the rain and evaporation from surface water is accounted for.

Rainfall on, and evaporation from the soil surface is not ignored in MSDM, however. How they are treated is outlined briefly in Sec. 2.2.6 below, and discussed at length in Chap. 6.

2.2.4. Surface Water Irrigation

A major influence on the surface-water pool occurs when the root zone dries out. If a farm is equipped for irrigation, it will then impose a demand for water upon the surface water system. Since large areas of the Netherlands are presently equipped for irrigation, and more irrigation equipment is expected to be installed, a large and growing demand can thus be imposed. Extractions of water by agriculture for irrigation are represented by irrigation variables.

In order to properly represent irrigation variables, certain auxiliary computations must be carried out. These computations concern the interactions of the subsoil and the root zone moisture pools with surface water. For much of the Netherlands, these computations can be done separately from MSDM, and the results simply provided to MSDM as input data. But for the part of the Netherlands that is irrigated from surface water--and hence may be influenced by managerial tactics--the computations must occur inside MSDM, interactively with the determination of the managerial strategy (see Chap. 6).

2.2.5. Leakage and Drainage

Significant amounts of water can be lost from large canals, especially in the highlands, and from small ditches, especially in low-lying areas, by leakage into the subsoil. Losses from large canals must be replaced if the desired flows are to be maintained, or the desired deliveries of water made to its downstream end. Losses from small ditches must be replaced if their water levels are to be maintained, a process not unnaturally called level control. The reverse can also happen, water draining out of the subsoil into the surface water.

In MSDM, these water movements are specified as inputs or outputs of water at each node in each decade. They are not represented by variables. Leakage from large canals in the highlands is considered to be the same in every decade, so a constant is specified for each of the several nodes from which this leakage occurs. For areas that are not irrigated or are irrigated from groundwater, the leakage and drainage rates are precalculated for each decade by a model called DISTAG (see Vol. XII and Chap. 6 of this volume), assigned to appropriate nodes, and presented to MSDM as inputs. For areas irrigated from surface water, the calculation of leakage and drainage is internal to MSDM, as we describe in Chap. 6.

2.2.6. Water Flows Indirectly Affecting Surface Water

There are several flows shown in Fig. 2.2 that only indirectly affect the surface water pool, and that have therefore not been discussed above. However, these flows do have an indirect influence on the surface water, through their influence on the root zone and subsoil moisture contents. The amount of moisture in the root zone affects the amount of surface water irrigation that will be done by agriculture, while the groundwater level, which depends on the total water content of the subsoil, affects leakage and drainage.

The flows in this category that affect the root zone moisture are infiltration, capillary rise, groundwater irrigation, rain, and evapotranspiration (evaporation from plants). Infiltration occurs when the root zone becomes saturated with water, and the excess moisture descends by gravity into the subsoil. Capillary rise is the reverse motion, water rising like oil in a wick from a wet subsoil to a dry root zone. Groundwater irrigation is the deliberate pumping of water from the subsoil for use as irrigation water, in place of surface water (the two are never used on the same plot of ground).

Infiltration, capillary rise, and groundwater irrigation also affect the subsoil moisture. In addition, the subsoil is affected by groundwater flow, the movement of water horizontally through the soil, sometimes over long distances. In the lowlands, such a flow brings highly saline water from the North Sea many kilometers inland. Here, the phenomenon is called seepage, and is more important for the salt it brings than for the water.

Water also flows into and out of the country under ground. This is similar to the underground flow of saline water from the North Sea to the lowlands, but here we refer to the underground flow of fresh water. The amount of water transported in this way is generally thought to be small, but is recognized to be significant in some locations (e.g., in Limburg). Moreover, it can be argued (App. C of Vol. XII) that significant losses and gains of water occur by underground flow even on a national scale.

All of these water flows are calculated either by DISTAG (Vol. XII) or by the procedure discussed in Chap. 6. If calculated by

DISTAG, the flows are never presented to MSDM as inputs. Rather, they are only reflected by their effect on leakage and drainage, which are in turn provided as inputs to MSDM, as discussed in Sec. 2.2.5 above.

2.3. WATER WITHIN THE SURFACE WATER SYSTEM

Water that remains within the surface water system during a decade either moves from node to node or is stored at a node. In MSDM, two kinds of variables represent these possibilities. Movement is represented by flow variables, and storage by level variables.

A flow variable is associated with a link having two nodes as endpoints. One node is designated as the point at which the flow originates, and the other the point at which the flow terminates. Some flow variables may take on negative values, in which case the flow is understood to reverse its direction. Then the node specified as the terminus functions as the origin, while the node specified as the origin functions as the terminus.

A level variable is associated with a single node. It is called a level variable because the MSDM model reports the water level at a node at the beginning and end of each decade, although it actually calculates the water volume. The volume is required to be nonnegative, although the water level can be either negative or positive, compared to the reference level (usually taken to be NAP). The water level at the beginning of the decade must be either specified exogenously or remembered from the end of the previous decade. The model then calculates the water volume and level at the end of the decade.

2.4. REQUIREMENT FOR WATER CONSERVATION

At each node, MSDM requires that water be conserved. This means that the inputs to a node from discharge and flow variables, and from rain and groundwater flow, must equal the outputs from the node via sink, extraction, irrigation and flow variable, and evaporation and groundwater flow, plus any increase in storage that has occurred during the time period. That is, we assume that all the gains and losses of water are accounted for and that they take place at nodes. We call such an equation a water balance constraint for the node in question.

At best, this assumption is only an approximation. Losses do take place along the lengths of numerous Dutch canals, as some of the water they carry leaks into the ground. Rain falls on canals and rivers, and water evaporates from them. However, in MSDM we calculate these gains and losses from each link, and assign them to the nodes that constitute the endpoints of the link.

In MSDM, we also assume that the links have zero volume. (From the viewpoint of water quantity, this is the same as assuming that the volume of water in a link remains at any constant level, although the two assumptions have different implications for water quality.) However, water levels in rivers and canals do change from time to time, thus changing the amount of water stored in links, and making the flow entering the upstream end of a link unequal to the flow leaving the downstream end. We have not adjusted the water conservation equations for this possibility in MSDM, because we consider that the error thus introduced is negligible.[2] This is because it is Dutch policy (and for very good reason) to maintain the water levels in canals constant if at all possible. In the rivers, where level changes do occur, the amount of water flowing through a link in a decade is so much larger than any change in storage in the link, that the change in storage need not be considered.

2.5. HYDROLOGIC CONSTRAINTS

At two locations in the network, the flows are partially determined by the hydrological characteristics of the rivers,[1] and are not under the complete control of the water manager. These locations are upper rivers, where the Rijn divides into three branches as it enters the country (the Waal, the Neder-Rijn, and the IJssel, represented by links 1 WAAL1, 6 NEDRIJN2, and 35 IJSSEL1), and lower rivers, where the Maas, the Waal, and the Lek enter a fairly complex network of river branches that lead to the North Sea via the Nieuwe Waterweg and the Haringvlietsluizen.

At upper rivers, the division of water between the three branches depends on the flow in the Rijn entering the Netherlands, and on the setting of the weir at Driel. Data supplied by Upper Rivers Department of the RWS enabled us to derive relationships between these two quantities and the flows in the three branches (see Vol. XI). Happily, these relationships can be expressed as functions that are linear in all flows except the Rijn flow into the Netherlands. It has therefore been quite simple to incorporate these relationships into MSDM, as described in App. A of Vol. VA.

The lower rivers hydrological relationships were developed by statistically fitting equations to the output of a model called IMPLIC. The IMPLIC model calculates the instantaneous flows in the lower rivers section of the network (see Fig. 2.3). This part of the network is affected by the tides in the North Sea, and hence show a periodic motion. We averaged the instantaneous flows in the various links of the lower rivers network over a period of two tide cycles (about 25 hours), and related the average flows thus obtained to the inputs into the lower rivers network of the Maas, the Waal, and the Lek. Other, smaller inputs and outputs to lower rivers were also considered, such as the flow in the Hollandsche IJssel, the discharges at the Volkerak locks and the Haringvlietsluizen, and various extractions occurring in the area. These relationships and the process of their derivation are described in detail in Vol. XI, and the manner in which they are incorporated in MSDM is described in App. A of Vol. VA.

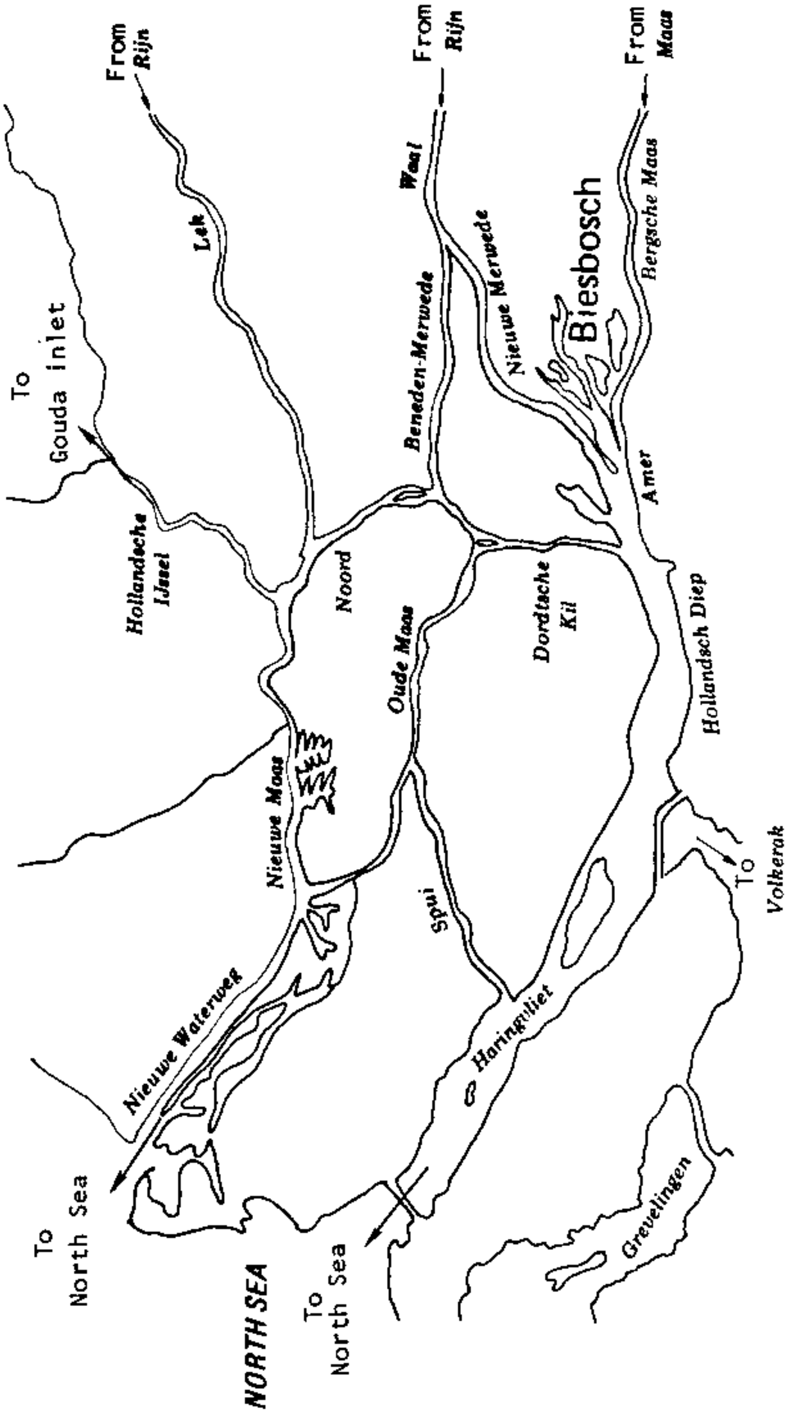

Fig. 2.3--Lower Rivers network

2.6. LINK CAPACITIES, REQUIREMENTS, AND PUMPING COSTS

Most links in the PAWN networks have both upper and lower bounds on their allowed flows. Similarly there are upper and lower bounds on the water levels allowed at nodes with storage. Some of these bounds derive from technical, physical, or safety limitations on the amount of water a link can carry or a node can store. Bounds of this type are called capacities. Other bounds derive from a need to maintain certain minimum flows in designated links or minimum levels at nodes with storage. These bounds are requirements. It is usually, but not always true, that upper bounds are capacities and lower bounds are requirements.

Some examples of capacity limitations are:

- The water level in the IJssel lakes (node 10 IJSLAKES) must not exceed NAP - 20 cm, because otherwise certain nature areas will be partially inundated, and certain harbors and boatworks along its shore will be rendered unable to operate.
- The locks at IJmuiden, at the end of the Noordzeekanaal (link 34 NZKANLSL), cannot discharge more that 230 m^3/s into the North Sea, because the sizes of the locks and associated sluices, combined with the water levels inside and out, prevent more than 80 m^3/s from draining out by gravity, while the pumping station there has a capacity of 150 m^3/s.
- Some capacities must be expressed in a more complex fashion. For example, when the IJssel lakes are at a level of NAP - 40 cm, only 93 m^3/s may be brought into Friesland province via the Prinses Margrietkanaal (link 44 MARGKAN1). If the lake level is at NAP - 20 cm, the capacity increases to 135 m^3/s. The increase in lake level increases both the area through which water can flow from the lake into the canal, and the head difference that provides the driving energy.

Some examples of requirements are:

- The water level in the IJssel lakes must not be allowed to fall below NAP - 40 cm, because certain areas drawing on the IJssel lakes for their critical water needs would then be unable to withdraw water. Their intake works would no longer be submerged.
- A minimum flow of 7 m^3/s is required to flow from North Holland into the Noordzeekanaal at Zaandam (link 42 ZAAN), in order to dilute the substantial discharges of industrial waste that enter the water near Zaandam.
- A treaty with Belgium requires that, of the flow in the Maas at Monsin (in Belgium), a total of 22.4 m^3/s is reserved for use by Belgium. From a purely technical standpoint, the Netherlands could reduce Belgium's "take" considerably, since

much of the water must flow into the Netherlands before reaching the diversion points to Belgium.
- Locks at various locations will leak some unavoidable amount.
- Industrial withdrawals are considered requirements because it is so costly to reduce them in the short term (although not in the longer term).

It may be necessary to pump water along some links in order to achieve some of the flows of which they are capable. But pumping requires energy, and energy costs money. Some examples of links with associated pumping costs are:

- Pumping is required to lift water from the IJssel River into the Twenthekanaal (link 57 TWENKAN1). Without pumping, there is a leakage of 1.75 m^3/s through the shipping locks there, and much more can be allowed to drain out of the canal if desired. But to reduce the flow from the canal into the river below 1.75 m^3/s, or to reverse the flow, costs 0.00187 Dfl/m^3 for pumping energy.
- As mentioned above, up to 80 m^3/s can flow by gravity from the Noordzeekanaal into the North Sea at IJmuiden (link 34 NZKANLSL). If the water manager wishes to increase the flow beyond 80 m^3/s, he must pump, at a cost of 0.00036 Dfl/m^3.

A complete list of capacities, requirements, and pumping costs can be found in Vol. XI and in App. A of Vol. VA. Reasons for these constraints are to be found in Vol. XI.

2.7. THE WATER DISTRIBUTION PROBLEM IN MSDM

In MSDM, we have formulated the problem of choosing the optimal water distribution as:

$$(2.1)\qquad \begin{cases} \text{Minimize (Total cost proxy objective function)} \\ \\ \text{Subject to: WATER BALANCE CONSTRAINTS} \\ \qquad\qquad\quad \text{UPPER RIVERS HYDROLOGIC CONSTRAINT} \\ \qquad\qquad\quad \text{LOWER RIVERS HYDROLOGIC CONSTRAINTS} \\ \qquad\qquad\quad \text{REQUIREMENTS} \le \text{FLOWS} \le \text{CAPACITIES} \end{cases}$$

In the earlier sections of this chapter, we have briefly discussed the constraints in Problem (2.1), and further discussion can be found in App. A of Vol. VA.

The total cost proxy objective function provides a measure of the costs and losses incurred by all water users and water uses. As this function depends on the flows in network links and levels at nodes with storage, it distinguishes among water distributions. In Problem (2.1), the task is to choose a distribution for which the value of the objective function is as small as possible, consistent with the constraints.

The total cost proxy objective function includes terms for direct costs of distributing water, such as the cost of energy for pumping. It also includes terms for losses suffered by such users as shipping, agriculture, and power plants. Some of the users' losses are directly related to the water distribution. For example, agriculture suffers if deprived of water for irrigation, and shipping suffers if flows in rivers are inadequate to maintain shipping depths. Other users' losses are less direct. For example, power plants may have to reduce heat discharges at a node, in order to meet the thermal water quality standard. As explained in Chap. 5, this entails additional power generation costs. But these costs could be avoided if the water distribution brought sufficient cooling water past the power plant. It is possible to infer a function that relates the extra generating costs to the flows describing the water distribution.

In addition, water quality standards are imposed on pollutants other than heat. The existence of a standard at a particular node may indirectly impose a requirement that water be available to dilute a discharge of pollutants into that node. Even though no actual direct costs are involved, it is possible to represent such quality-related requirements as terms in the objective function.

Later chapters discuss the objective function in detail. In particular, Chap. 8 discusses how water quality aspects come to be included in the water distribution problem.

2.8. DEGREES OF FREEDOM FOR WATER DISTRIBUTION

The water manager is free to choose any water distribution that satisfies both the constraints that water be conserved at each node, and the hydrologic constraints at Upper and Lower Rivers. Once one such water distribution has been found, however, the water manager can exercise his considerable degrees of freedom to realize others. For example, he can increase the amount of water diverted from the Waal at Tiel (node 23). Accompanying this change must be a compensating decrease in the flow in the Waal downstream of Tiel (links 3 and 4, and also link 8). Also, the diverted water must go somewhere, either down the Lek (links 7 and 9) or up the Amsterdam-Rijnkanaal (links 29, 30, and 32) and out to the North Sea via the Noordzeekanaal (links 33 and 34). These degrees of freedom are shown in Figs. 2.4a and 2.4b.

We define a degree of freedom to be any set of adjustments to the flows in the links and the water levels at nodes with storage that does not change the compliance of the water distribution with the water conservation constraints or the hydrologic constraints. That is, if a particular constraint was satisfied before the adjustment, it will be satisfied after, while if a constraint was violated

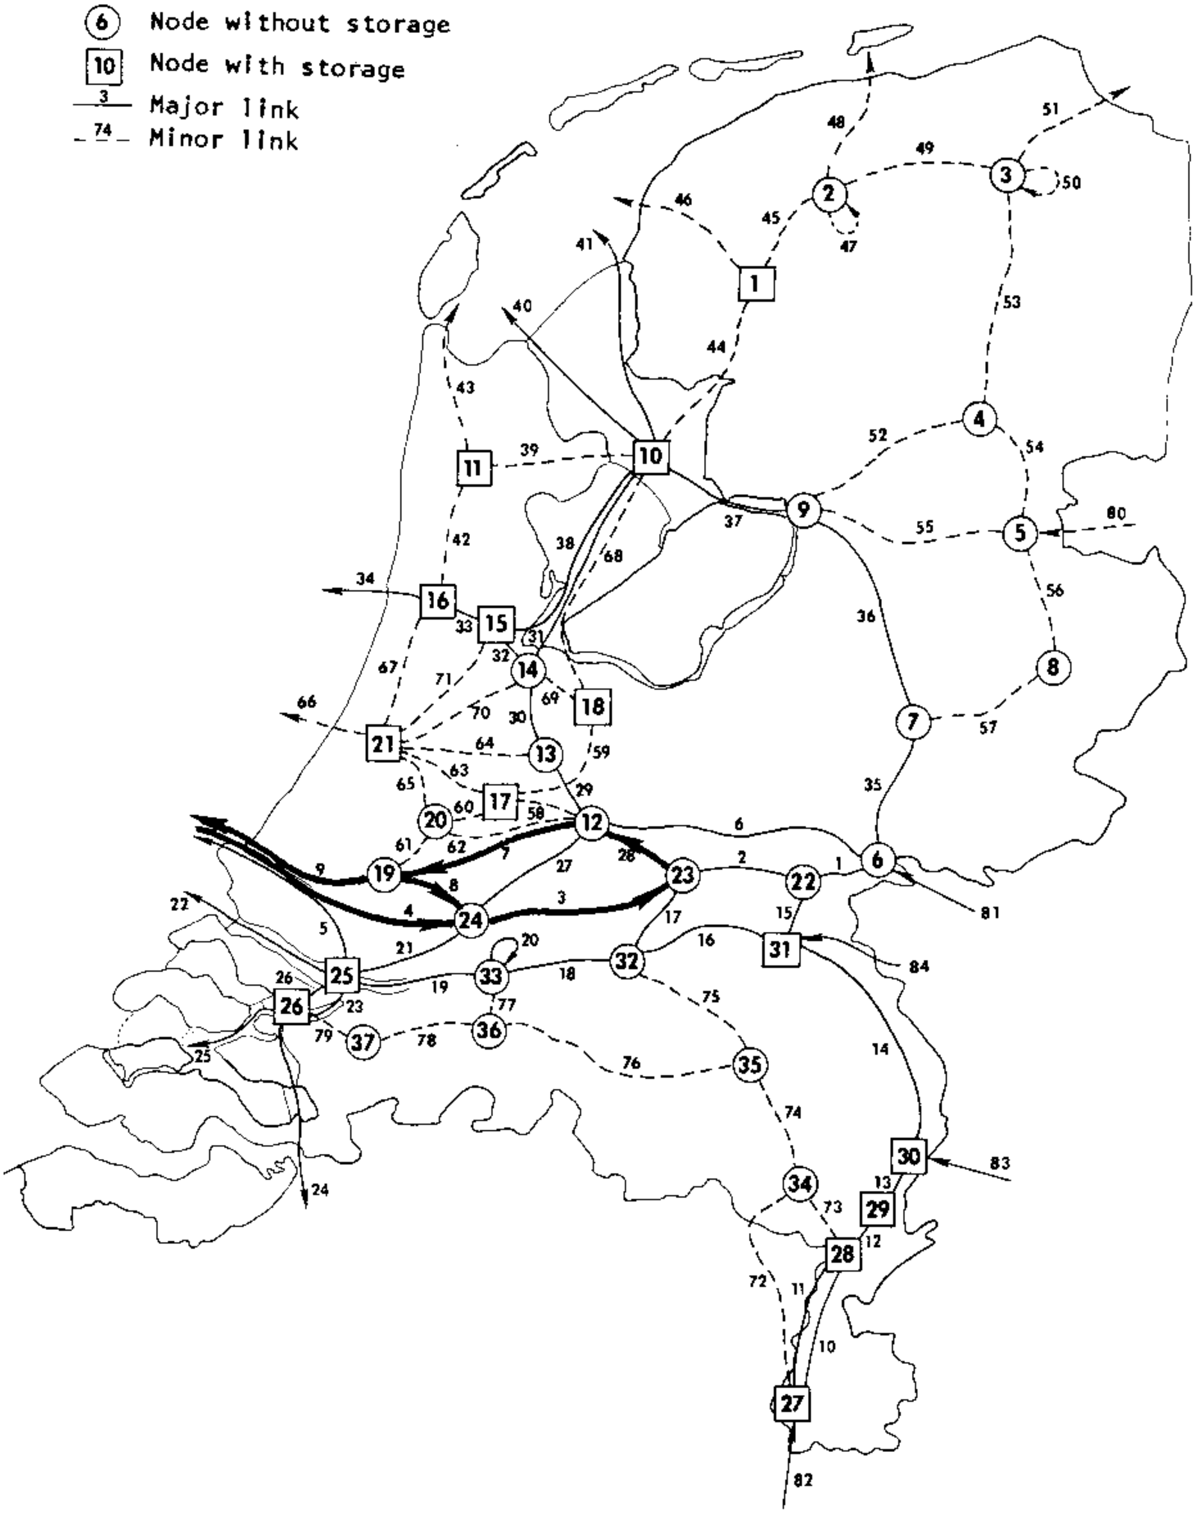

Fig. 2.4a--Degrees of freedom involving withdrawals at Tiel: increase flow in the Nieuwe Maas (link 9) at the expense of the Oude Maas (link 4)

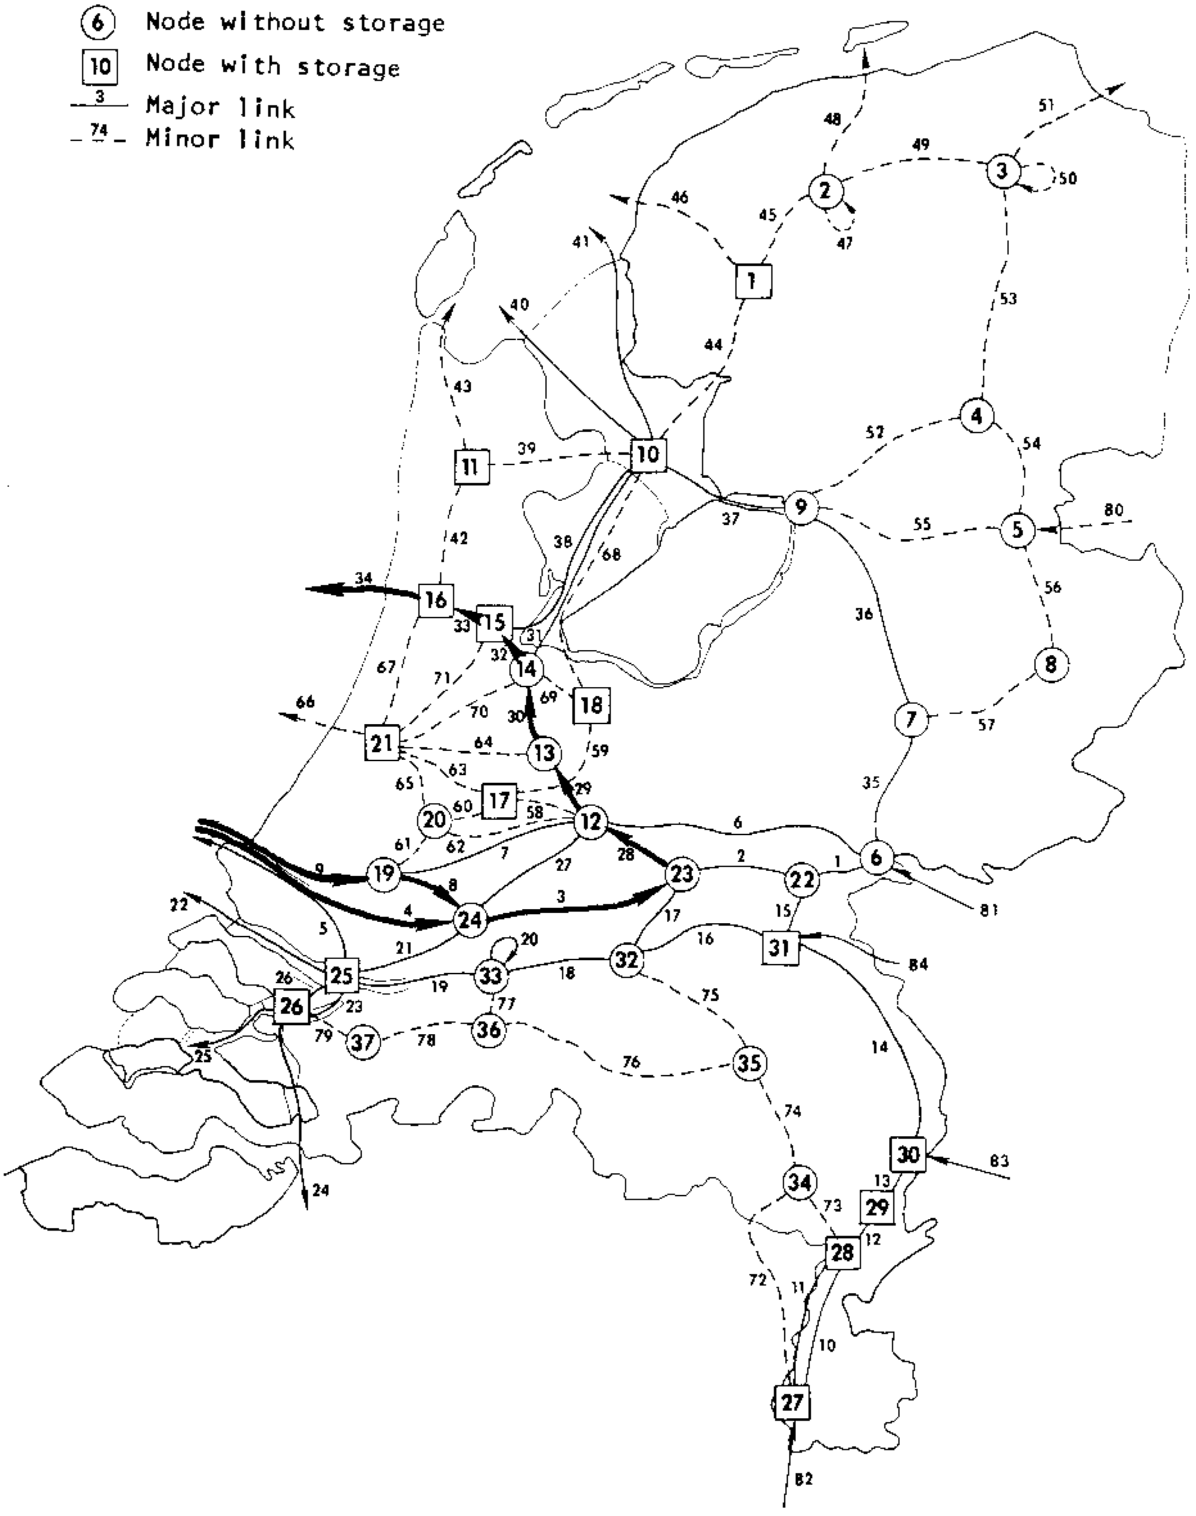

Fig. 2.4b--Degrees of freedom involving withdrawals at Tiel (node 23): increase flow in Amsterdam-Rijnkanaal (links 29, 30, 32) and Noordzeekanaal (links 33, 34) at the expense of the Oude Maas (link 4) and Nieuwe Maas (link 9)

before, it will also be violated after. A little thought will convince the reader that a degree of freedom corresponds to a path or a combination of paths through the network that takes water either from some point outside the network or from a node with storage and delivers it to a similar point or points. Thus, the two example degrees of freedom shown in Figs. 2.4a and 2.4b take water that would otherwise flow into the North Sea (outside the network) via the Waal and the Oude Maas, send it inland along the Waal to Tiel, north from Tiel to the Lek, and from there either west along the Lek, through the Rotterdam Waterway, and into the North Sea, or north along the Amsterdam-Rijnkanaal, into the Noordzeekanaal, and into the North Sea. Of course, the total flow in the Waal is never actually inland. But the degrees of freedom described here involve adjustments that reduce the flow in the Waal, and hence may be thought of as sending some water upstream to "cancel" a similar increment of flow downstream.

From the two degrees of freedom we have used as examples we can construct a third. This one would take water from the Rotterdam Waterway, send it inland along the Lek (i.e., reduce the flow on the Lek), and thence north along the Amsterdam-Rijnkanaal, into the Noordzeekanaal, and out to sea. This new path is the difference between the first two. That is, we can operate the path that takes water from the Waal and sends it north along the Amsterdam-Rijnkanaal forward, and simultaneously operate the path that takes water from the Waal and sends it west along the Lek backwards, and the result will be the new path.

Notice that the example degrees of freedom both appear simple, in that they involve only a small fraction of the links in the network. It is possible, of course, to define degrees of freedom that are not simple. We could, for example, operate both of the example degrees of freedom simultaneously in a forward direction, circulate water through the Northeast Highlands (links 52, 54, and 55), through the Eastern Highlands (links 57, 56, 55, and 36), and so forth. But such complex paths would not help one to understand the possibilities open to the water manager. Furthermore, it is often true that a simple degree of freedom can be exercised by the water manager in practice by a simple action. The water manager need only open a sluice gate on the Amsterdam-Rijnkanaal between Tiel and the Lek to exercise the Oude Maas-Waal-Lek-Rotterdam Waterway degree of freedom. He need only open this same sluice gate, plus another on the Amsterdam-Rijnkanaal to the north of the Lek, and one on the Noordzeekanaal at IJmuiden, to exercise the Oude Maas-Waal-Amsterdam-Rijnkanaal-Noordzeekanaal degree of freedom.

The two example degrees of freedom are important for water management, since they can control large amounts of water and they have (as we shall see later) significant consequences for water users. But there are other important degrees of freedom as well. In the southern part of the country, the water manager may deplete the weir ponds (stuwpanden) along the Maas and letting their water run downstream. The weir ponds are represented by nodes 27, 28, 29, 30, and 31. The water flows via links 10 (or 11), 12, 13, 14, 16, 18, 19, and 5 to the North Sea. This degree of freedom is exercised by opening the several weirs with which the Dutch have canalized the river Maas.

The water manager can flush the Zoommeer by exercising a degree of freedom. This requires that he reduce the flow north from the Haringvliet in link 5, and increase the flow south through the Volkerak (link 23) and through either the Kreekrak locks (link 24) or the Philipsdam locks (link 25). This degree of freedom is exercised by opening the sluices in the Volkerak locks and in either the Kreekrak or Philipsdam locks, as desired.

Another important degree of freedom involves flushing the Noordzeekanaal with water from the IJssel lakes. Water is obtained by lowering the level of the IJssel lakes (node 10), and sent into the Noordzeekanaal by any of a number of pumping stations, the largest of which are Zeeburg and Schellingwoude (represented by link 38). Then the water flows west through the Noordzeekanaal and into the North Sea (links 33 and 34). This degree of freedom is controlled by the pumps on link 38, and by the sluices and pumps at IJmuiden at the mouth of the Noordzeekanaal (represented by link 34).

Exercising yet another degree of freedom, the water manager can lower the level of the IJssel lakes (node 10) by discharging water through the Afsluitdijk into the North Sea (links 40 and 41). This is accomplished simply by opening some sluice gates.

One of the most important degrees of freedom involves increasing the flow in the Neder-Rijn, and reducing the flows in the Waal and IJssel rivers. This degree of freedom is exercised by opening the weir at Driel on the Neder-Rijn, as well as the other weirs along that river. The Upper Rivers hydrologic constraint dictates that if the weir opens to allow an additional quantity of water to flow down the Neder-Rijn, the IJssel River flow will be reduced by 42 percent, and the Waal flow by 58 percent, of that quantity. The reduction in the IJssel River flow lowers the level of the IJssel lakes (node 10), and reduces the flow in links 37, 36, and 35. The reduction in the Waal flow reduces the flow into the North Sea via the Oude Maas (link 4), and reduces the flow on the links of the Waal (links 3, 2, and 1). The water obtained from these two sources flows west along the Neder Rijn (link 6) and the Lek (link 7) through the Rotterdam Waterway (link 9) and into the North Sea. Link 8 is also involved in this degree of freedom because of the Lower Rivers hydrologic constraints.

These five additional degrees of freedom are shown in Figs. 2.5a-2.5e.

These seven are by no means all of the degrees of freedom, but between them they offer the greatest possibilities to control flows in the major waterways. In addition, they are the degrees of freedom most likely to be under the control of a national water manager. The other degrees of freedom are more likely to be at least partly controlled by local or regional interests. This does not mean that

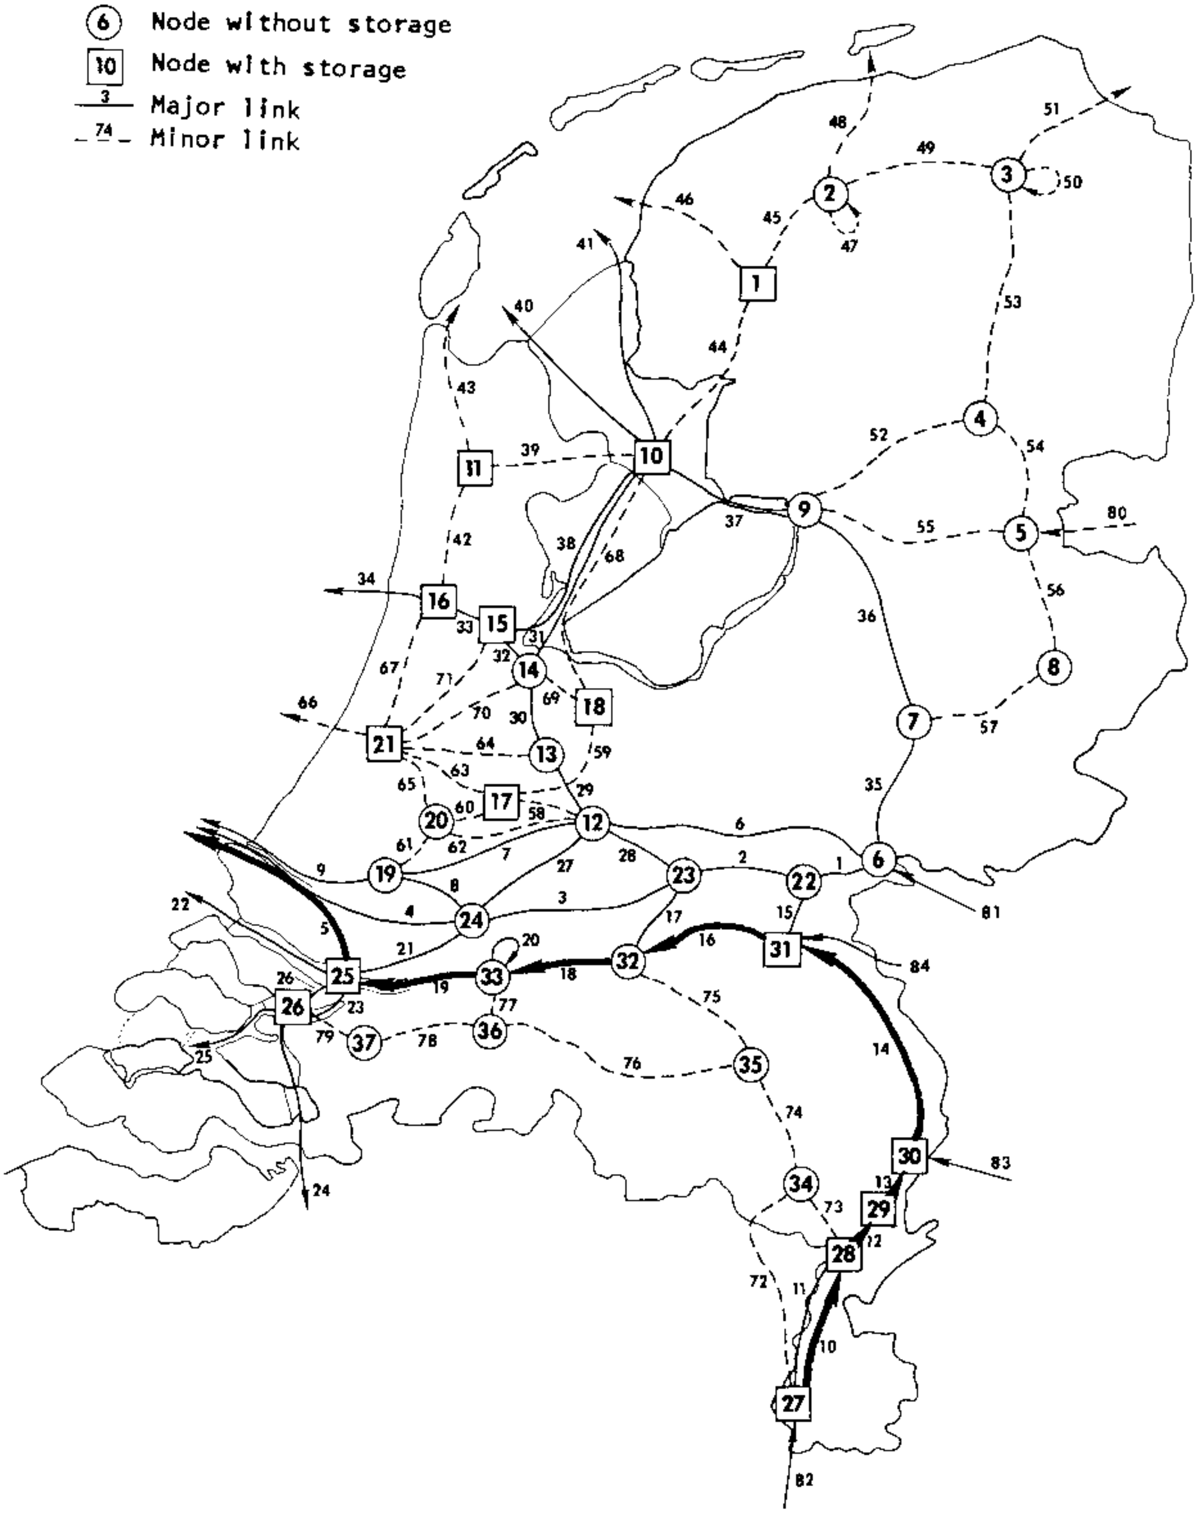

Fig. 2.5a--Five additional degrees of freedom: deplete stuwpanden in the Maas (nodes 27, 28, 29, 30, 31) to augment flow farther downstream (links 16, 18, 19, 5)

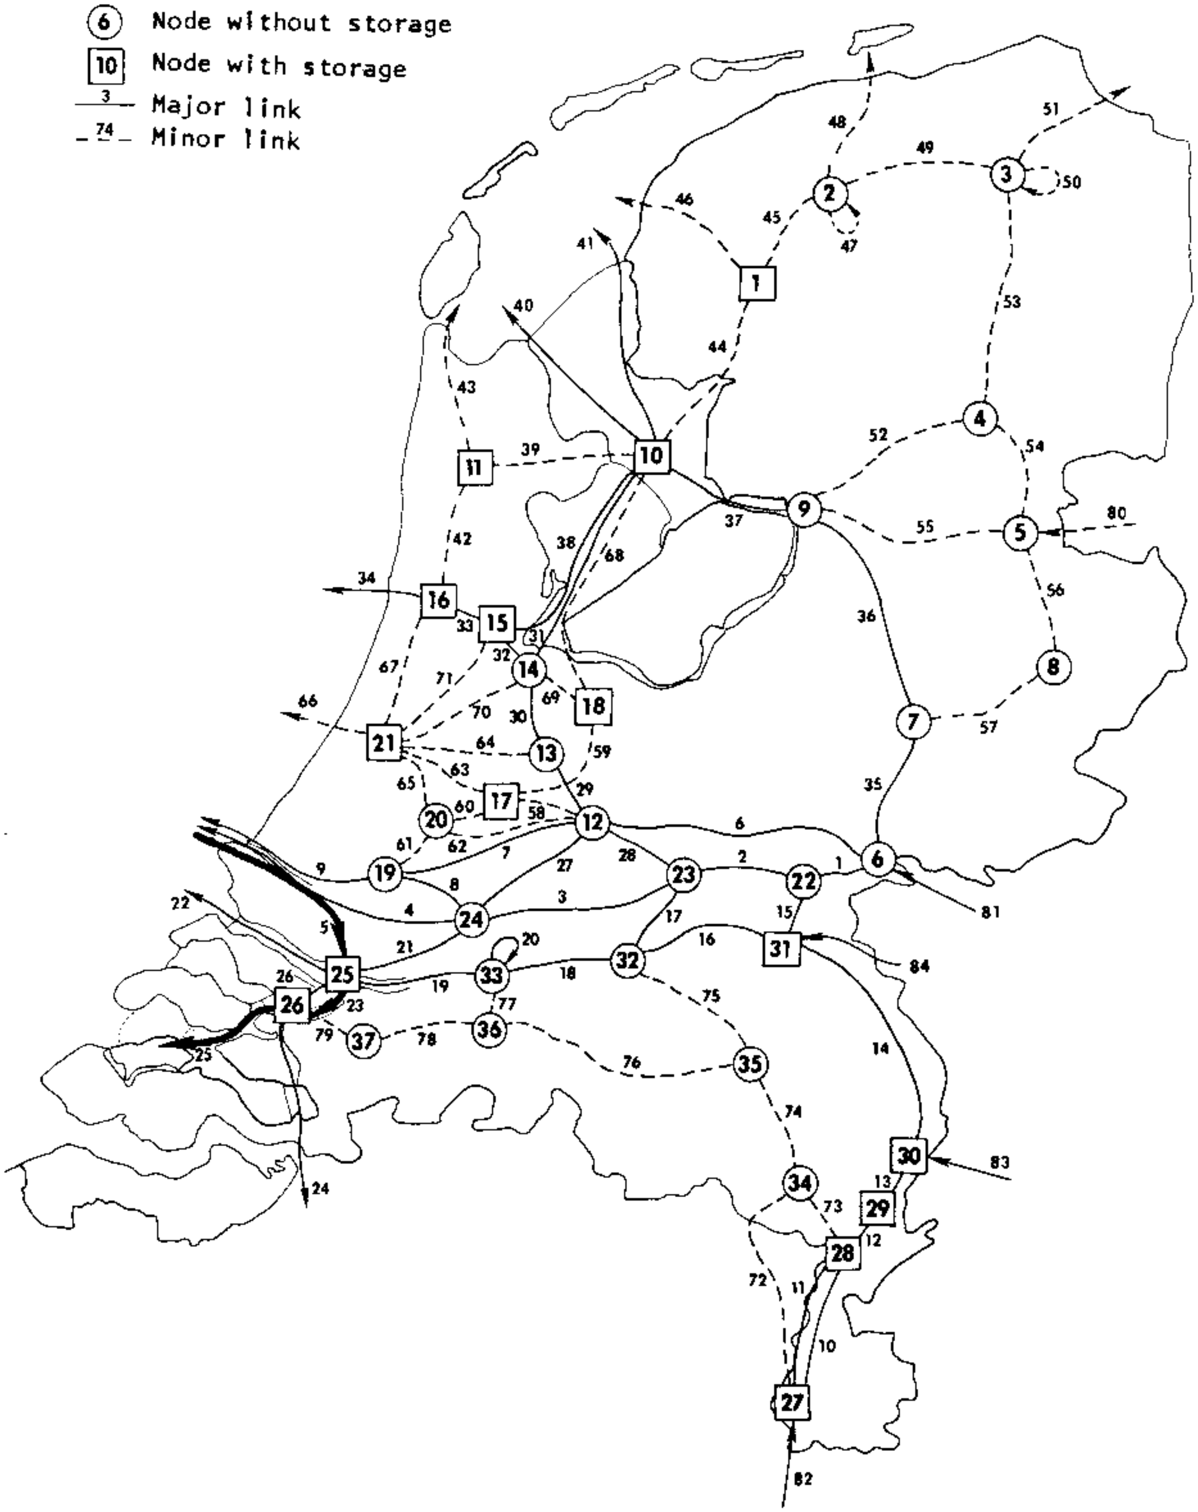

Fig. 2.5b--Five additional degrees of freedom: flush the Zoommeer (node 26)

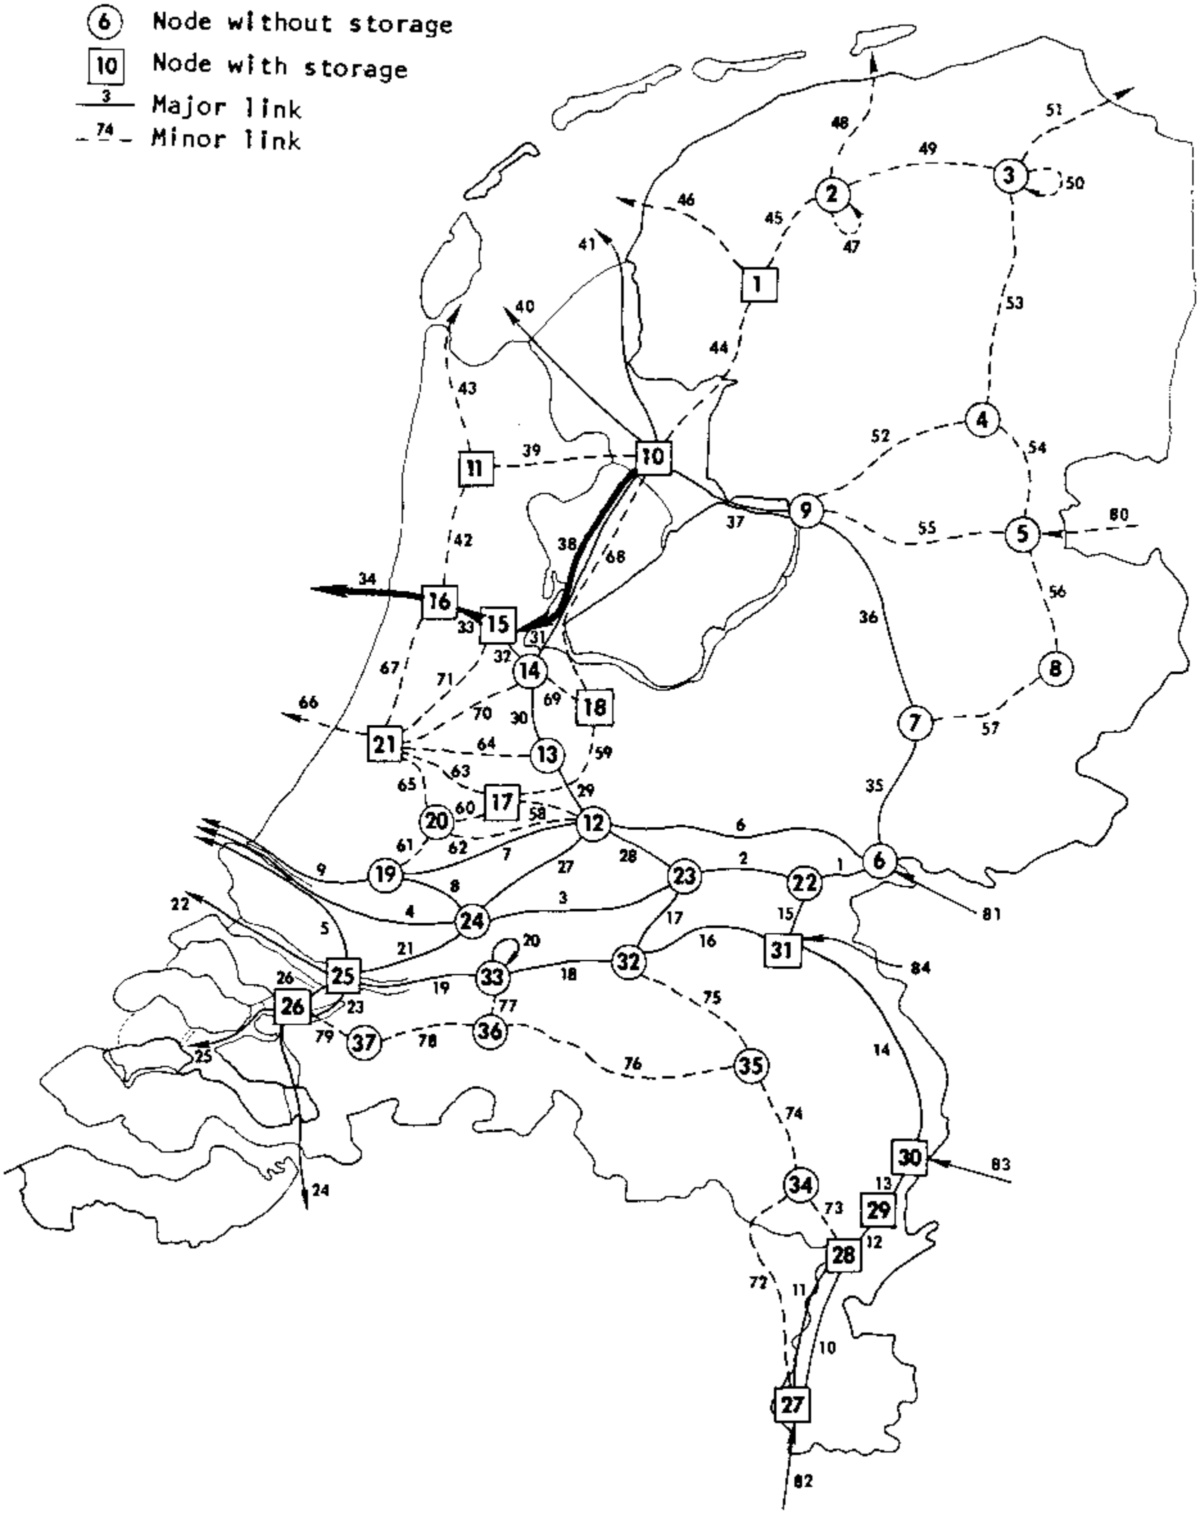

Fig. 2.5c--Five additional degrees of freedom: flush the Noordzeekanaal (links 33, 34) with water from the IJssel lakes (node 10)

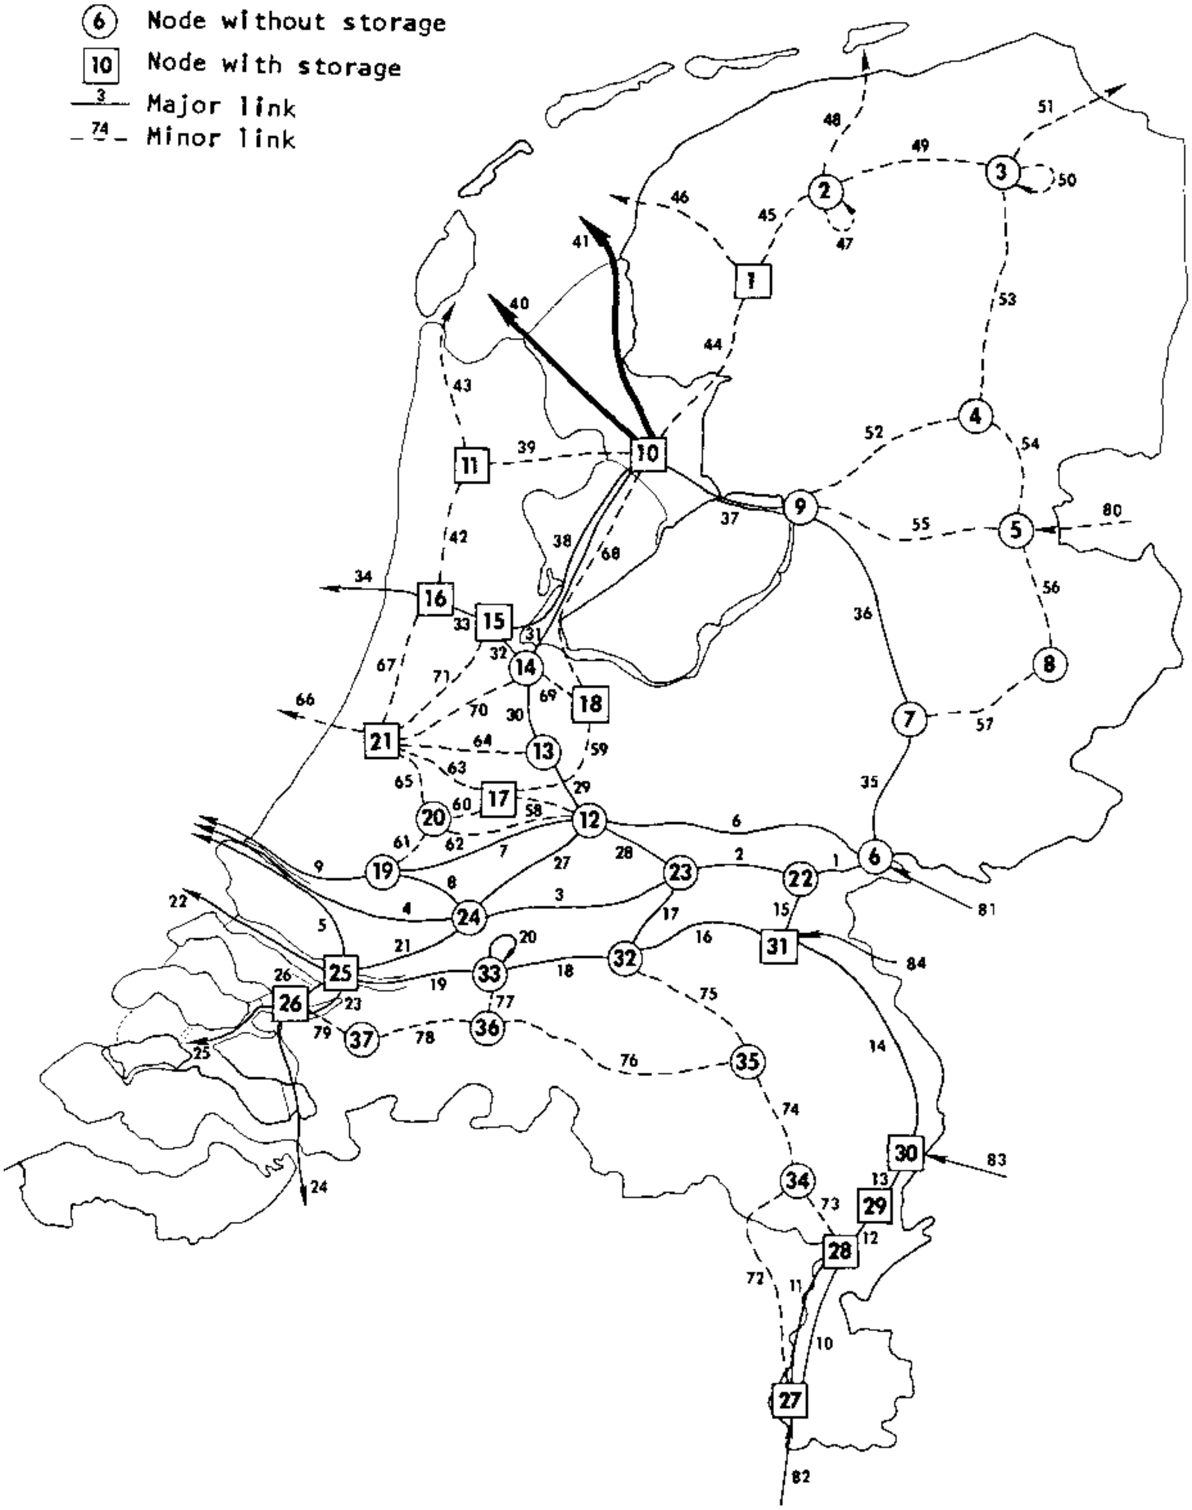

Fig. 2.5d--Five additional degrees of freedom: lower the water level in the IJssel lakes (node 10) by discharging through sluices in the Afsluitdijk (links 40, 41)

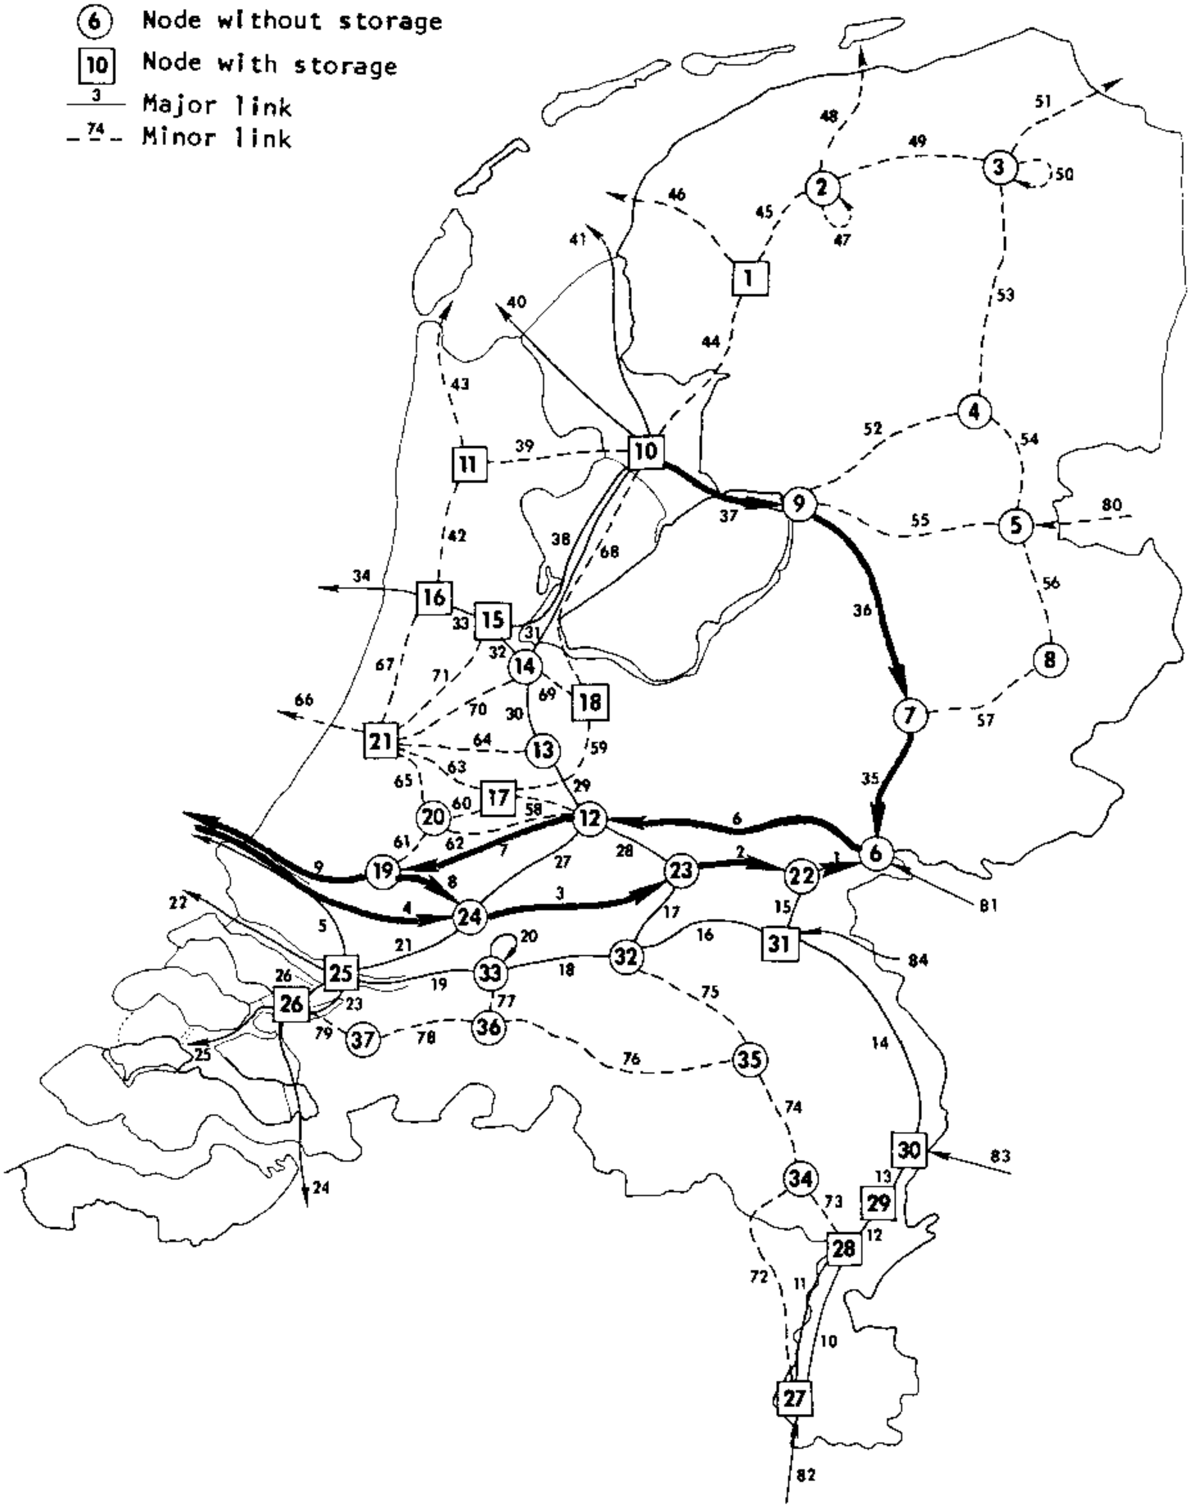

Fig. 2.5e--Five additional degrees of freedom: reduce flow in the Waal (links 1, 2, 3, 4) and the IJssel (links 35, 36, 37) to augment flow in the Lek (links 6, 7, 9) by opening weir at Driel

the more local degrees of freedom are unimportant, but only that they contribute less to both an overall understanding and an overall control of the water distribution.

No degree of freedom can be exercised arbitrarily. If the IJssel lakes are at their minimum allowed level, the water manager cannot lower them further, either to discharge water through the Afsluitdijk or to flush the Noordzeekanaal. Even if the IJssel lakes are at a high level, the pumps at Zeeburg and Schellingwoude (link 38) cannot be operated at more than their combined capacity (93 m^3/s). In general, upper and lower bounds on flows in various links and levels at nodes with storage will limit the extent to which the water manager can exercise each degree of freedom.

This observation is especially important when considering the effects of new infrastructure. A new canal might add entirely new degrees of freedom to those presently available to the water manager. Suppose, for example, that the St. Andries Connection were built between Tiel (node 23) and Den Bosch (node 32). This possible new link is represented in the present network as link 17, which presently has a capacity of zero. If the connection were made, increasing the capacity to twenty or fifty or even one hundred cubic meters per second, the water manager would be enabled to divert water from the Waal downstream of Tiel, in order to augment the flow in the Maas downstream of Den Bosch. If in addition pumping stations were installed along the Wilhelminakanaal (link 76) or the lower portion of the Zuid-Willemsvaart (link 75), as well as perhaps the middle portion of the Zuid-Willemsvaart (link 74), the flow in the Waal could be reduced in favor of increasing the deliveries of water to the Southern Highlands (nodes 34, 36, and 36). These new possible degrees of freedom are shown in Figs. 2.6a-2.6c.

Although it is useful to view new infrastructure in terms of the degrees of freedom that it adds to the present network, there are other viewpoints as well. Some new infrastructure merely increases the range over which a present degree of freedom might be exercised. This will often be the result of adding to existing capacity. Other new infrastructure may be designed to reduce or eliminate some element of the cost associated with exercising an existing degree of freedom. An example of this is building a structure to narrow the Waal below Tiel so that diversions of water north from the Waal at Tiel will not reduce the depth, and hence harm shipping. However, the notion of degrees of freedom is a very useful one to keep in mind when examining both the present infrastructure and proposed changes to it.

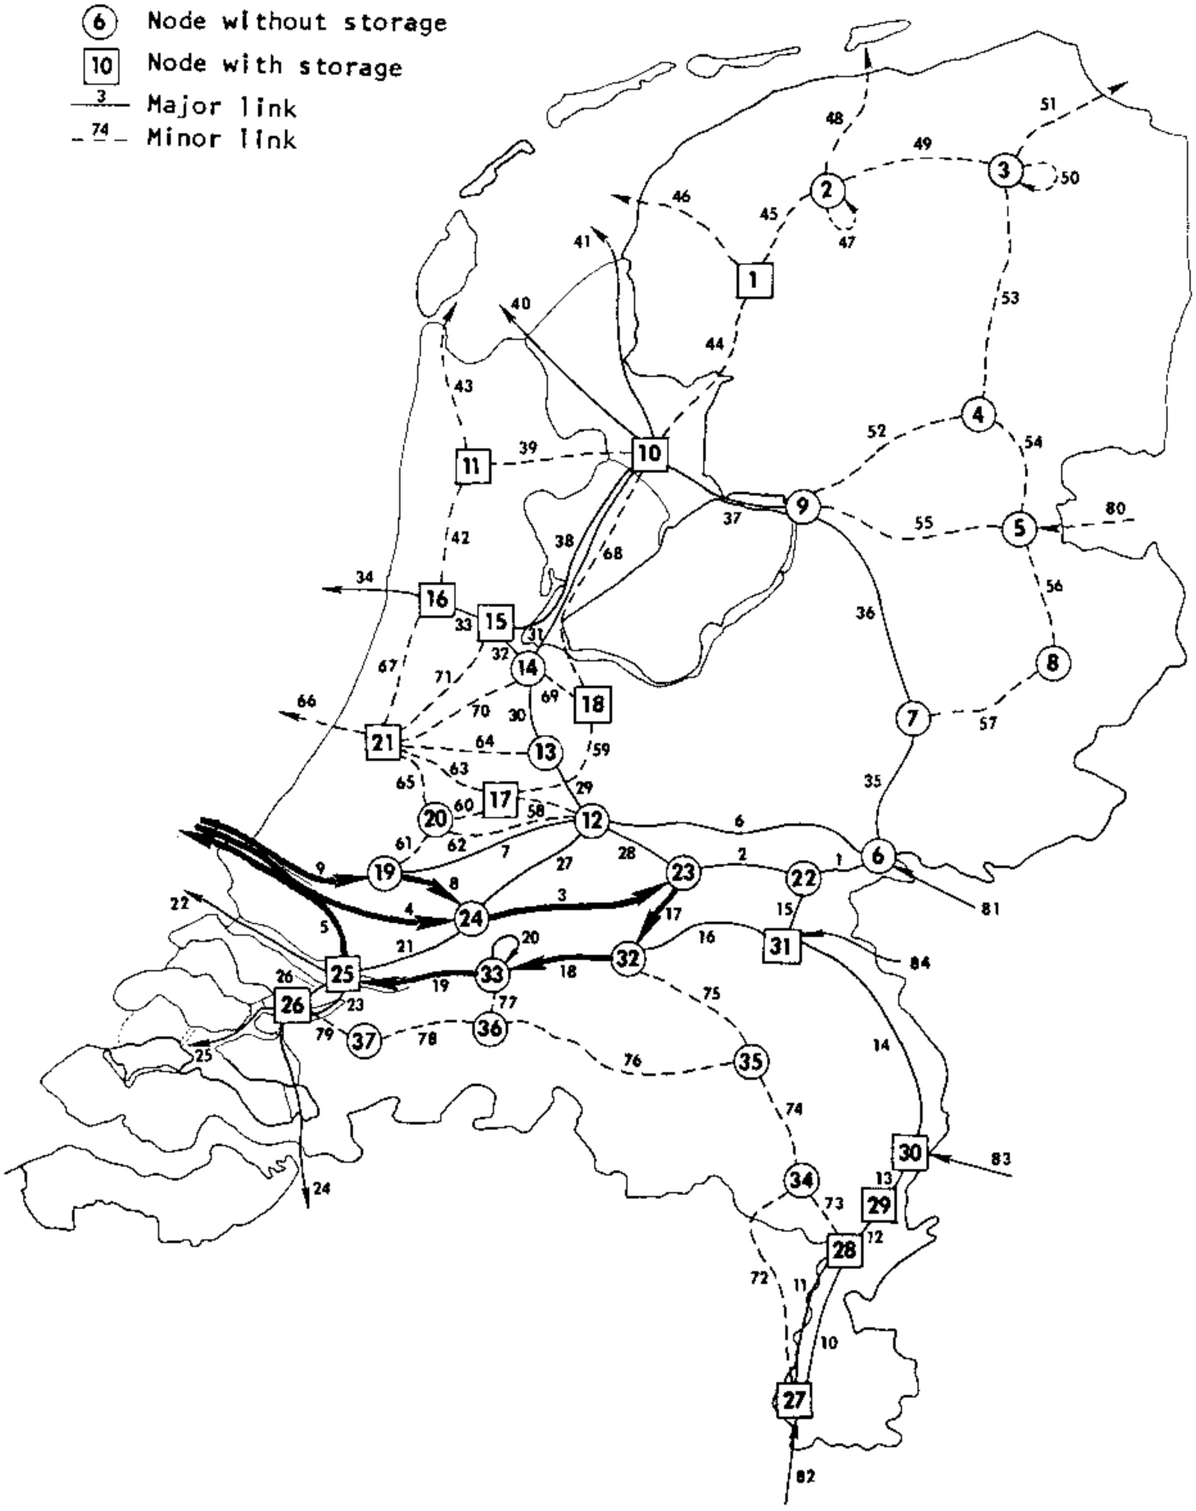

Fig. 2.6a--Degrees of freedom added by new infrastructure: add St. Andries Connection (link 17)

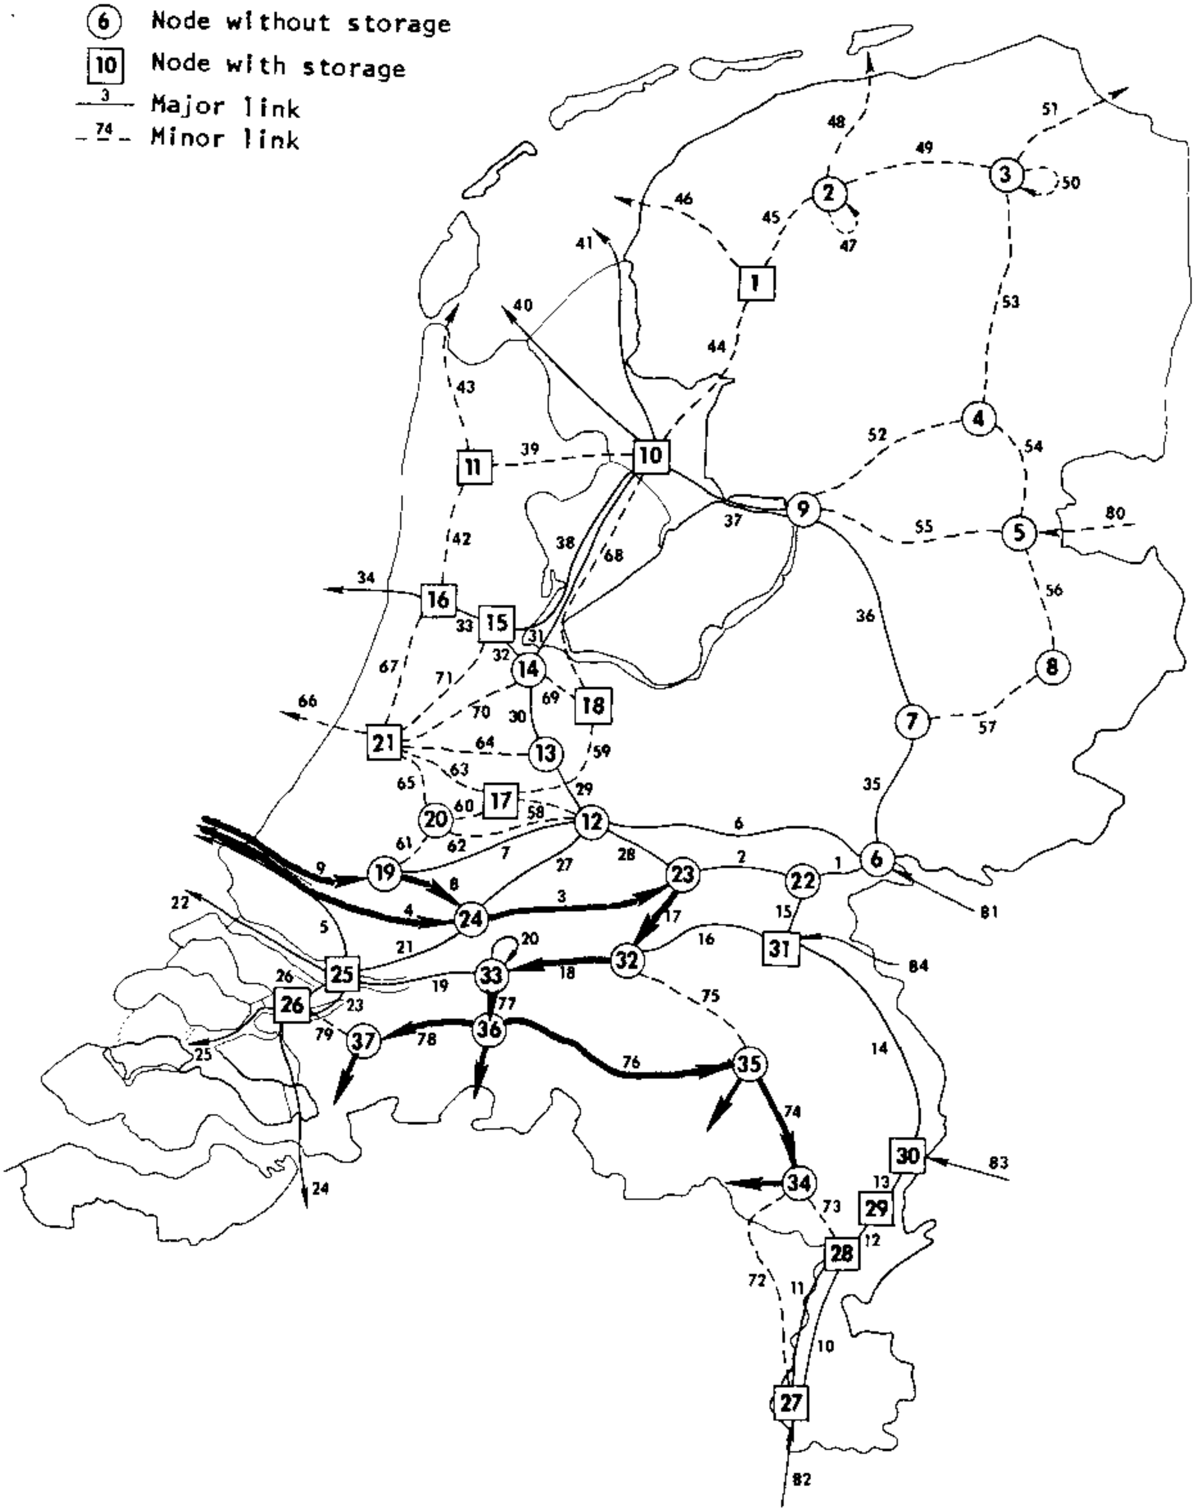

Fig. 2.6b--Degrees of freedom added by new infrastructure: add St. Andries Connection (link 17); add pumping on Wilhelminakanaal (link 76) and upper Zuid-Willemsvaart (link 74)

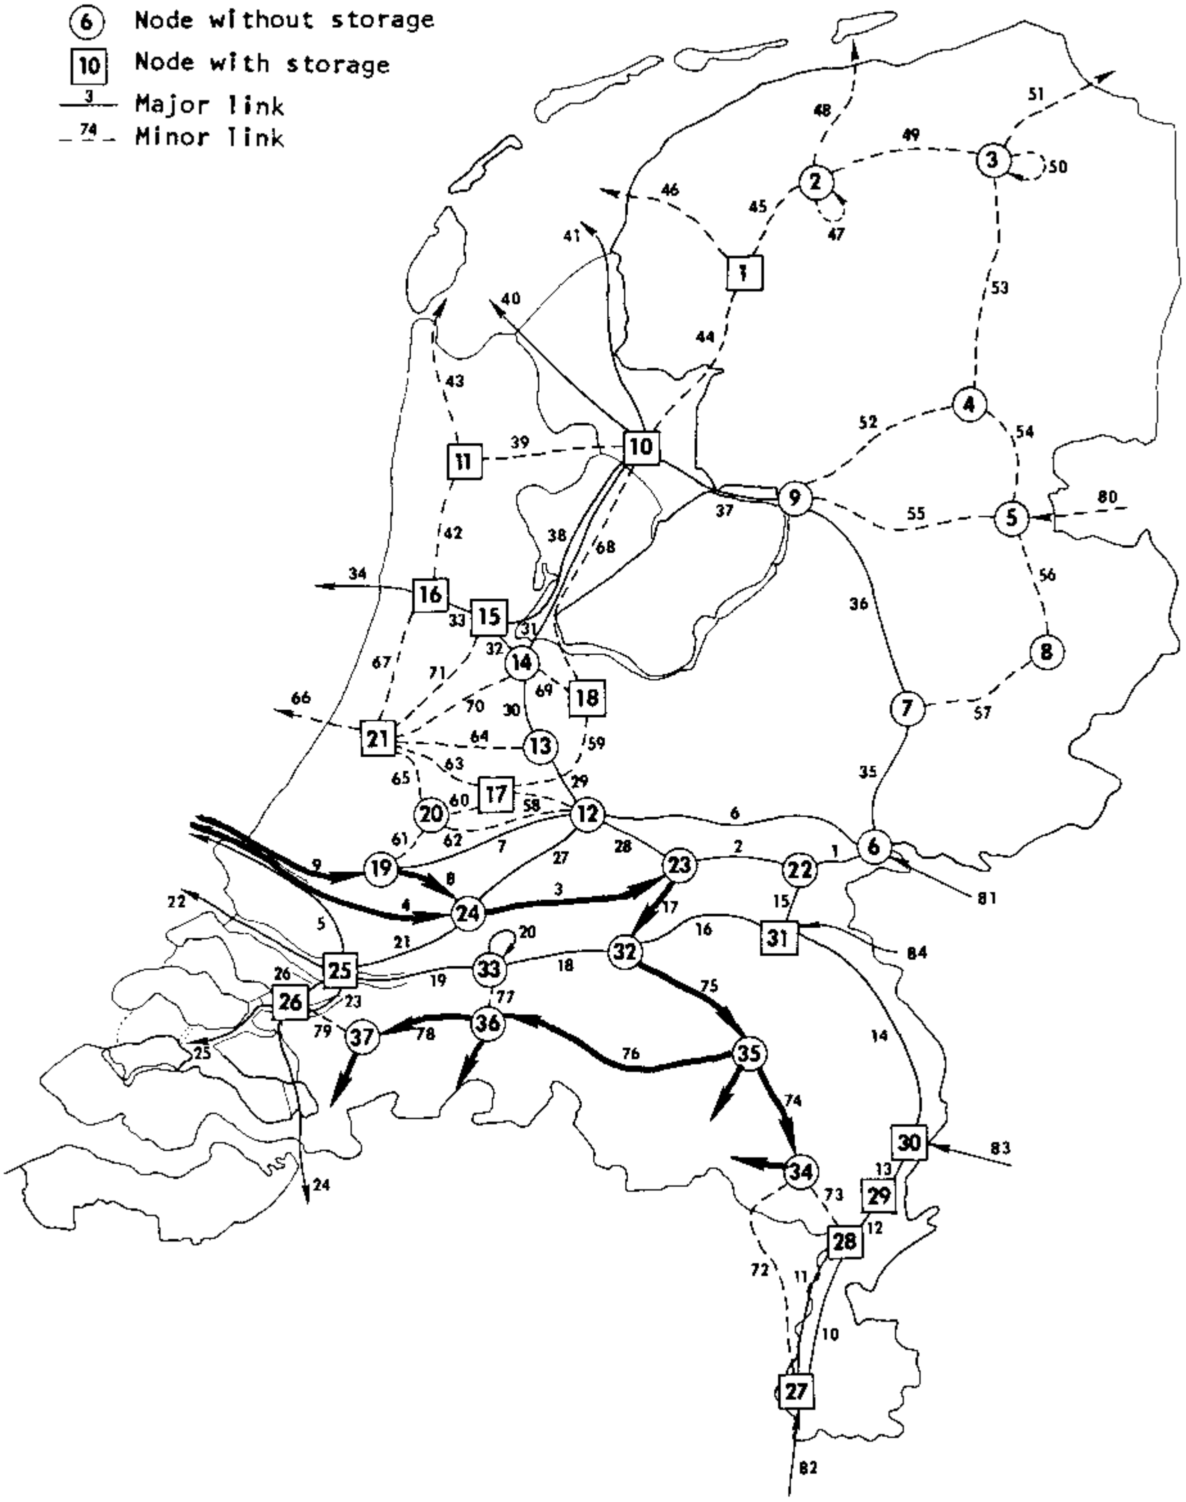

Fig. 2.6c--Degrees of freedom added by new infrastructure: add St. Andries Connection (link 17), add pumping on lower Zuid-Willemsvaart (link 75) and upper Zuid-Willemsvaart (link 74)

NOTES

1. According to Webster's New Collegiate Dictionary, 1975, hydraulics is "a branch of science that deals with practical applications (as the transmission of energy or the effects of flow) of liquid (as water) in motion," whereas hydrology is "a science dealing with the properties, distribution, and circulation of water on the surface of the land, in the soil and underlying rocks, and in the atmosphere." The main distinction seems to be that the term "hydraulics" and related terms (e.g., hydraulic, hydraulically) refer to man-caused effects on water, while terms such as hydrology, hydrological, and the like, refer to natural effects. The hydrological characteristics referred to here are the shapes, resistances to flow, and gradients of the different sections of the river channels. Although man has tried to influence these characteristics in selected locations (e.g., by dredging), they are for the most part determined naturally--hence the use here of the term "hydrological" rather than "hydraulic."
2. We do take into account changes in water levels when calculating shipping costs. For this purpose, however, we express the water level (and hence the depth) directly as a function of flow. We do not try to calculate the water level by estimating the difference between the flows into and out of a link, and determining what effect the consequent change in water volume will have on the water level. See Vol. IX for the derivation of the functions relating depth to flow in selected links.

REFERENCES

2.1. MW-197 (unpublished PAWN memorandum), "Cost to Industries of Short-term Cutbacks of Groundwater Intake," July 1979.

2.2. MW-127 (unpublished PAWN memorandum), "Screening of Pricing and Regulation Tactics on Shortage, Excess, and Salinity Issues," July 1978.

2.3. De Vries, J. J., "Groundwater Flow Systems and Stream Nets in the Netherlands," Rodopi, N. V., Amsterdam, 1974.

Chapter 3

WATER MANAGEMENT AND ENVIRONMENTAL QUALITY

3.1. INTRODUCTION

Water is a vital factor in determining the state of the environment. The quality of water in a lake or stream is one of the most important factors in determining what species of plant and animal life can thrive there. But water is almost equally important to terrestrial environments. The quantity and quality of water supplied to an area of land strongly influences the composition of the plant community growing there. The influence extends indirectly to animals, who ultimately depend on the plants for food and shelter.

Water quality also affects the uses that man may make of water. Drinking water, for example, must be free of toxic substances, and must meet certain standards of salinity, odor, and taste. Some industrial processes also require water to meet rigorous quality standards, especially water used in food processing and in certain chemical processes. Fisheries depend on the fish being healthy, abundant, and free of accumulated pollutants that would make the fish unsalable. Waterborne recreation is encouraged by clean, clear water, and discouraged where the water is turbid, or choked with aquatic weeds. If the water is sufficiently polluted, it may even become a hazard to human health.

Since changes in water quality evidently have such widespread implications, it behooves the water manager to investigate whether his actions may, inadvertently or not, affect water quality. However, it is impossible to consider all potential water quality and environmental issues in MSDM. The list is open-ended, and even that part of the list that could be readily written down is so long as to make a complete analysis hopeless. The best we can do is to choose a representative sample of water quality and environmental problems for consideration, and to hope that additional or alternative choices would not have revealed important phenomena that our actual choices have left hidden. Two measures of water quality, salinity and temperature, are treated elsewhere, in Chaps. 4 and 5, respectively. In this chapter, we discuss how various other aspects of water quality and the environment were chosen for, and are dealt with in MSDM.

Changes in the quantity of water can also have environmental impacts. We cannot speak here of making deserts bloom--there are no deserts in the Netherlands--but of changes in the environment at the border between the aquatic and the terrestrial. For example, a tactic such as increasing the water level of the IJsselmeer in summer will flood beaches, marshes, etc., that are presently available for recreation, or for feeding or nesting birds. As another example, increased extractions of groundwater will lower the groundwater levels in various locations. But there are many locations where the plant communities have evolved to accommodate very shallow groundwater tables, which would be replaced by other types of plant communities if the groundwater level were to drop permanently. In other locations, where the groundwater level might rise as a result of artificial infiltration, as has been proposed in the Veluwe region of the province Gelderland, the reverse might happen.

Changes of this kind have only a limited and temporary environmental effect unless they recur repeatedly over the long term. Since MSDM looks only at short term intervals, these changes are of limited significance in the MSDM context. Instead, the changes tend to be considerations that must accompany the decision to implement--or not to implement--one or another technical tactic. Accordingly, we shall omit discussion of them in this volume.

3.2. THREE TYPES OF WATER QUALITY PARAMETERS

Leaving aside salinity and thermal pollution, which we discuss in Chaps. 4 and 5, it is possible to distinguish three types of water quality changes, each with its own effects on aquatic environments. Different water quality parameters reflect each type of change. The three different types are oxygen-consuming substances, eutrophication, and toxic substances. In MSDM we have chosen to model one parameter from each category: BOD, an index of oxygen-consuming substances; phosphate, a nutrient that contributes to eutrophication; and chromium, a heavy metal with toxic properties. Other models in PAWN, discussed elsewhere, consider additional parameters. A detailed discussion of how we chose these parameters to study can be found in App. C of Vol. VA.

3.2.1. Oxygen-Consuming Substances

The first type of parameter involves water pollution by oxygen-consuming substances. These substances are not themselves toxic, and may indeed serve as a food source for aquatic organisms. But for the most part, these substances are degraded by bacterial action to inorganic molecules, in the process of which oxygen is consumed. If the oxygen is sufficiently depleted, fish may die, and bad odor and appearance may result.

There are two common indices for measuring the degree of pollution of this type. One is simply the amount of oxygen present in the water, sometimes expressed as a percentage of the amount that would be present at saturation. The shortcoming of this index is that it measures only the present state of the water, and not its potential for further oxygen depletion. The other index attempts to remedy this shortcoming. It is a direct measure of the amount of oxygen consumed by the bacterial digestion of substances in the water in a standard amount of time (usually five days) at a standard temperature (usually 20 degrees Celsius). This index is given the name BOD, with the relevant time and temperature given as a subscript and superscript, respectively.

In PAWN, we have modeled both dissolved oxygen and BOD. Dissolved oxygen, however, is only considered in lakes and reservoirs, and is not modeled at all in MSDM or the Distribution Model (DM). It is discussed in Vol. VI. The other index, BOD, is modeled in both DM and MSDM.

3.2.2. Eutrophication

The second type of parameter measures a class of effects called eutrophication. Technically, eutrophication is defined as the aging of bodies of water, and as such involves a constellation of effects including the buildup of sediment on a lake bottom, the accumulation of the inorganic nutrients needed for plant and algae growth, the growth of water plants and algae, and the appearance of fish and insects characteristic of the altered environment. As we use it here, the meaning of eutrophication will be confined to the accumulation of nutrients and the subsequent growth of algae in the body of water.

Algae require a number of inputs from their surroundings in order to grow. Nutrients, especially phosphate and nitrogen, and for some species of algae silicon, are necessary. Solar energy for photosynthesis is also required. In addition, the temperature must be in the proper range, carbon dioxide must be available, and the acidity of the water cannot be too high or too low. Any of these factors can set a limit on the growth of algae if it strays out of the proper range, but the factor most often singled out as potentially under mankind's control is phosphate. (We point out in Vol. III that this may be misleading, and that phosphate control programs, in the absence of complementary measures, could be expensive failures as algae control measures.)

Heavy growths of algae, called algae blooms, are objectionable for many reasons. They may lend a greenish or brownish tinge to the water, and cause bad odors. They sometimes form a scum on the water surface, thus rendering the water less attractive. The algae reputedly may be toxic to domestic animals. They certainly may clog filters and interfere with water treatment. And when the algae bloom dies off, its mass can be a major contribution to the BOD load of the water body, resulting in the problems of oxygen depletion mentioned above.

There are several indices of the degree of eutrophication. One of the most common is the concentration of chlorophyll in the water. This has the advantage that it is easy to measure, but it has two shortcomings. First, different species of algae contain different proportions of chlorophyll, so the chlorophyll concentration is a poor indicator of total biomass, or of the objectionable features of an algae bloom--such as the potential contribution to oxygen depletion problems--that are essentially proportional to biomass. Second, the chlorophyll concentration is a measure only of the present state of the water, and not of its potential to support large blooms. Indices that attempt to remedy the latter shortcoming are nutrient concentrations, usually phosphate and sometimes nitrogen.

In PAWN we have modeled both algae blooms and the nutrients phosphate, nitrogen, and silicon on which they depend. The behavior of algae and nutrients in lakes and reservoirs is discussed in Vols. III and IV. The transport of phosphate through the rivers and canals of the water management infrastructure is modeled in both DM and MSDM, and the similar transport of nitrogen in modeled only in DM.

3.2.3. Toxic Substances

The third type of water quality index concerns toxic substances. Potentially there are many of these. The Dutch, in an agreement with other countries, have created two lists of water pollutants, called the Black and Grey lists, respectively, depending on their level of toxicity. The Black List contains substances that are highly toxic, persist in the water for a substantial time, and accumulate in plants, marine life, or body tissues. At present, fifteen substances have been placed on the Black List, as shown in Table 3.1.

Table 3.1

BLACK LIST SUBSTANCES

Mercury and its compounds
Cadmium and its compounds
Dieldrin |
DDT |
Endrin > pesticides
Aldrin _|
Chlordane
Heptachlor
Heptachloroepoxide
Hexachlorabenzene
Hexachlorcyclohexane
Polychlorinated biphenols (PCB's)
Endosulphan
Hexachlorbutadiene
Pentachlorphenol

The Grey List contains substances from the same chemical families as the Black List substances, but for which no standards have been agreed upon based on toxicity, persistence, and bio-accumulation. Also on the Grey List are substances and categories of substances shown in Table 3.2.

Table 3.2

GREY LIST SUBSTANCES

Category	Substance
Metals	Zinc
	Copper
	Nickel
	Chromium
	Lead
	Selenium
	Arsenic
	Antimony
	Molybdenum
	Titanium
	Tin
	Barium
	Beryllium
	Boron
	Uranium
	Vanadium
	Cobalt
	Thallium
	Tellurium
	Silver
Biocides and derivatives	
Substances influencing the taste or smell of water used by man	
Toxic or persistent organic compounds of silicon	
Inorganic phosphorous compounds and elemental phosphorous	
Nonpersistent mineral oils and hydrocarbons	
Cyanides and flourides	
Substances affecting the oxygen balance	Ammonia Nitrates

In PAWN we have modeled one parameter, chromium, from this category. It appears in both DM and MSDM.

3.3. OVERVIEW OF PAWN'S WATER QUALITY METHODOLOGY

3.3.1. Two Separate Water Quality Methodology Components

For modeling purposes, it is convenient to deal separately with pollution in waterways (rivers and canals with running water) and water basins (lakes and reservoirs containing still water). The water quality phenomena occurring in waterways consist largely of the transport of pollutants from one location to another. In Dutch circumstances, relatively little happens to pollutants as they traverse a waterway. In particular, it appears unnecessary to model interactions among pollutants in waterways.

By contrast, virtually nothing but interactions occur to pollutants in water basins. Algae grow, incorporating nutrients and producing BOD. The BOD in turn consumes oxygen, which may be reduced thereby to problem levels. If oxygen is sufficiently reduced, nutrients may be made available from the bottom sediments.

Accordingly, our methodology for dealing with pollutant transport in waterways is implemented as part of one set of models, and our methodology for dealing with water quality problems in water basins is implemented in an entirely different set of models. The models that will deal with pollutant transport in waterways are MSDM and DM. The models that deal with water quality problems in water basins are described elsewhere in this series of reports (Vols. III and VI).

DM and MSDM deal with pollutant transport in waterways. They calculate the water flows in the links of the PAWN network, and water storage in those PAWN network nodes specified to have storage. Both models also calculate pollutant concentrations at the nodes of their respective networks, including those with storage (see App. B of Vol. VA). It is in these models that the interactions between water quality and water quantity are largely represented. The difference between the two models is that the DM considers the water management policy to be exogenously specified, and calculates its consequences. Thus, the DM will calculate how often water quality standards may be violated, and how much the farmers and other water users may be damaged by high salinity. By contrast, MSDM calculates the water management policy that will meet the standards as closely as possible, while minimizing the damage done to all users.

3.3.2. Problem Formulation

As explained in detail in App. B of Vol. VA, this simple description of pollutant transport through the network can be expressed mathematically as a system of simultaneous linear equations. The unknowns in the equations are the pollutant concentrations at each node. The coefficients of these unknown concentrations are calculated from the flows of water in the network links, and the amounts of water stored at the nodes. As explained in Chap. 8, whenever we make this calculation, we will know, at least provisionally, the flows of water in every link of the network. The right-hand sides, or constant terms in the equations, are the discharges of pollutant into each node of the network. The sources of these discharges, and their possibilities for control, are discussed in the next section. Mathematically, then, the calculation of pollutant concentrations in MSDM can be expressed as:

$$(3.1)\quad \begin{cases} \text{MATRIX} * \text{CONC} = \text{DISCH} \\ \qquad\quad \text{CONC} \le \text{STD} \end{cases}$$

In Problem (3.1), DISCH is a vector of pollutant discharges, one corresponding to each node. CONC is a vector of pollutant concentrations, and STD a vector of water quality standards, each vector having one element for each network node.

MATRIX is the matrix of coefficients in the simultaneous linear equations that relate the pollutant concentrations to the discharges. As mentioned above, the individual entries in the matrix depend on the water flows in network links, and water storage at nodes. Details of how these entries are determined can be found in App. B of Vol. VA.

3.3.3. Remedy for Overdetermination

One important feature of Eq. (3.1) is that it is overdetermined. That is, if we ignore the existence of water quality standards, the remaining equations are sufficient to specify all the concentrations. This means that, if we specify standards that are insufficiently lenient, Eq. (3.1) will have no solution at all.

In DM, overdetermination presents no real problem. Whenever a water quality standard is violated at a node, the fact is merely noted. No effort is made to adjust the managerial strategy in order to eliminate the violation. One of the outputs of DM, in fact, is the frequency with which standards are violated. But in MSDM we wish to devise, if possible, a strategy that will meet the standards, and so a different action must be taken. Thus, we modify the problem, and instead of demanding that the concentration at each node satisfy its standard, we allow the standard to be violated, but we exact a penalty for the privilege. Recalling Sec. 2.7, this penalty offers a way to represent water flow requirements imposed by water quality standards in the water distribution problem of MSDM. The problem MSDM solves for each pollutant is:

$$
(3.2)\quad \begin{cases} \text{Minimize:} & \sum_{j=1}^{n} [\max\{CONC_j - STD_j, 0\}] \\ \\ \text{Subject to:} & MATRIX * CONC = DISCH \end{cases}
$$

In Problem (3.2), the subscript "j" indicates which node is referred to. The total number of nodes is "n," so the summation covers all nodes. For each node, the term appearing in the function to be minimized (the objective function) is the excess of the concentration over its standard.

3.3.4. Are We Studying the Right Problem?

At this point, the reader may question whether we are solving the proper problem. For example, a high BOD level is not by itself objectionable; only if other conditions are suitable will a high BOD level cause the oxygen to be depleted. Similarly, the fact that phosphate or nitrogen is abundant does not insure that algae blooms will be large. Even in the case of heavy metals, which are toxic in their own right, interactions with the sediments and the biota make the concentration in water a poor measure of the harm done.

Indeed, more sophisticated models of the effects of these pollutants would take other factors into account. Our models of eutrophication in lakes (see Vol. VI) in fact do so. A model that would do the same for chromium could in principle be built, although the practical difficulties are enormous.

However, the pollutant concentration is the only simple, quick measure available of the degree of pollution. Any measure taken to improve water quality will first produce an effect on the concentration, and only later will the more sophisticated measures of pollution reflect the change. As a practical matter, therefore, water quality must be measured in terms of pollutant concentrations (although this does not preclude the use of other measures in parallel).

In addition, the pollutant concentration is often a good measure of the risk of eventual harm. A high concentration need not of itself cause damage, but a high concentration in conjunction with other conditions can. Eventually, those other conditions will occur, although nobody can predict when. Thus, to avoid the harm it may be necessary to keep the concentration low at all times. Indeed, the standards for pollutant concentrations are set with this in mind.

3.4. WATER QUALITY STANDARDS

The water quality standards we have selected for use in MSDM derive from actions taken under the Dutch "Pollution of Surface Waters Act of 1970." This act created the Commissie Uitvoering Wet Verontreiniging Oppervlaktewateren (Committee for the Execution of the Pollution of Surface Waters Act, CUWVO) to formulate an updated plan each five years for maintaining and improving surface water quality. This plan, called the Indicatief Meerjarenprogramma (Prospective Multi-Annual Program, IMP) established water quality standards, requirements for permits for pollutant discharges, etc. The first plan was published in 1975 [3.1]; a new plan was published during the preparation of this volume. The 1975 IMP, which we call the "old IMP," contains two sets of water quality standards. One places provisional limits on various water quality parameters, and the other establishes target values for the same parameters. The provisional limits are uniformly more lenient than the target values.

We have chosen to express pollutants in the following units. For BOD and phosphate, concentration units will be milligrams per liter, abbreviated mg/l. For chromium, because its concentration is so much lower, we use micrograms per liter, abbreviated ug/l. Milligrams per liter is the same as parts per million (ppm) by weight; micrograms per liter is the same as parts per billion (ppb).

Other water quality standards could have been chosen. In App. C of Vol. VA we present some alternative standards, as well as discussing reasons for the standards. Here we mention only one other possible source of standards, RIN.

The Rijksinstituut voor Natuurbeheer (State Institute for Nature Management, RIN) has the task of advising the Dutch Government on measures for managing and preserving nature areas--i.e., areas presently subject to minimal human influence. To that end they have proposed water quality standards expressly intended to ensure environmental preservation [3.2]. For BOD, they recommend a standard of 5 mg/l, the same as the old IMP provisional limit. For total phosphate, their standard is 0.05 mg/l, the same as the old IMP target value. But for chromium, they recommend a standard of zero.

RIN's standard of zero for chromium reflects their opinion that regardless of how low the concentration of chromium (or of heavy metals generally) becomes, there is still some benefit to be realized by reducing it further. That is, RIN is recommending that we minimize the concentration of chromium.

In general, it is not reasonable to minimize the concentration of chromium, or of any other pollutant, without regard for the costs this may impose on the various water uses or users (industry, households, etc.). Minimizing the chromium concentration must be interpreted as including a term in the objective that represents the damage done by the presence of chromium at a range of different concentrations. Other terms in the objective would represent costs of reducing the chromium concentration, either by treatment or by reducing or eliminating some activities of industry, shipping, recreation, housholds, etc. But we have no damage function for chromium, nor, indeed, for any other pollutant. It is for this reason that we rely on standards such as those presented in Table 3.3, standards which, in the judgment of the parties involved in setting them, constitute a reasonable balance between the damage done by pollutants and the costs of improving water quality.

Table 3.3

OLD IMP WATER QUALITY STANDARDS

	Parameter		
Source of Standards	Cr (ug/l)	BOD (mg/l)	Tot-P (mg/l)
Old IMP, provisional limits	50.0	5.0	0.3
Old IMP, target values	--	3.0	0.05

However, it is of some interest to attempt to minimize the concentrations of a pollutant at the various nodes. At the least it indicates what managerial tactics may improve water quality, and how these actions will affect other water users and uses. We have included several computer runs of MSDM of this kind in our study, and we report on them in Chap. 9, where we begin discussing results.

3.5. POLLUTANT SOURCES AND POSSIBILITIES FOR THEIR CONTROL

This section summarizes the sources of chromium, BOD, phosphate, and nitrogen pollution, and how they may be controlled. We consider pollutants brought into the Netherlands in border-crossing rivers separately from pollutants discharged into waterways inside the country, because of the different possibilities for control. A more complete discussion may be found in App. C of Vol. VA.

We develop two scenarios for both external and internal sources of pollution. One reflects the situation of 1976, and the other a future situation, which we suppose will occur at approximately 1985. In 1976, the Netherlands was in the midst of a program to construct sewage treatment plants throughout the country. By the mid-1980s, it is planned that this program will be complete, and in excess of 90 percent of all municipal and industrial discharges will be treated. At that time, the remaining untreated discharges will be so small and scattered that it will not be worthwhile attempting to treat them.

In addition, the nations along the Rijn, notably Germany, are taking their own steps to reduce pollutant discharges, which will reduce the amount of pollutant brought into the Netherlands via the Rijn. The 1985 scenario anticipates the effect that these programs may eventually have.

3.5.1. Pollutants in Border-Crossing Rivers

In DM and MSDM, we consider six border-crossing rivers. They are the Rijn, which is the most important, followed by the Maas. There are three small rivers flowing into the Maas from Germany, namely the Roer, the Swalm, and the Niers. Finally, the Overijsselsche Vecht flows into the Northeast Highlands from Germany.

In the 1976 scenario we assume that the concentrations of chromium, BOD, total phosphate are as they were observed to be in 1976 in each of the border-crossing rivers. In order to avoid much of the variability in concentrations that exists in the actual measurements, we averaged the measurements by month, and assumed that the pollutant concentrations were constant throughout each month. The concentrations used in the 1976 scenario, month by month, appear in Tables 3.4 through 3.6. Note that for the rivers in which chromium was not measured, we assumed there was zero chromium, except that in the Roer we assumed 5 ug/l. Note also that we assumed the pollutant concentrations in the Swalm to be the same as the concentrations in the Niers. These assumptions are discussed in App. C of Vol. VA.

In the 1985 scenario we assumed that reductions in pollutant discharges outside the Netherlands would reduce pollutant concentrations in border-crossing rivers whenever those concentrations had exceeded certain maxima in 1976. For chromium the maximum was 10 ug/l, so the 1985 scenario for chromium is the same as the 1976 scenario in any month and any river in which the 1976 concentration is less than 10 ug/l. But in any month and river for which the 1976 concentration exceeds 10 ug/l, the corresponding concentration in the 1985 scenario is set equal to 10 ug/l. For the other pollutants, the maxima are: BOD, 10 mg/l; and total phosphate, 0.3 mg/l (see App. C of Vol. VA for a discussion of these maxima). Tables 3.7 through 3.9 present the concentrations used in the 1985 scenario.

These two scenarios are clearly keyed to the year 1976, and to use them in conjunction with river flows from other years is clearly a suspect procedure. Our reasons for relying on these two scenarios alone are lack of data relating to other years of interest, and lack of a suitable procedure for generating more reasonable, though artificial, scenarios.[1] Some further discussion may be found in App. C of Vol. VA.

Table 3.4

CHROMIUM CONCENTRATIONS (MICROGRAMS/LITER) IN BORDER-CROSSING RIVERS: 1976 SCENARIO

Month	Rijn at Lobith	Over-IJsselsche Vecht	Maas at Eijsden	Roer at Roermond	Swalm	Niers at Gennep
Jan	36.0	0.0	28.0	5.0	0.0	0.0
Feb	43.0	0.0	17.0	5.0	0.0	0.0
Mar	65.5	0.0	7.0	5.0	0.0	0.0
Apr	71.0	0.0	8.0	5.0	0.0	0.0
May	55.0	0.0	9.0	5.0	0.0	0.0
Jun	43.0	0.0	5.0	5.0	0.0	0.0
Jul	55.66	0.0	1.0	5.0	0.0	0.0
Aug	45.0	0.0	3.0	5.0	0.0	0.0
Sep	61.0	0.0	2.0	5.0	0.0	0.0
Oct	51.5	0.0	6.0	5.0	0.0	0.0
Nov	57.5	0.0	5.0	5.0	0.0	0.0
Dec	52.0	0.0	10.0	5.0	0.0	0.0

Table 3.5

BOD CONCENTRATIONS (MILLIGRAMS/LITER) IN BORDER-CROSSING RIVERS: 1976 SCENARIO

Month	Rijn at Lobith	Over-IJsselsche Vecht	Maas at Eijsden	Roer at Roermond	Swalm	Niers at Gennep
Jan	6.31	14.50	2.72	10.79	12.89	12.89
Feb	7.21	20.33	5.09	15.84	14.24	14.24
Mar	9.11	15.00	3.89	10.0	14.99	14.99
Apr	10.35	6.75	5.69	9.59	30.49	30.49
May	12.39	15.25	4.84	11.99	13.34	13.34
Jun	9.76	7.50	7.57	10.89	23.76	23.76
Jul	14.47	7.38	9.34	7.64	36.49	36.49
Aug	13.19	7.33	7.01	5.09	66.99	66.99
Sep	11.64	5.00	4.59	10.19	73.99	73.99
Oct	9.51	5.25	3.84	7.49	21.29	21.29
Nov	12.49	9.25	4.47	8.99	36.49	36.49
Dec	7.14	0.0	4.72	20.99	21.32	21.32

Table 3.6

TOTAL PHOSPHATE CONCENTRATIONS (MILLIGRAMS/LITER) IN BORDER-CROSSING RIVERS: 1976 SCENARIO

Month	Rijn at Lobith	Over-IJsselsche Vecht	Maas at Eijsden	Roer at Roermond	Swalm	Niers at Gennep
Jan	0.79	0.80	0.45	3.0	1.99	1.99
Feb	0.74	1.34	0.88	2.64	1.58	1.58
Mar	0.96	0.82	0.47	2.28	1.59	1.59
Apr	1.06	0.35	0.49	2.30	2.86	2.86
May	1.05	0.59	0.68	1.99	2.99	2.99
Jun	0.86	0.34	0.97	2.21	1.69	1.69
Jul	1.02	0.34	0.90	2.49	2.14	2.14
Aug	0.80	0.29	1.16	2.19	2.49	2.49
Sep	0.99	0.19	1.62	2.21	5.34	5.34
Oct	1.03	0.60	1.40	2.65	3.97	3.97
Nov	1.15	1.11	1.26	2.04	3.55	3.55
Dec	0.92	0.0	0.74	2.66	2.22	2.22

Table 3.7

CHROMIUM CONCENTRATIONS (MICROGRAMS/LITER) IN BORDER-CROSSING RIVERS: 1985 SCENARIO

Month	Rijn at Lobith	Over-IJsselsche Vecht	Maas at Eijsden	Roer at Roermond	Swalm	Niers at Gennep
Jan	10.0	0.0	10.0	5.0	0.0	0.0
Feb	10.0	0.0	10.0	5.0	0.0	0.0
Mar	10.0	0.0	7.0	5.0	0.0	0.0
Apr	10.0	0.0	8.0	5.0	0.0	0.0
May	10.0	0.0	9.0	5.0	0.0	0.0
Jun	10.0	0.0	5.0	5.0	0.0	0.0
Jul	10.0	0.0	1.0	5.0	0.0	0.0
Aug	10.0	0.0	3.0	5.0	0.0	0.0
Sep	10.0	0.0	2.0	5.0	0.0	0.0
Oct	10.0	0.0	6.0	5.0	0.0	0.0
Nov	10.0	0.0	5.0	5.0	0.0	0.0
Dec	10.0	0.0	10.0	5.0	0.0	0.0

Table 3.8

BOD CONCENTRATIONS (MILLIGRAMS/LITER) IN BORDER-CROSSING RIVERS: 1985 SCENARIO

Month	Rijn at Lobith	Over-IJsselsche Vecht	Maas at Eijsden	Roer at Roermond	Swalm	Niers at Gennep
Jan	6.31	10.0	2.72	10.0	10.0	10.0
Feb	7.21	10.0	5.09	10.0	10.0	10.0
Mar	9.11	10.0	3.89	10.0	10.0	10.0
Apr	10.0	6.75	5.69	9.59	10.0	10.0
May	10.0	10.0	4.84	10.0	10.0	10.0
Jun	9.76	7.50	7.57	10.0	10.0	10.0
Jul	10.0	7.38	9.34	7.64	10.0	10.0
Aug	10.0	7.33	7.01	5.09	10.0	10.0
Sep	10.0	5.00	4.59	10.0	10.0	10.0
Oct	9.51	5.25	3.84	7.49	10.0	10.0
Nov	10.0	9.25	4.47	8.99	10.0	10.0
Dec	7.14	0.0	4.72	10.0	10.0	10.0

Table 3.9

TOTAL PHOSPHATE CONCENTRATIONS (MILLIGRAMS/LITER) IN BORDER-CROSSING RIVERS: 1985 SCENARIO

Month	Rijn at Lobith	Over-IJsselsche Vecht	Maas at Eijsden	Roer at Roermond	Swalm	Niers at Gennep
Jan	0.3	0.3	0.3	0.3	0.3	0.3
Feb	0.3	0.3	0.3	0.3	0.3	0.3
Mar	0.3	0.3	0.3	0.3	0.3	0.3
Apr	0.3	0.3	0.3	0.3	0.3	0.3
May	0.3	0.3	0.3	0.3	0.3	0.3
Jun	0.3	0.3	0.3	0.3	0.3	0.3
Jul	0.3	0.3	0.3	0.3	0.3	0.3
Aug	0.3	0.29	0.3	0.3	0.3	0.3
Sep	0.3	0.19	0.3	0.3	0.3	0.3
Oct	0.3	0.3	0.3	0.3	0.3	0.3
Nov	0.3	0.3	0.3	0.3	0.3	0.3
Dec	0.3	0.3	0.3	0.3	0.3	0.3

3.5.2. Pollutant Discharges in the Netherlands

The source of our data on pollutant discharges inside the Netherlands is Ref. 3.3. This document was compiled by a team from RIZA, DHL, and WW from preliminary returns from the new IMP survey on existing and planned wastewater treatment plants, and annual reports and water quality plans of the different regional water quality boards. In Ref. 3.3 there is given a line of data for each discharge source.

Entries in each line identify the DM node or PAWN district into which the discharge takes place, a name of the source, and the annual discharge for 1976 and 1985. For each node and district, we accumulated the discharges assigned thereto.

The discharges into DM nodes were assigned to MSDM nodes using the correspondence between nodes in the DM network and nodes in the MSDM network (see App. A of Vol. VA). Some DM nodes correspond to no MSDM nodes. We ignore these discharges in MSDM because, with two exceptions, the water and pollutants entering these DM nodes is washed directly into the Waddenzee or the North Sea. The exceptions (DM nodes 83 BETUWE and 84 TIELWARD) receive such small pollutant discharges and have such low pollutant concentrations that no significant error is introduced by ignoring their pollutant discharges.

In DM, surface water in the districts is explicitly represented, so that pollutants can be discharged directly into the district waters, and only make their way into the network when the districts drain excess water. In MSDM, however, the districts are not explicitly represented. We have chosen, therefore, to discharge the district pollutants directly into MSDM nodes.

We assign district discharges to MSDM nodes as follows. First, for each district we have specified to which DM nodes, and in what proportions, that district will discharge its drainage water. We apportion the pollutant discharges into districts among the DM nodes as specified by those drainage proportions. Any proportion of the drainage from a district that is not assigned to a DM node results in that proportion of the district pollutant discharges remaining unassigned to any MSDM node. We assign the resulting pollutant discharges into DM nodes to the MSDM nodes by the same rule discussed above. The results of these assignments are shown for the two scenarios in Tables 3.10 and 3.11.

3.5.3. Control of Pollutant Discharges

Generally, pollutant discharges can only be controlled by changes in industrial processes or by the construction of various kinds of treatment facilities. These tactics require long-range planning and investments. In fact, we have discovered only two tactics that might be used to reduce pollutant discharges at a few days' notice, as is

Table 3.10

POLLUTANT DISCHARGES INTO MSDM NODES: 1976 SCENARIO

	Node	Chromium (g/s)	BOD (kg/s)	Total Phosphate (kg/s)
1.	FRIELAND	0.0634	0.4109	0.0409
2.	B'GUMMER	0.0	0.0	0.0
3.	GRONETAL	0.1744	0.6417	0.0546
4.	NEHIGH	0.0	0.1821	0.0246
5.	OVIJVECH	0.1237	0.1093	0.0245
6.	UPRIVER	0.0332	0.2067	0.0199
7.	TWENMOND	0.1190	0.2195	0.0255
8.	TWENTEND	0.0095	0.0114	0.0022
9.	IJSLEND	0.0	0.2061	0.0173
10.	IJSLAKES	0.3012	0.9469	0.0955
11.	NORHOLL	0.0	0.3670	0.0307
12.	A-R.MOND	0.0	0.1714	0.0131
13.	UTR+MAAR	0.0	0.0078	0.0027
14.	DIEMEN	0.0	0.0097	0.0031
15.	AMSTEDAM	0.0	0.0843	0.0044
16.	HAL+IJMU	0.0190	0.2989	0.0274
17.	LOPIKWAR	0.0	0.0149	0.0014
18.	VECHT	0.0634	0.2338	0.0308
19.	IJSLMOND	0.0	2.1836	0.0416
20.	GOUDA	0.0159	0.0401	0.0064
21.	MIDWEST	0.0507	0.3805	0.0405
22.	NIJMEGEN	0.0	0.3990	0.0281
23.	TIEL	0.0	0.2456	0.0178
24.	GOR+DOR	0.0	0.2332	0.0208
25.	LOWRIVER	0.0	0.0917	0.0079
26.	DLTALAKE	0.0	0.0059	0.0011
27.	MAASLOO	0.0	0.0303	0.0040
28.	BORN+PAN	0.0	0.1468	0.0189
29.	LINNE	0.0	0.2107	0.0123
30.	ROER+BEL	0.0	0.1820	0.0221
31.	SAMGRALI	0.0	0.3040	0.0231
32.	DBOS+BOX	0.1140	0.1745	0.0157
33.	GERTRUID	0.0412	0.0340	0.0021
34.	WEER+MEY	0.0	0.0	0.0
35.	HELMOND	0.0	0.0	0.0
36.	OSTRHOUT	0.1333	0.0422	0.0064
37.	FIJNAART	0.0	0.0071	0.0037
38.	(unassigned)	0.0	3.0430	0.0888

Table 3.11

POLLUTANT DISCHARGES INTO MSDM NODES: 1985 SCENARIO

	Node	Chronium (g/s)	BOD (kg/s)	Total Phosphate (kg/s)
1.	FRIELAND	0.0	0.3809	0.0414
2.	B'GUMMER	0.0	0.0	0.0
3.	GRONETAL	0.0	0.1181	0.0247
4.	NEHIGH	0.0	0.1017	0.0182
5.	OVIJVECH	0.0	0.0917	0.0235
6.	UPRIVER	0.0	0.0854	0.0219
7.	TWENMOND	0.0	0.1695	0.0237
8.	TWENTEND	0.0	0.0114	0.0022
9.	IJSLEND	0.0	0.0692	0.0117
10.	IJSLAKES	0.0	0.7020	0.0835
11.	NORHOLL	0.0	0.1068	0.0146
12.	A-R.MOND	0.0	0.1071	0.0112
13.	UTR+MAAR	0.0	0.0029	0.0024
14.	DIEMEN	0.0	0.0051	0.0030
15.	AMSTEDAM	0.0	0.0487	0.0025
16.	HAL+IJMU	0.0	0.1978	0.0229
17.	LOPIKWAR	0.0	0.0105	0.0013
18.	VECHT	0.0	0.1086	0.0270
19.	IJSLMOND	0.0	2.0214	0.0354
20.	GOUDA	0.0	0.0236	0.0059
21.	MIDWEST	0.0	0.3121	0.0375
22.	NIJMEGEN	0.0	0.0508	0.0111
23.	TIEL	0.0	0.0192	0.0040
24.	GOR+DOR	0.0	0.1165	0.0198
25.	LOWRIVER	0.0	0.0463	0.0060
26.	DLTALAKE	0.0	0.0035	0.0010
27.	MAASLOO	0.0	0.0107	0.0027
28.	BORN+PAN	0.0	0.0804	0.0202
29.	LINNE	0.0	0.0428	0.0082
30.	ROER+BEL	0.0	0.0386	0.0117
31.	SAMGRALI	0.0	0.0731	0.0161
32.	DBOS+BOX	0.0	0.1358	0.0141
33.	GERTRUID	0.0	0.0340	0.0021
34.	WEER+MEY	0.0	0.0	0.0
35.	HELMOND	0.0	0.0	0.0
36.	OSTRHOUT	0.0	0.0422	0.0064
37.	FIJNAART	0.0	0.0069	0.0035
38.	(unassigned)	0.0	2.6566	0.0775

required for a tactic to be considered in MSDM. One of these tactics is to let a treatment plant stand idle except when water quality standards are violated--a rather silly tactic, and better characterized as the tactic of increasing pollutant discharges when the standards are more than met. The other tactic is to sequester pollutants in a special holding pond when it is necessary to reduce discharges, only to release them later, when more water is available for dilution.

The second tactic is used to some extent in the Ruhr district of Germany [3.4]. There, the "holding ponds" are entire canalized rivers, including the Wupper, Ruhr, Emscher, Lippe, and Niers. In a nation with as little land as the Netherlands, it may not be feasible to provide holding ponds with significant capacity. In any case, such holding ponds do not now exist, and the PAWN project has not considered building them. Thus, MSDM possess no tactics that could reduce pollutant discharges within the Netherlands.

Pollutants in border-crossing rivers cannot be controlled unilaterally by the Dutch. Treaties and other international agreements are needed between the Dutch and the nations whose discharges pollute the border-crossing rivers. Even should such agreements be reached, the same objections regarding day-to-day control of discharges exists in countries other than the Netherlands as we have pointed out for the Netherlands itself.

Accordingly, discharges of chromium, BOD, and phosphate are taken to be entirely part of the scenario presented to MSDM. They are exogenously specified constants, in no degree capable of modification within MSDM.

The other major source of both BOD and phosphate is from internal loadings in lakes. Internal loading of BOD is due to the growth of algae. In fact, algae growth is the sole source of BOD problems in the Netherlands, except for some extremely localized situations. Internal loading of phosphate is from the release of phosphate from the bottom sediments of the lakes. Although it is true that under present circumstances the lake bottoms absorb more phosphate than they release, this could change if less phosphate were introduced into the lakes from outside. These internal loadings are discussed at length in Vol. VI.

3.6. CALIBRATION OF POLLUTANT TRANSPORT

3.6.1. The Calibration Process

DM and MSDM must be calibrated for chromium, BOD, and phosphate. Because the networks involved in the two models are so closely related, the calibration could be carried out in DM, and transferred with little change directly to MSDM. This section only summarizes the calibration process and results. They are reported in detail in App. C of Vol. VA.

Calibrating DM for a pollutant means determining a decay rate[2] of that pollutant in each network link and in each network node with storage. We provide as inputs the 1976 inventory of pollutant discharges at each node, and the DM provides estimates of flows and lake levels averaged over decades (ten-day periods). Then the best decay rates are those that lead to the closest decade-by-decade and node-by-node match of the observed and calculated 1976 pollutant concentrations.

No choice of decay rates will eliminate all differences between the observed and estimated pollutant concentrations. The observed concentrations exhibit a considerable amount of variation from one observation to the next. Further, no node was blessed with many measurements taken during any one decade. Thus, the measured decade averages had a considerable variation that appeared to us to be random. Whether this variation was due to errors in measurement, or nonuniformities in the concentration profile of a pollutant in the cross-section of a waterway, or a real difference in concentration due to unknown changes in pollutant discharges, we cannot say. Figure 3.1 shows an example of the variability in the concentration of BOD at Lobith on the Rijn, perhaps the most-measured pollutant at the most-measured point in the Netherlands.

In addition to this "random" deviation between estimated and observed concentrations, there are some systematic differences that cannot be eliminated by any reasonable choices for the decay rates. These systematic differences are probably due to errors in the inventories of pollutant sources.[3] When we have detected such errors, we have sought additional data, opinion, or speculation, that would justify altering the inventory. In App. C of Vol. VA, we discuss for what pollutants and for which nodes we have made such adjustments.

3.6.2. Available Data

For calibration we had data on pollutant concentrations from a variety of sources. Data on pollutants at 193 locations in state waters were obtained from RIZA, the Rijksinstituut voor Zuivering van Afvalwater (State Institute for Wastewater Treatment). At a few locations, measurements were taken weekly; at most locations they were biweekly. In regional waters--i.e., waters under the jurisdiction of provincial or lower-level governmental organizations--we have data at perhaps 100 locations measured monthly. Characteristics of these data are discussed at greater length in App. C of Vol. VA.

3.6.3. BOD and Phosphate Calibration in Links

Table 3.12 shows the decay rates we have estimated for total phosphate and BOD, based on pollution concentration and discharge data for 1976.

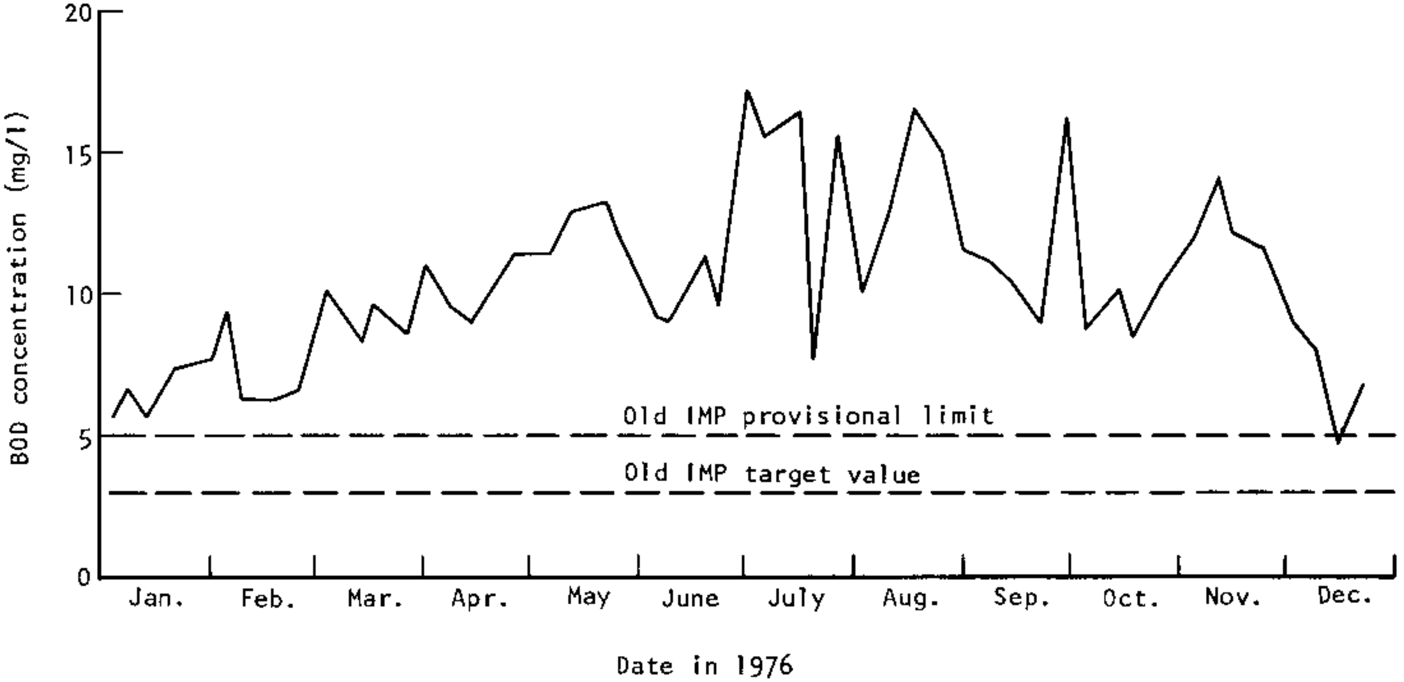

Fig. 3.1--BOD concentration in the Rijn at Lobith throughout 1976

Table 3.12

ESTIMATED DECAY RATES FOR BOD AND PHOSPHATE

	Fraction Lost per Day of	
Waterway	BOD	Phosphate
Waal	0.25	0.15
Neder-Rijn and Lek	0.10	0.15
Nieuwe Waterweg	0.50	0.15
Maas	0.10	0.10
Haringvliet, Hollandsch Diep and Zoommeer	0.10	0.10
Amsterdam-Rijnkanaal	0.10	0.10
Noordzeekanaal	0.10	0.15
IJssel rivers	0.20	0.15
IJsselmeer, Markermeer, Randmeren	0.10	0.10
Groningen (highlands)	0.05	0.0
Zwarte Meer	0.05	0.10
Drenthe	0.05	0.0
Midwest	0.10	0.0
Linge	0.20	0.0

Typical decay rates for BOD found in the literature range between 0.1 and 3.0 per day, with the usual rate lying between 0.1 and 0.6 per day [3.5, p. 97]. The rate depends on water temperature, type of BOD (carbonaceous BOD decays faster than nitrogenous BOD), and a multitude of other factors. Models of BOD decay frequently ignore such factors, often because the data on such factors is absent, and use a single average decay rate instead. For example, DHL's Rijn water quality model [3.6] uses a single rate of 0.3 per day.

From Table 3.12, we see that PAWN's estimates of the decay rates for BOD in waterways are within the generally accepted range. The estimates vary by waterway, the highest being 0.5 in the Rotterdam Waterway, followed by 0.25 in the Waal. The lowest values are 0.05 for numerous highland waterways.

In the minor waterways, data on concentrations and water flows are so sparse that for many links, no calibration was possible. Examples are the waterways in the low-lying parts of Friesland and Groningen, and in the southern highlands. In those links, we made the conservative assumption that no BOD decay occurs. This assumption is conservative in that it results in higher estimates of BOD concentration than more realistic assumed decay rates.

Our efforts at calibrating DM and MSDM for phosphate indicate that between ten and fifteen percent per day of the dissolved and suspended phosphate in waterways (rivers and canals) is lost. This loss is probably due to sedimentation of the fraction of phosphate that is adsorbed on particulate matter.

Loss rates in this range were applied to all major waterways. Again, data on concentrations and water flows in minor waterways were often so sparse that no calibration was possible. Again, therefore we made the conservative assumption that no losses of phosphate occurred in minor waterways.

3.6.4. BOD and Phosphate Calibration in Water Basins

In lakes and reservoirs, phosphate and BOD behavior is more complex. There, the growth and subsequent death and settling to the bottom of algae tends to strongly influence both the phosphate and BOD concentration in the water basin. For both of these pollutants, we found it necessary to introduce internal sources of pollutant in most of the water basins. The average rates at which these sources provide pollutants to the different water basins are given in Table 3.13.

The PAWN study of nutrients in lakes (Vol. VI) has demonstrated that a steady release of phosphate from the bottom occurs. In addition, conditions may occasionally occur which promote the massive and explosive release of phosphate from the bottom sediments. The existence of this source of phosphate was also suggested by early attempts to calibrate phosphate in the Distribution Model. Without

Table 3.13

INTERNAL RELEASE RATES OF BOD
AND PHOSPHATE IN WATER BASINS
($g/m^2/day$)

Water Basin	BOD	Phosphate
Haringvliet, Hollandsch Diep	0.0	0.09
Zoommeer	0.0	0.0
Amsterdam-Rijnkanaal	0.0	0.05
Noordzeekanaal	0.0	0.0
IJsselmeer	2.5	0.09
Markermeer	1.3	0.06
IJmeer	1.2	0.06
Veluwemeer	1.15	0.035
Gooimeer	1.85	0.13

such a source, it was necessary to assume decay rates of phosphate from lakes much lower that the decay rates from waterways, when a consideration of the processes believed to be involved (sedimentation of algae and of inorganic particulate matter) suggested that loss rates should be at least as high in the still waters of lakes. Therefore, we were forced to include a source of phosphate in each of the major lakes in order to model phosphate in DM and MSDM.

In many lakes, the major source of BOD appears to be due to the growth and subsequent death of algae in the lake itself, BOD from sewage contributing only a small fraction to the total BOD. This is confirmed by the PAWN investigation of dissolved oxygen in a number of Dutch lakes (Vol. VI), as well as by our own calibration attempts. Even assuming no decay of BOD in water basins, we were unable to reproduce the observed BOD concentrations in the more eutrophic lakes, especially those observed in the summer, when algae blooms are large. Accordingly, for each lake we assumed the BOD decay rate was the same as the decay rate in nearby waterways. Then we introduced an internal source of BOD which produced at a rate sufficient to elevate the estimated BOD concentrations approximately to the observed values. The internal BOD production was assumed to take place throughout the year in proportion to the available solar energy, as one would expect if it were due to production of the photosynthetic algae.[4]

3.6.5. Chromium Calibration

Such concentration and discharge data as we have for chromium indicate that a decay rate of five percent per day is suitable everywhere. In view of the uniformity of the estimates of chromium decay rates, we feel justified in assuming the same rate applies where data are exceptionally scarce or absent, notably on minor waterways. We feel constrained to point out, however, that assuming this decay rate merely makes our estimated chromium concentrations match the observed concentrations on the average. Considerable scatter remains between individual observations and the corresponding estimates.

NOTES

1. We gave some thought to the possibility that using pollutant loads from 1976 in other years would have been preferable to using pollutant concentrations (the load is calculated as the concentration times the river flow). But some analysis of the available data showed that the pollutant load has considerable random variation, although less than the pollutant concentration. We felt that using loads in place of concentrations would improve the analysis only marginally, and hence devoted our efforts to other activities with (we felt) higher payoff. Our simple data analysis is reported in App. C of Vol. VA.
2. The decay rates are intended to represent various mechanisms by which different pollutants are lost from the water over time. For example, heat exchanges with the air; bacteria digest BOD, converting it into nitrogen, phosphorous, etc., and consuming oxygen; and chromium and phosphorous are adsorbed on small particles which sink to the bottom. There is no analogous process for salt, and hence its decay rate is zero.
3. Some systematic differences could also be due to the simplicity of our approach. For example, observed concentrations in a small stream could be very high, due to a combination of low streamflow and high pollutant discharge. But if that stream were only one of many waterways represented by the same link in MSDM or DM, neither model would calculate such a large concentration. In cases such as this, we did not adjust the inventory of pollutant discharges; instead, we reconsidered which observed concentrations should be compared with the concentration our models calculated.
4. It is true, of course, that the more calibrating, adjusting, and general tinkering one must do to a model, the more circumspect one should be in drawing conclusions from the model's output. In the present case, the need to introduce internal sources of phosphate and BOD in lakes casts doubt on the predictive value of the model's calculated concentrations in circumstances other than those to which the model was calibrated; and this is true in spite of the fact that we have identified and verified that

processes do occur in lakes that are responsible for just such internal releases of pollutants (see Vol. VI). The simplicity of our model, its degree of aggregation, and its lack of any process other than first-order decay for the removal of pollutants, are other factors contributing to a need for caution in the drawing of conclusions, as are various shortcomings of the data. We believe, however, that we have been appropriately cautious in our use of the model's outputs, and that our conclusions regarding water quality are solid.

REFERENCES

3.1. *De Bestrijding van de Verontreiniging van het Oppervlaktewater: Indicatief Meerjarenprogramma 1975-1979* Ministerie van Veerkeer en Waterstaat (The Combat Against Surface Water Pollution in the Netherlands: Prospective Multi-annual Program 1975-1979) Ministry for Transport and Public Works, The Hague, 1975.
3.2. Unpublished information from the Rijksinstituut voor Naturbeheer (State Institute for Nature Management, RIN) reviewing water quality standards for the protection of the natural environment (PAWN file DW-472).
3.3. Unpublished information from various Dutch sources on pollutant discharges into Dutch waters (PAWN files DW-398 and DW-398A).
3.4. Fair, G. M., "Pollution Abatement in the Ruhr District," in *Comparisons in Resource Management*, H. Jarrett, ed., Johns Hopkins Press, Baltimore, Maryland, 1961.
3.5. Thomann, R. V., "Systems Analysis and Water Quality Management," *Environmental Research and Applications*, New York, 1971.
3.6. *Waterkwaliteitsmodel Rijn* (Water Quality Model of the Rhine), Delft Hydraulics Laboratory, Reports R 1056-II/S 321-I, R 1056-II/S 321-II, R 1056-II/S 321-III, 1977.

Chapter 4

SALT POLLUTION

4.1. INTRODUCTION

We have studied salinity in PAWN because it is a long-standing and serious problem in the Netherlands. Salt water from the sea invades the low-lying part of the Netherlands through estuaries, harbors and shipping locks, and also, driven by hydraulic pressure, through the subsoil. The water and salt budgets of these low-lying areas must be carefully controlled lest the salt concentration in the ditches and canals become too high to permit irrigation. This is especially a problem during rainless periods, and when fresh, clean river water is not available. But low-salinity river water has become less and less available as the Rijn has become more and more polluted. The salt load measured in the Rijn at Lobith has risen from an annual average of about 125 kg/s chloride in 1930 to about 300 kg/s chloride in 1976.

The salinity of the Rijn is aggravated during periods of low flow. In 1976, the average Rijn flow was only 1350 m^3/s, with an average chloride concentration of over 200 ppm. There were times during the year when the concentration exceeded 400 ppm. By contrast, in 1977 the chloride load was higher than in 1976 (about 350 kg/s), but because the average flow was also higher (2200 m^3/s), the average chloride concentration was lower (160 ppm) [4.1, 4.2].

A high salt concentration in the water supply causes damage to agriculture, industry, households, public health, and the environment. The adverse effects of salinity on agriculture takes many forms. Yields may be reduced in weight, or in the number of physical units of product, or in quality. For example, high quality lettuce grown in Dutch glasshouses may have small brown edges at some of the leaves rendering the product almost worthless.

The adverse effects of salt on industrial, commercial, and household use of water is due to the corrosive effect of saline water. Dissolved solids in water used for these purposes speed up the deterioration of equipment and furnishings (pipes, faucets, fabrics, appliances, etc), increasing user's costs by requiring more maintenance and cleaning, and earlier replacement.

At high concentrations, total dissolved solids (which include chloride) are objectionable in drinking water because of possible physiological effects (e.g., high blood pressure), unpalatable tastes, and unappealing appearance and odor.

In environmental terms, seawater intrusion into formerly nonsaline surface waters can damage the existing freshwater flora and fauna. However, most of the intrusion presently occurs in locations already "lost" from an environmental perspective (e.g., the Rotterdam Waterway and Noordzeekanaal). Salinity from the Rijn may have more long term and more widespread impacts.

It is interesting to note that in some areas a high salinity is considered desirable. Brackish-water ecological communities have developed along the shores of the Volkerak that environmentalists are anxious to preserve. These areas are threatened by the plans to dam off the Volkerak and the eastern end of the Oosterschelde, and convert them from the present, tidal basins to a freshwater lake (the Zoommeer).

No satisfactory way has been discovered to monetize the effects of salt on health or the environment. This makes it impossible to compare these impacts directly with the impacts on agriculture, industry, drinking water companies, and the like. In order to consider both the monetizable and nonmonetizable impacts, we have included the monetizable ones in our MSDM objective function, and imposed water quality standards to reflect the nonmonetizable ones.

In the remainder of this chapter, we will expand on the material presented in this first section. In Sec. 4.2, we will discuss our general approach to modeling salinity. In Secs. 4.3 through 4.6 we will deal with the various sources of salt in the Netherlands, including intrusion through locks, the Rotterdam salt wedge, and the Rijn. In Secs. 4.7 and 4.8 we will derive expressions for the losses suffered by agriculture and others as a result of highly saline waters. In Sec. 4.9, we discuss salinity standards for public health and environmental reasons. In the final section, 4.10, we estimate the value of water used to combat salt pollution in various ways, in order to compare these uses of water with others.

4.2. APPROACH TO MODELING SALT IN PAWN

Historically, the Dutch have measured salinity in terms of the chloride content of the water. This used to be more sensible than it is now. When the Rijn contained relatively little salt, the concentration of chloride was a good index of the concentrations of all salts in the water, because the source of the salts was the sea, whose salt composition is nearly constant. Now, however, the salt composition of water in the Netherlands is determined to an important degree by the salts in the Rijn, which are present in different ratios to chloride than in the sea. A measure such as conductivity, total hardness, or total dissolved solids, would be preferable to chloride, because it would be a better index of the potential for damage possessed by saline water [4.3]. However, most Dutch data are expressed in terms of chloride, and so we have been forced to use chloride as our measure of salinity in PAWN as well. However, we have chosen to refer to the problem as "salt pollution" rather than "chloride pollution."

We calculate the concentration of chloride at each node of the network in each decade. As explained in Chap. 8, whenever we make the calculation, we will know, at least provisionally, the flow in each link of the network. We assume that the salt in a link is carried passively downstream with the flow of water. When the water reaches a node, both the water and salt mix completely with any water and salt that may be stored there, or that may enter from another link, or that may be discharged into the network at that node from an external source. Water flowing out of the node carries salt with it at the concentration of the total mixture of inflows to the node.

As explained in detail in App. B of Vol. VA, this simple description of salt transport through the network can be expressed mathematically as a system of simultaneous linear equations. The unknowns in the equations are the salt concentrations at each node. The coefficients of these unknown concentrations are calculated from the (known) flows of water in the network links, and the amounts of water stored at the nodes. The right-hand sides, or constant terms in the equations are the discharges of salt into each node of the network. The sources of these discharges of salt, are the subjects of the next several sections.

Discharges of salt are not all specified as inputs to MSDM. Some may be changed by employing managerial tactics, and hence are calculated internally. These tactics will cost money, however, and MSDM will only choose to employ them to the extent that they reduce the monetary damage done by high concentrations of salt, or to the extent necessary to meet water quality standards. Indeed, when MSDM calculated salt concentrations at all the nodes, it does so by solving a minimization problem whose goal is to find the salt discharges, and corresponding salt concentrations, that meet the standards for the smallest sum of tactic costs (for reducing salt discharges) plus salinity damages. Mathematically we can write this problem as:

$$(4.1) \quad \begin{cases} \text{Minimize: (Cost of reducing DISCH) + (Damage from CONC)} \\ \\ \text{Subject to: } \text{MATRIX} * \text{CONC} = \text{DISCH} \\ \\ \qquad\qquad\qquad\quad \text{CONC} \le \text{STD} \end{cases}$$

In Problem (4.1), DISCH is a vector of salt discharges, one corresponding to each node. For any node at which no salt discharge occurs, the corresponding element of DISCH is zero. For any node at which there are no managerial tactics available for changing the rate of salt discharge, the corresponding element is a constant. But at nodes where such tactics exist, the corresponding element of DISCH is a function of the tactics employed, and by imputation a function of their costs.

CONC is a vector of salt concentrations, and STD a vector of salinity standards, each vector having one element for each network node.

MATRIX is the matrix of coefficients in the simultaneous linear equations that relate the salt concentrations to the discharges. As mentioned above, the individual entries in the matrix depend on the water flows in network links, and water storage at nodes. Details of how these entries are determined can be found in App. B of Vol. VA.

4.3. SALT INTRUSION THROUGH LOCKS

A significant amount of salt enters the Netherlands from the sea through shipping locks along the coast such as the locks at IJmuiden, Den Helder, locks in the Afsluitdijk, etc. Salt passing the locks at IJmuiden enters the Noordzeekaanal. There are several locks separating the Noordzeekaanal from smaller canals in the midwest and North Holland, through which Noordzeekaanal water can intrude. In Fig. 4.1, we show the points of seawater intrusion through the locks considered in MSDM.

It should be noted that the points at which salt intrusion takes place are determined to some extent by the infrastructure. In particular, the Dutch are presently engaged in damming off the eastern end of the Oosterschelde estuary, and the estuary's extension to the north (the Volkerak), to make this currently saltwater area into a freshwater lake called the Zoommeer. As long as the area remains salt, the Volkerak locks (at node 25 LOWRIVER) will be a salt intrusion point, but neither the Philipsdam locks nor the Kreekrak locks (at node 26 DLTALAKE) will allow salt to intrude into fresh water. Once the fresh Zoommeer has been constructed, however, the situation will be reversed. The Philipsdam and Kreekrak locks will be salt intrusion points, and the Volkerak locks will not.

4.3.1. Description of a Lock

A lock complex maintains two bodies of water separate, while allowing ships to pass from one to the other. The complex consists of a barrier (e.g., a dam or weir) that separates the two bodies of water, pierced by a lock that permits ships to pass from one side to the other. Lock complexes exist in the highlands, where they separate bodies of water at two different levels, and in the lowlands, where they separate bodies of water with both different levels and different salinities. In this chapter, we will discuss only the lowland (also referred to as salt-fresh) complexes. Highland (fresh-fresh) complexes are discussed in Chap. 7.

Figure 4.2 shows a diagram of a typical lock in two views, top and side. (From this point we will omit further mention of the barrier component of the lock complex, since it is through the lock that salt intrusion occurs.) Ships approaching the lock from either direction will moor while waiting to enter the lock. When the lock is empty and

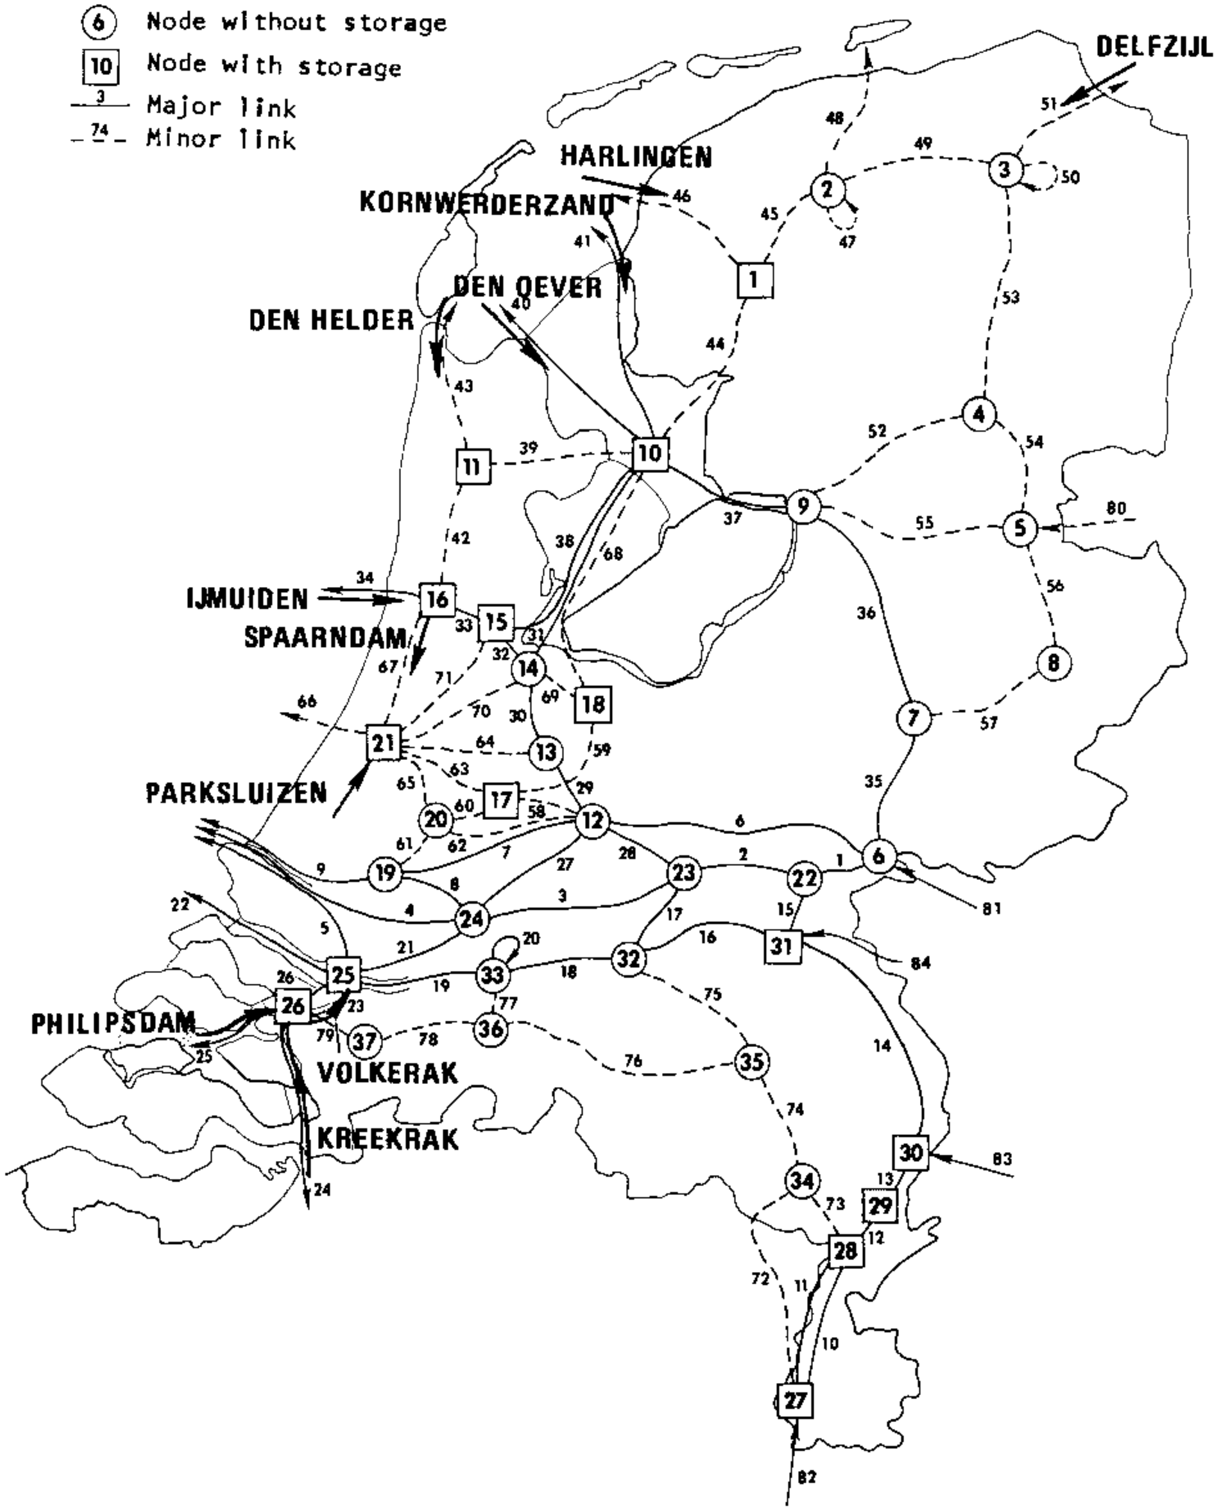

Fig. 4.1--Locks with salt intrusion

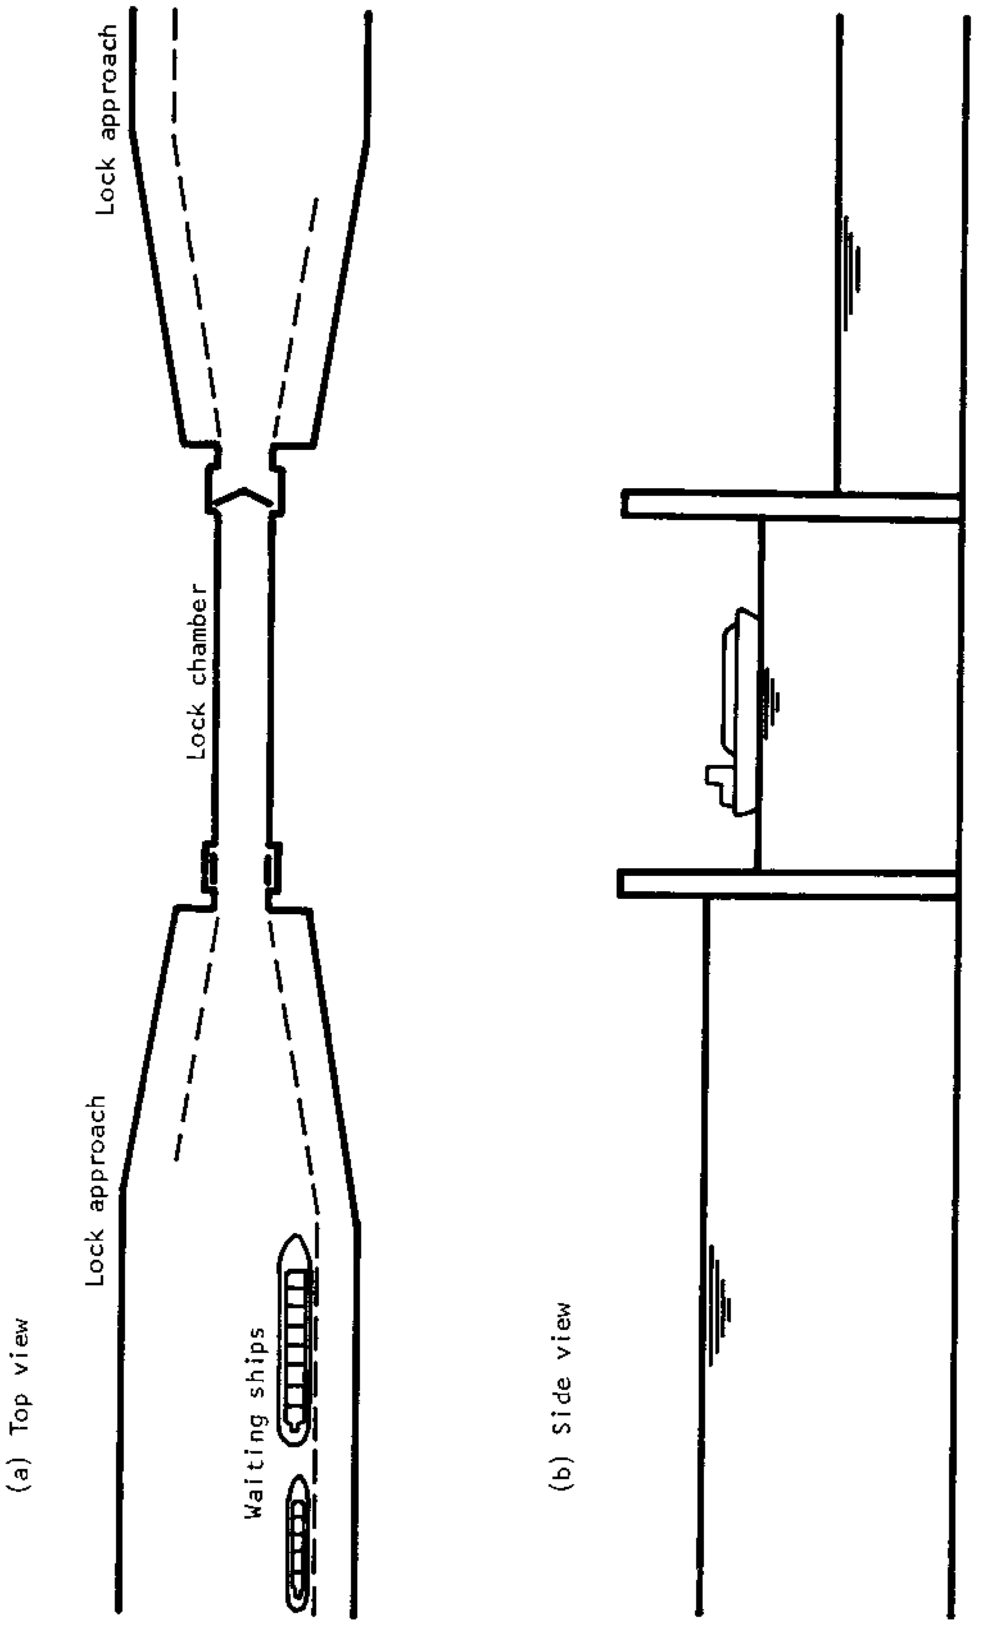

Fig. 4.2--Diagram of a lock

open, the ships enter and tie up inside the lock, until all ships have entered or the lock is full. The doors are then closed and the water level within the lock raised or lowered until it equals that of the waterway on the opposite side. When these levels have been equalized, the opposite doors are opened and the ships leave. The lock can then be cycled back in the other direction in the same manner.

4.3.2. Water Management Functions at Locks

Three water management functions have priority at salt-fresh locks: (1) preventing salt water from contaminating the fresh water, (2) isolating the fresh water from tidal influence, and (3) preventing fresh water from contaminating the salt water. The mere presence of the lock ensures that the fresh water will be tideless, however, and contamination of the salt water by fresh is a potential problem only at Volkerak and Philipsdam locks. Accordingly, in this volume we consider only the problem of salt intrusion.

Without additional measures, the mere presence of a lock does not prevent salt intrusion. Salt water will pass the lock for two reasons. First, when the lock opens to the salt side, salt water will mix with the fresh water in the lock chamber. When the lock subsequently opens to the fresh side, the salt water in the chamber can then contaminate the fresh water. Second, when the saltwater level is higher than the fresh (e.g., at high tide for salt-fresh locks located on the North Sea), salt water will be used to bring the levels to equilibrium during the locking process. The greater the level difference, the greater the amount of salt water that will enter the chamber. This water will mix in the chamber and be discharged to the fresh side during the next return cycle. How much salt enters the fresh water is a complex problem, depending on relative water levels, mixing characteristics, and other factors. As salt enters the fresh water, however, it diffuses back into the freshwater body, spreading the contamination.

4.3.3. Tactics for Reducing Salt Intrusion at Locks

At many salt-fresh locks the salt intrusion caused by this mixing process may be acceptable. The flow of fresh water out the lock and nearby sluice dilutes the salt and carries it back out to the sea. Salt concentrations in the fresh water away from the lock do not become objectionable.

But these concentrations are objectionable at other locks, as may be the loss of fresh water. Under these circumstances, additional measures are needed to further reduce the salt contamination and fresh water loss. These additional measures can be divided into two categories, technical and managerial. Technical tactics change the infrastructure or equipment at the lock, and managerial tactics change the operation of this equipment or of the lock itself.

We considered three technical tactics to be promising. A full description of these devices can be found in Vol. IX. Pneumatic barriers or bubble screens pump air into the water to create a continuous upward convection current at an entrance of the lock. The current creates a barrier that retards the mixing of salt and fresh water. Excavation and selective withdrawal systems use a pit excavated on the freshwater side of the lock. Salt water entering through the chamber, being more dense than fresh, will collect in the pit, from which it can be later pumped back out to the salt side of the lock. Kreekrak system locks are specially constructed to simultaneously pump the denser salt water out at the bottom and the lighter fresh water in at the top during the lock cycle. The fresh and salt water are kept separate through careful design and controlled inlet and outlet conditions. Unlike the first two technical tactics, the complex cycle of a Kreekrak-type lock requires more time than a normal lock, thus delaying ships.

It is also possible to manage the lock in such a way that salt intrusion is reduced. As described in detail in Vol. IX, we can classify these tactics into four general groups. First, operators may choose whether to use the salt intrusion technology installed at the lock. If the costs of using the technology exceed the costs of salt intrusion, the technology should be left idle. Second, the operator may flush the lock with fresh water, either through the lock or an adjacent sluice. This will reduce the amount of salt water that remains on the fresh side. Third, the operator may reduce the number of lock cycles, and hence reduce salt intrusion. One way this may be done is to require a specific number of ships (or fraction of lock capacity) in the lock before cycling. Fourth, the operator may reduce the salt intrusion per cycle. Some ways this may be done are: use the smallest lock at a complex, whenever safety permits; close lock doors whenever ships are not entering or leaving locks; and use intermediate doors in locks, whenever possible, to reduce chamber volume.

4.3.4. Representing Salt-Fresh Locks in MSDM

From the point of view of MSDM, there are four elements in the description of a salt-fresh lock. The first is the location of the lock, which we always take to be a node in the MSDM network. The other three are the rate of salt intrusion, the rate of water loss (also called the flushing rate), and the cost. Using the lock simulation model described in Vol. IX, one can establish two relationships. For some locks, we are interested in a relation that predicts the rate of salt intrusion as a function of the flushing rate and the cost, thus:

$$\text{SALT} = f_1(\text{FLUSH},\text{COST}) \qquad (4.2a)$$

where SALT is the salt intrusion rate in kg/s, FLUSH is the flushing rate in m^3/s, and COST is the cost of the salt intruson abatement tactics employed, in Dflm/dec, excluding the direct cost of flushing. COST includes both the operating cost of the technical tactics employed, if any, plus the cost of any delay incurred by shipping. SALT is one of the components of the vector DISCH that appears in Problem (4.1). The component of which it is a part is the component that discharges into the node at which the lock is located. The locks at which we use this relation are those at IJmuiden (node 16 HAL+IJMU), Den Oever and Kornwerderzand (node 10 IJSLAKES), the Volkerak locks (node 25 LOWRIVER), and the Philipsdam and Kreekrak locks (node 26 DLTALAKE). Salt intruding at IJmuiden is not confined to node 16 HAL+IJMU, but migrates upstream in the Noordzeekaanal as far as node 15 AMSTEDAM. Thus, two functions (4.2a) are associated with the locks at IJmuiden, one discharging salt into node 16 HAL+IJMU, the other into node 15 AMSTEDAM.

The other function we obtain from the lock simulation model predicts the incremental salt concentration at the lock, again as a function of the flushing rate and the cost. The incremental salt concentration is the difference between the concentration that is actually observed, including the effect of salt intrusion, and the concentration that would be observed if there were no salt intrusion (the background concentration). We can write this function as:

$$(4.2b) \quad \mathrm{CONC} = f_2(\mathrm{FLUSH},\mathrm{COST})$$

where CONC is the incremental salt concentration at the lock in mg/l of chloride. This relation is used for the Parksluizen, and for the locks at Spaarndam, Den Helder, Harlingen, and Delfzijl. Salt intruding at these locks is confined by even very modest flushing rates to very local areas, so it is inappropriate to discharge the intruding salt into a node from which it can spread widely through the network. Accordingly, we treat these locks differently from those mentioned earlier. Further discussion may be found in Sec. 4.7.3 below.

Functions (4.2a-b) are obtained by simulating a number of cases using the lock simulation model. The cases differ in the managerial tactics used at the lock, but not in the technical tactics available (since MSDM is concerned with tactics that may be employed in the very short term, and installing new technical devices at a lock would require longer than a very short term). Each case yields another data point, consisting of values for SALT, CONC, FLUSH, and COST. We are able to interpolate among these points to form the functions (4.2a-b). For more discussion, see App. D of Vol. VA.

MSDM makes use of Eq. (4.2a) in two ways. When calculating the concentrations of salt at all the nodes, the value of FLUSH is known, at least provisionally, since it is the flow on a link of the network. With FLUSH fixed at its known value, Eq. (4.2a) becomes one of the terms in the objective function of Problem (4.1), a part of (Cost of reducing DISCH).

Less straightforward is the use made of Eq. (4.2a) when the water flows, including FLUSH, are to be calculated. Here it will suffice to say that an estimate is made of the reduction in the objective function of Problem (4.1) that would occur if FLUSH were changed, and that this estimate appears in the objective function of the problem that MSDM solves to find the water flows (see Chap. 2).

4.4. THE ROTTERDAM SALT WEDGE

The intrusion of seawater into the Rotterdam Waterway has been the subject of a great deal of study in PAWN (see Vol. XIX). Because the sea water intrudes by forcing itself underneath the lighter fresh water in the shape of a wedge, this phenomenon is called the Rotterdam salt wedge. The salt wedge can extend many kilometers inland. The fresh and salt layers are easily distinguished at locations close to the sea, but as the seawater intrudes farther upstream, mixing takes place, reducing the distinction between the two layers.

Salt intrusion in the Rotterdam Waterway plays a major part in the problems of water management in the Netherlands. The water supply to the midwest part of the country is provided mainly through the intake point at Gouda (node 20) on the Hollandsche IJssel River (link 61). The Hollandsche IJssel connects directly to the Nieuwe Maas (link 9), part of which coincides with the Rotterdam Waterway.

Salt intrusion into the Hollandsche IJssel has implications for both water consumers and for the various water boards. The present infrastructure offers limited possibilities to transport water to the midwest by other routes (emergency supplies). In the event that the emergency supplies are insufficient the water boards are forced to take water from the Hollandsche IJssel in spite of its salinity, in view of the requirement for level control in the boezems and the polder areas of the midwest. Raising the salinity in the boezems and polders in the region seriously affects regional water consumers, particularly those cultivating vegetables and flowers in glasshouses (see Sec. 4.7).

4.4.1. Points At Risk from the Salt Wedge

Gouda is the most important point influenced by the Rotterdam salt wedge, but it is not the only one. Some small withdrawals occur at the mouth of the Hollandsche IJssel. The effect of the salt wedge on the salinity here is considerably greater than the effect at Gouda, but less water is withdrawn, and fewer consumers depend upon it than upon withdrawals at Gouda. Thus, the influence of the salt wedge on the salinity at Gouda is much more important. Should the flows in links 4, 5, and 9 become low enough, the salt wedge could intrude into the Noord (link 8), where there is the inlet to Nederwaard, and farther into the Lek (link 7), perhaps as far as the inlets to Krimpenerwaard or Alblasserwaard.

The salt wedge intrudes into the Oude Maas as well as the Nieuwe Maas. If flows in links 4, 5, and 9 are low enough, it can reach the mouth of the Spui River, and possibly (although not likely) contaminate the Haringvliet. At still lower flows, it can intrude all the way to Dordrecht, where the Oude Maas separates from the Dordtsche Kil. This would contaminate water extracted there for drinking and industrial uses. We should note that while the salt wedge has been observed to intrude as far as the mouth of the Spui, usually during conditions of very high tide aggravated by a storm surge, it has never intruded nearly as far as Dordrecht. Figure 4.3 shows the locations that might be affected by the salt wedge.

In MSDM, we have ignored all points that are potentially affected by the salt wedge except the mouth of the Hollandsche IJssel (node 19) and the inlet at Gouda (node 20). The other points are hardly ever at risk.

4.4.2. Representing the Salt Wedge in MSDM

The model of the salt wedge used in the distribution model is described in Vol. XIX. There, the effects of the salt wedge at Gouda (MSDM node 20) and at IJsselmonde (MSDM node 19) are expressed as incremental chloride concentrations that depend on the flows in three waterways. The three waterways are the Nieuwe Maas (MSDM link 9), the Oude Maas (MSDM link 4), and the Nieuwe Waterweg (with no corresponding link in the MSDM network, but whose flow is the sum of the flows in the Nieuwe Maas and the Oude Maas). The dependence of the incremental concentrations on the flow in the Nieuwe Maas is much stronger than the dependence on the flow in the Oude Maas, and for simplicity we assumed in MSDM that the effects depended solely on this flow.

One problem was that, for flows of interest in the Nieuwe Maas, the incremental concentration at IJsselmonde is always greater than that at Gouda. This is true even though there is a flow from IJsselmonde (node 19) to Gouda (node 20) along the Hollandsche IJssel (link 61). In real life the explanation for this lies in the effect of the tides, but MSDM is ignorant of tides. As explained in Sec. 4.2, MSDM assumes that salt is carried passively downstream along with the flow of water, an assumption which implies that the incremental concentration at Gouda must equal that at IJsselmonde.

A second problem was that the formulation of the effects of the salt wedge as increments to concentrations did not harmonize with the formulation of Problem (4.1) in terms of salt concentrations (not increments) and salt discharges.

We solved these problems by introducing two functions into MSDM. One was a salt discharge at IJsselmonde which was a function of the flow

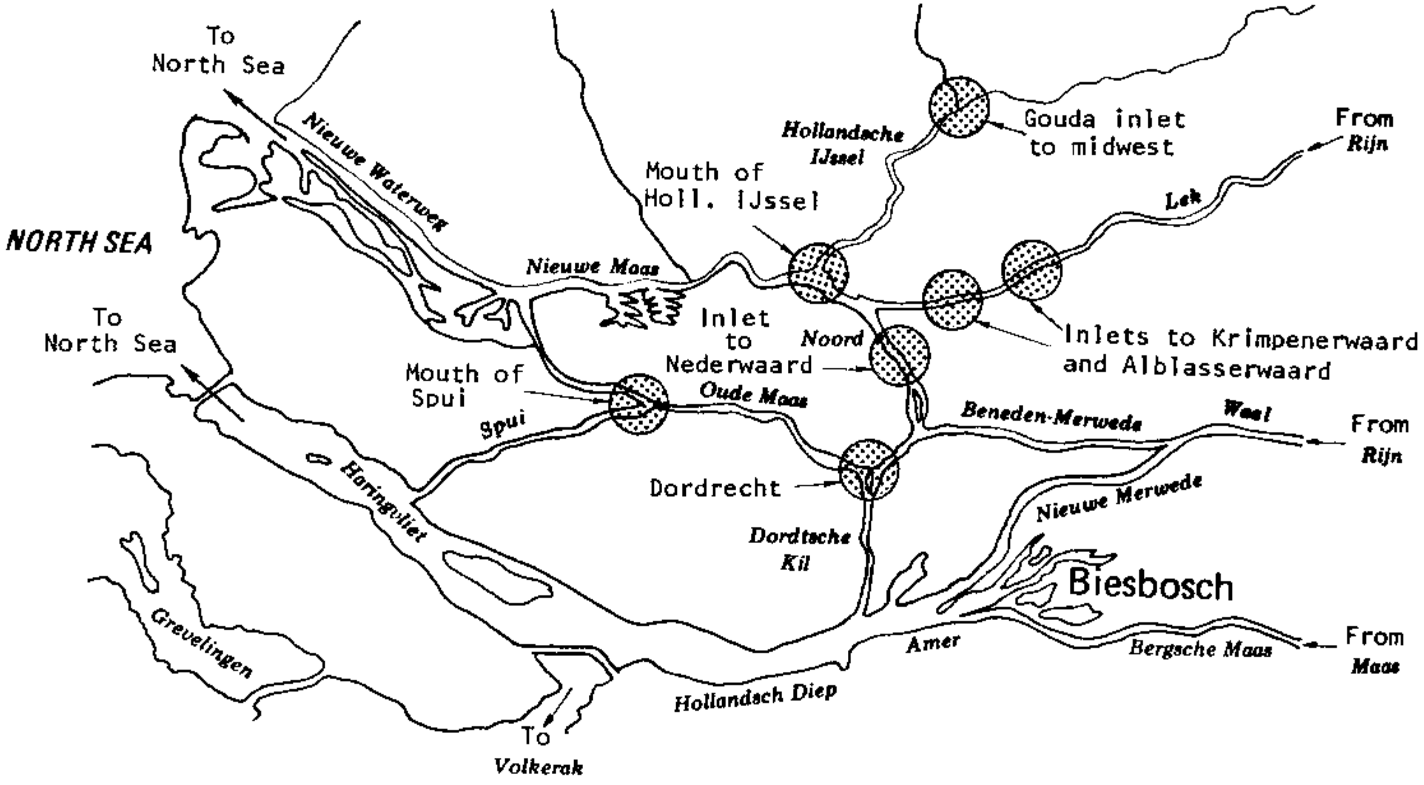

Fig. 4.3--Points potentially affected by the Rotterdam salt wedge

in the Nieuwe Maas. The salt thus discharged becomes part of the appropriate component of the vector DISCH in Problem (4.1), and is transported passively through the network by the flow of water. We call this the transportable component of the salt wedge. We adjusted this function to reproduce, as nearly as possible, the incremental salt concentration at Gouda predicted by the salt wedge model of Vol. XIX.

With only the transportable component of the salt wedge included, the incremental concentration at IJsselmonde will be too low. Accordingly, we introduced a second function into the MSDM model, which describes the remaining nontransportable component of the salt wedge as a function of the flow in the Nieuwe Maas. This variable serves to elevate the concentration of salt at the node 19 IJSLMOND by an internally calculated amount, without affecting the concentrations at any other node. We adjusted this function to make it as nearly equal as possible to the difference between the incremental concentrations at IJsselmonde and Gouda. For additional details, see App. D of Vol. VA.

4.4.3. Tactics Affecting the Salt Wedge

Under present circumstances, a freshwater flow of 625 m^3/s into the North Sea is considered the minimum adequate flow to prevent the salt wedge from reaching IJsselmonde (node 19 IJSLMOND). However, a number of technological devices have been proposed for reducing this requirement, the designs and costs of which are described in Vol. XVI. One proposal, making the upstream end of the Rotterdam harbor area shallower, was actually implemented, but has had a disappointingly small effect. Other proposals have been to construct groins or install bubble screens in the waterway to promote mixing of the fresh and salt water layers (see App. D of Vol. VA), or to build weirs or dams across one or more of the lower river branches (see App. A of Vol. VA) that could channel the water where it is most effective during periods of low Rijn flow. Weirs and dams would obstruct shipping, resulting in losses due to delays. These impacts on shipping are discussed in Vol. IX. Another type of proposal is to move the intake point for the midwest farther upstream, or to provide an alternate intake point (see App. A of Vol. VA), so that the salt wedge can intrude farther inland without damage than it does presently.

4.5. SALT IN BORDER-CROSSING RIVERS

Six rivers are represented in MSDM as crossing the border into the Netherlands. They are the Rijn (link 81), the Maas (link 82), the Roer and the Swalm (combined as link 83), the Niers (link 84) and the Overijsselsch Vecht (link 80). Each of the rivers contains some salt.

4.5.1. Salt in the Rijn

4.5.1.1. Locations of Salt Pollution. Another major source of salt in the Netherlands is the Rijn. In Fig. 4.4 we see the main sources of Rijn chloride. This figure is taken from Ref 4.4, App. 30.6, and represents chloride measurements taken along the length of the Rijn from Basel to Rotterdam between June 24 and July 1, 1974.

At Basel, the chloride concentration is less than 15 ppm, and this condition is maintained until just downstream of kilometer 200. At this point, the French discharge the waste salt from their Alsatian potash mines, just north of Mulhouse. This discharge increases the chloride concentration from 15 ppm to approximately 130 ppm. As the Rijn flow at this point was about 1400 m^3/s at the time of the measurements, the chloride discharge from the French mines must have been about 160 kg/s.

The chloride concentration remains at 130 ppm from kilometer 200 to approximately kilometer 600, where the Mosel and the Lahn enter the Rijn. At that point, the chloride concentration begins to rise gradually, reaching about 180 ppm at kilometer 780. The chloride concentration jumps to 220 ppm at Duisburg, where the Ruhr enters the Rijn. From there, the chloride concentration remains constant until the Rijn flows into the North Sea.

Of course, the actual concentrations shown in Fig. 4.4 are specific to the date on which they were measured. On other dates, the concentrations would be different. But the general pattern described here--i.e., the locations at which the concentration shows changes, and the relative sizes of those changes--will be the same regardless of the date. We can conclude, therefore, that perhaps 90 percent of the salt in the Rijn today is there because of human activities, and of that 90 percent, over half is due to discharges from the Alsatian potash mines.

4.5.1.2. The Trend of Rijn Salt. In 1876, the average chloride concentration in the Rijn at Lobith was less than 20 mg/l. In 1976 it was over 200 mg/l. Year by year, the concentration fluctuates around a definitely rising trend line, rising above the trend in years with low river flows, and falling below in years with high flows.

It is possible to eliminate much of these fluctuations around the trend by looking at the salt load instead of at the salt concentration. The load is simply the mass of salt carried past Lobith by the Rijn. It may be calculated as the product of the concentration and the river discharge.

The best simple description of the Rijn salt trend combines elements of both views. This description assumes that the salt passing Lobith during a particular year consists of a natural or background salt content, whose concentration is constant, plus an amount of salt dumped into the Rijn by industry and other human activities, which

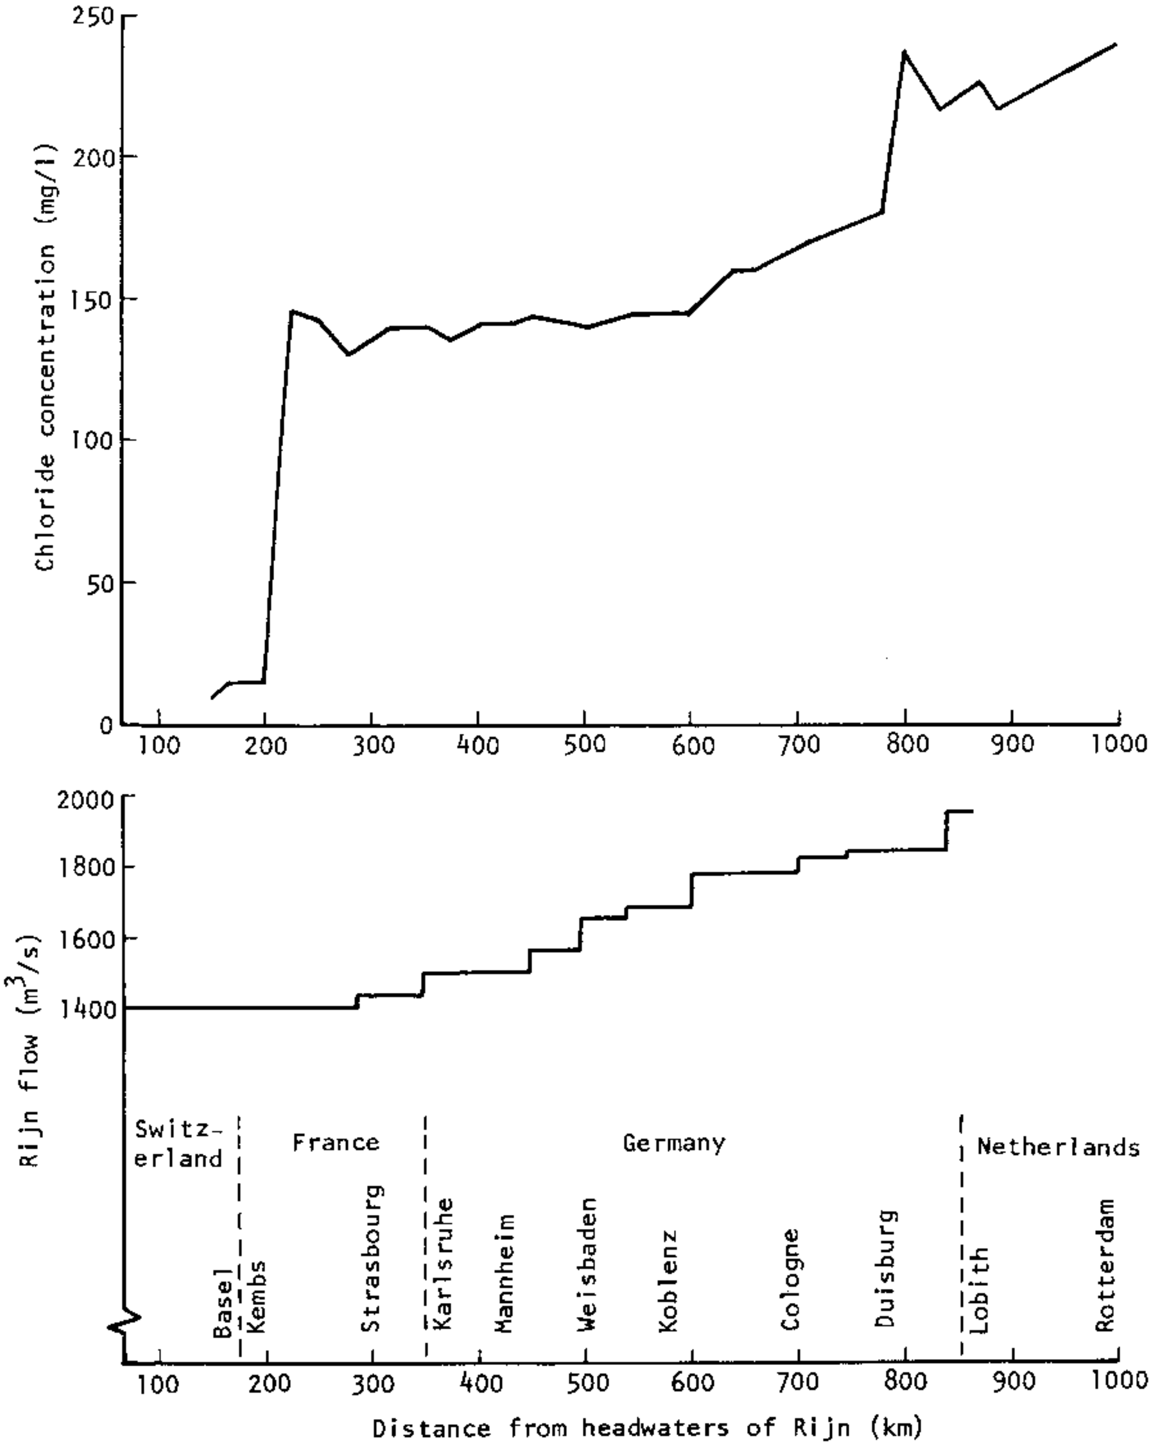

Fig. 4.4--Chloride concentrations along the Rijn measured between June 24 and July 1, 1974 (from "Internationale Kommission zum Schutze des Rhein Gegen Verunreinigung,") Postfach 309, 5400 Koblenz)

has increased at a constant rate in each year. Linear regression indicates that the background salt concentration is about 25 mg/l, close to the 1876 concentration of nearly 20 mg/l. Linear regression also indicates that the amount of salt dumped into the Rijn is increasing by 6 kg/s in each year. Figure 4.5 shows the salt dump as a function of year from 1930 through 1979 (see Vol. XI).

4.5.1.3. Measures to Reduce Rijn Salt. The Dutch border is at kilometer 862, so it can be seen from Fig. 4.4 that the increase in chloride in the Rijn is from French and German sources. Thus, the Dutch have no possibility for reducing the Rijn salinity by unilateral action. However, a treaty was agreed upon by the nations of the Rijn (Switzerland, Germany, France, and the Netherlands) to help the French finance a new method for disposing of the salt waste from their potash mines. The treaty envisioned a reduction of only 60 kg/s in the amount of chloride dumped by the French, so the reduction in the Rijn chloride content would be little better than marginal. In any case, the French, for internal political reasons, have felt unable to implement the treaty and at this writing it appears unlikely that the agreed reduction in chloride will take place soon.

4.5.2. Salt in Other Rivers

The salt content in the other border-crossing rivers causes less concern than the salt in the Rijn. They tend to be considerably less saline than the Rijn, with the occasional exception of the Roer, and their salinity is evidently not increasing with time.

The salt concentration in Maas water at Maastricht is frequently lower than 50 mg/l, and only rarely exceeds 125 mg/l of chloride. The Roer contains between 120 and 200 mg/l of chloride. Salt in the Niers varies between 70 and 120 mg/l. We have no direct measurements of salt in the Swalm and the Overijsselsche Vecht, but from observations at locations near the border, and probably not much compromised by water or salt from other sources, we guess that the Swalm contains approximately the same salt concentration as the Niers, and that the Overijsselsch Vecht contains approximately 150 mg/l of chloride.

4.5.3. Specification of River Salt in MSDM

The salt loads in the various border-crossing rivers are specified exogenously for each run of MSDM, and provided as an input. No managerial tactics are available to reduce the loads, so MSDM considers them to be inviolable constants.

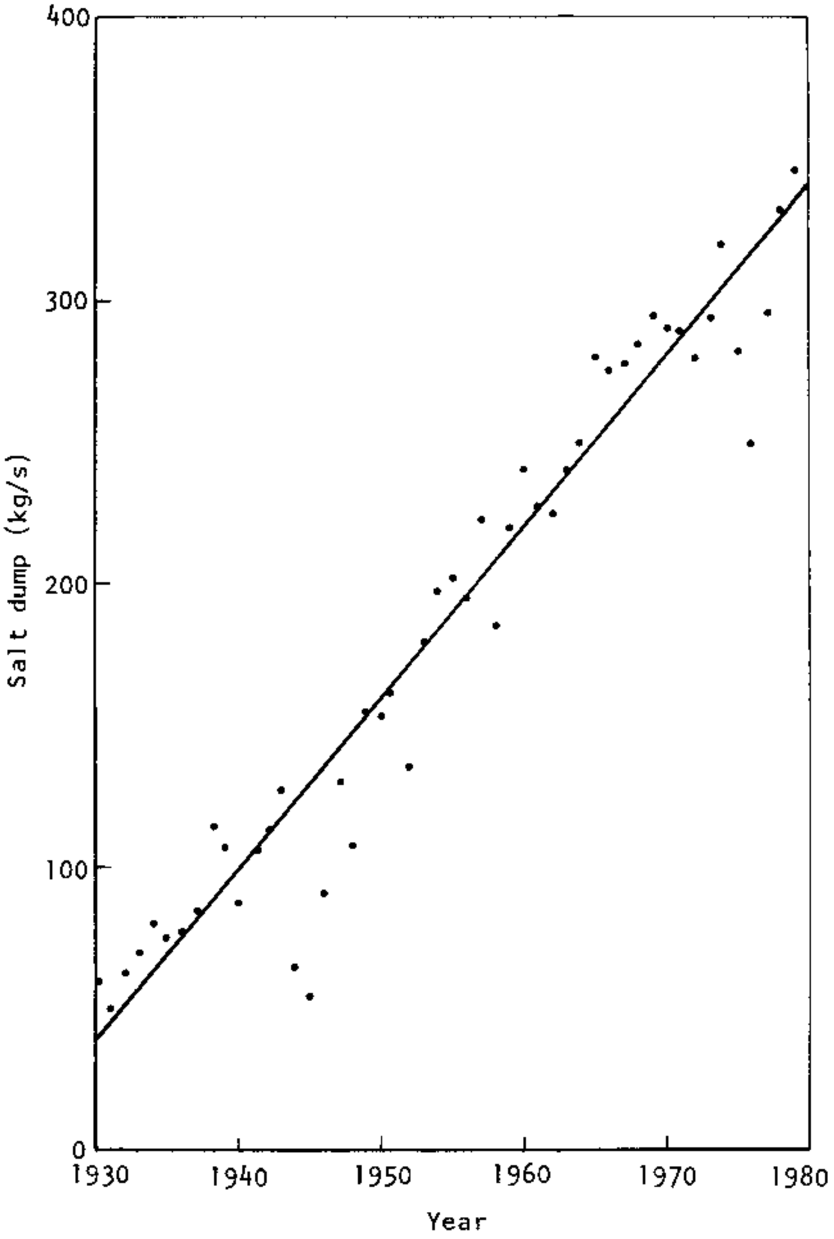

Fig. 4.5--Trend of the salt dump in the Rijn

4.6. SALT FROM DIFFUSE SOURCES

Diffuse sources of salt are sources that cannot be associated with specific locations. They are the opposite of point sources, which we have been discussing in the previous sections of this chapter. For salt, the diffuse sources are seepage, rain, and agricultural runoff.

4.6.1. Underground Seepage of Salt

Driven by hydraulic pressure, seawater seeps into the subsoil. From there it emerges in the surface water of the areas of the Netherlands that lie below sea level (the polder areas). The seepage is very slow, because the path over which it occurs is very long and the pressure difference that drives the flow is small. Further, by the time the seawater reaches the surface water in the polders, it has been much diluted by rain. The seepage rate is controlled by two long-term policies. First, the dunes along the coast are kept filled with fresh water. This creates a large, deep freshwater lens that blocks the landward movement of seawater. If that freshwater lens were to disappear, the seawater would have a much increased cross-sectional area through which to seep, and the seepage rate would accordingly increase greatly.

The second policy that controls the seepage rate is the policy of level control. Water levels in the canals and ditches of the polder areas are maintained within a few centimeters of the soil surface. We are told that a reduction in the water level of several decimeters might increase the seepage rate considerably, due to the reduction in the hydraulic pressure opposing seepage. (The water levels are kept 10-20 cm lower in winter than in summer, in anticipation of the greater excess of rain over evaporation, and hence greater required drainage capacity.) Figure 4.6 shows the seepage process.

In addition to these two policies, the salinity of the surface water in polder areas is further controlled by flushing. This is the policy of exhausting water from the ditch and canal system of the polders to the main rivers or canals, or directly to the sea, and replacing the discharged water by rain or, as needed, by fresh river water. The greater the flushing rate, and the fresher the makeup water, the lower will be the salinity in the surface water of the polders.

4.6.2. Other Diffuse Sources

Other diffuse sources are rain and agricultural runoff. There is a small amount of salt kept airborne by the wind, which is washed out of the air by the rain. In addition, rain falling on urban areas washes salt from the streets into the sewer systems, through which it runs off into the surface water. Most of this salt was spread on roads during winter to control ice.

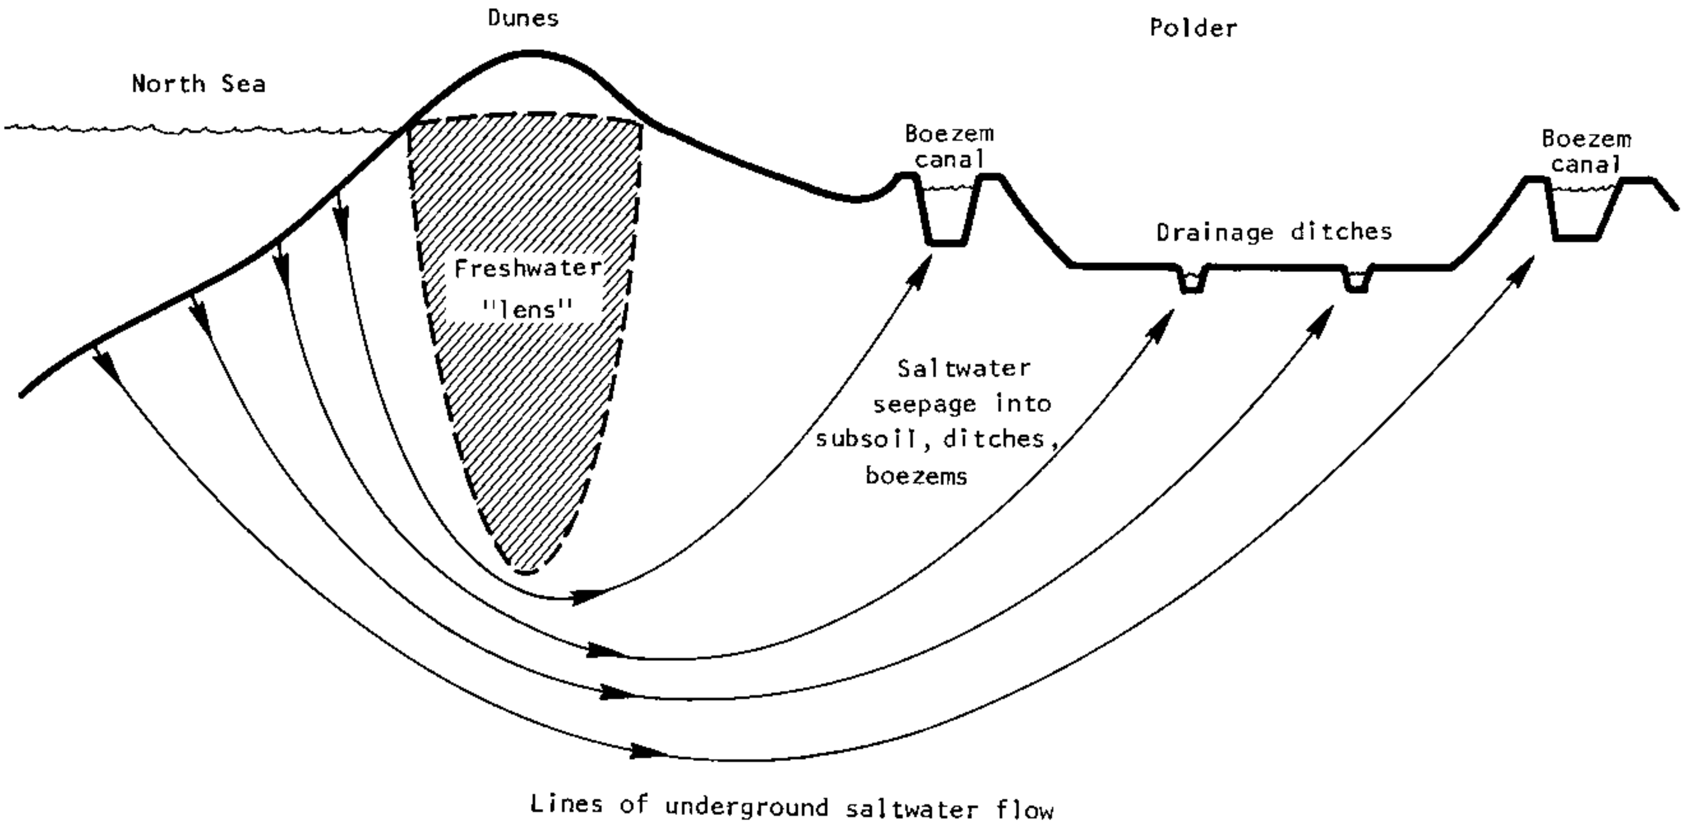

Fig. 4.6--Seepage of salt water into polder areas

Farmers in the Netherlands fertilize their land as heavily as farmers do anywhere, and much more heavily than most. But fertilizers contain salts of many kinds. These salts remain in the root zone for only a limited time, after which they are washed into the subsoil by rain or heavy irrigation, and eventually drain into the surface water as agricultural runoff.

4.6.3. Specification of Diffuse Sources of Salt in MSDM

Salt from diffuse sources is not a direct input to MSDM. As explained in Chap. 6, most discharges of salt (and water) from the land into the surface water are precalculated by a separate model called DISTAG (Vol. XII) and presented to MSDM as a stream of inputs. The only diffuse source discharges not treated in this way are those occurring from land irrigated with surface water. These discharges are calculated within MSDM.

4.7. SALINITY COSTS TO AGRICULTURE

Much of the Netherlands' agriculture suffers somewhat from salt pollution. For example, in the midwest, the physical yields of all the crops (vegetables and flowers under glass, open-air vegetables, flower bulbs and tubers, potted plants, ornamental shrubs and trees, and grassland farming) are lower than they would otherwise be as a result of prevailing salinity levels. The prevailing salinity levels in the region's rivers, canals, boezems, and ditches are about 200 to 300 mg/l in summers and, due to dilution, somewhat lower in winters. In dry years in these areas, summer salinity levels can be as high as 350 to 450 mg/l. Considerable variations from week to week, year to year, and location to location are possible.

The adverse affects of salinity on agriculture takes many forms. Yields may be reduced in weight, or in the number of physical units of product, or both. Some units, because of salt burn may have little or no marketable value. For example, high quality lettuce grown in Dutch glasshouses may have small brown edges at some of the leaves rendering the product almost worthless. Some crops may grow to harvestable size more slowly. This prevents farmers from obtaining favorable prices and reduces the number of harvests they can make per year from the same land. Finally, farmers may have higher costs for sprinkling because they try to leach out salts by heavy watering.

In this section, we consider the agricultural damage caused by high salinity levels. Most such damage is damage to crops grown under glass, which we discuss in Sec. 4.7.1. In addition, there are a number of very small areas that may be very strongly affected by local salt intrusion through locks. Because these areas are much smaller than the PAWN districts, we have treated this portion of salt damage to agriculture in a special way, which we discuss in Sec. 4.7.2.

4.7.1. Salt Damage to Crops Under Glass

Salt damage to vegetables and flowers grown under glass is by far the major portion of total salt damage to agriculture. Four factors contribute to this. First, the crops under glass are the most salt sensitive crops we consider. They begin to suffer damage when the chloride concentration in the root zone reaches 200 mg/l, whereas few of the crops grown in the open air are affected until the chloride level reaches 700 mg/l. Second, crops under glass derive little or no benefit from rain, which dilutes the root zone chloride for open-air crops. Their entire water needs are supplied from groundwater or surface water, which are much more saline than rain. Because of this factor, a crop grown under glass will suffer more salt damage than the same crop grown in the open air on a neighboring field. Third, most (nearly 80 percent) of the glasshouse cultivation is found in the midwest (Delfland, Rijnland, and Schieland), where the water supply is always highly saline due to seepage. Finally, crops grown under glass have the highest value per hectare of any crops we have considered. Even though they are grown on a relatively small area (6800 ha), their total annual value is about ten percent of the annual value of all agriculture in the Netherlands.

Our method for treating crops under glass is based on a rule of thumb we learned from Dr. Sonneveld at the Agricultural Research Station in Naaldwijk. The rule states that, in the long run, the chloride concentration found in the root zone of crops under glass will be approximately twice as high as the chloride concentration in the irrigation water. This rule arises from the sprinkling policy followed by the farmer. As the farmer irrigates his crops, water will be transpired and the salt it contained will be left behind in the root zone. If nothing is done, the root zone moisture will grow more and more saline, finally killing the crop. The farmer can flush the salt out of the root zone by adding more water to his crops than will be transpired. The excess will drain off, carrying with it the unwanted salt. Unfortunately, it will also carry with it the desirable nutrients that the farmer has added at some expense as fertilizer. The greater the amount of water added, the more of both salt and nutrients will be carried away. In balancing these two factors, the farmer has arrived at a happy medium, which is characterized by a root zone salinity about twice as high as the salinity of the irrigation water.

To represent crops under glass in MSDM, therefore, we attach a cost function to each of the twenty nodes that supplies water to crops under glass. The cost associated with a node depends only on the chloride concentration at that node. For any chloride concentration "c" at the node, the function assigns the damage the farmer's crops would suffer if there were a chloride concentration of "2c" in the root zones of all crops under glass supplied from this node. The salt damage function for flowers and vegetables under glass is shown in Fig. 4.7.

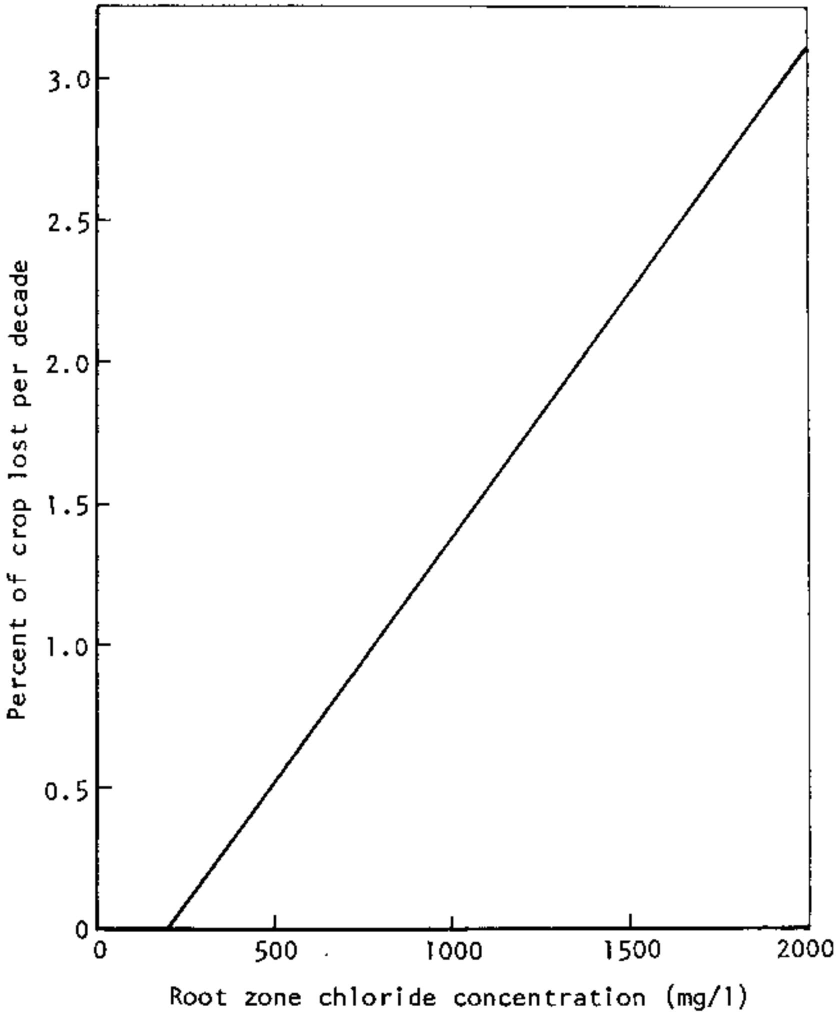

Fig. 4.7--Salt damage function for crops under glass

Figure 4.7 expresses damage as a fraction of the total value of the crop. To determine the damage in absolute terms, the fraction must be multiplied by the total value of glasshouse crops supplied from a particular node, and, for nodes with storage, by the length of time a molecule of salt can be expected to remain at the node to cause damage (the equivalent exposure time).[1] There are twenty nodes in the MSDM network that supply glasshouse crops with surface water. These nodes, together with the areas and values of glasshouse crops they supply, are shown in Table 4.1. Also in Table 4.1 is the expected equivalent exposure time of salt at each node. How the equivalent exposure time is calculated is discussed in App. B of Vol. VA.

4.7.2. Local Salt Damage Due to Salt Intrusion Through Locks

We have identified five locks at which salt intrusion takes place, but is confined to a small, very local area. These locks, whose locations are shown in Fig. 4.1, are Parksluizen, Spaarndam, Den Helder, Harlingen, and Delfzijl. It would be inappropriate to treat these locks as we have discussed in Sec. 4.2.4, because if we were to discharge the intruding salt into an MSDM network node, it would spread farther and affect more cultivated area than is the case in reality. Nevertheless, heavy salt damage may occur to crops supplied with water from points near these locks, so neither can we ignore these instances of salt intrusion. Therefore, we have developed an alternative procedure for dealing with these locks, which we describe in detail in App. D of Vol. VA.

The general layout of each of these locks, and its relation to the crops supplied with water from nearby points, is shown in Fig. 4.8. Salt intrudes from the saltwater side of the lock, entering the canal on the freshwater side. The flow of fresh water in the canal is always toward the lock, so the salt must move upstream in opposition to the flow in order to reach the extraction point for agriculture. If the salt reaches the extraction point, salt damage to locally grown crops will follow.

For each of our five locks, we have identified the extraction point nearest the lock (we ignore more distant extraction points), and the values and sensitivities of the crops supplied from that point. The distances from the locks to the associated extraction points is 3-4 km for most locks, but 9 km for Den Helder. The nearest major extraction point to Delfzijl appears to be so far away that intruding salt will never reach it, and we therefore dropped this lock from consideration.

As mentioned above, intruding salt must move upstream against the flow of water to reach the extraction point. This flow is called the <u>flushing rate</u>. Not surprisingly, the greater the flushing rate, the less salt will reach the extraction point, and the less damage will be done to locally grown crops. We have derived a rough relation between the flushing rate and the incremental salt concentration at the extraction point. This relation builds upon

Table 4.1

GLASSHOUSE CROPS IRRIGATED FROM SURFACE WATER

Supply Node	Area Supplied (ha)	Total Value (a) (Dflm/yr)	Open Water Volume (Million m^3)	Average Outflow Rate (m^3/s)	Equivalent Exposure Time (dec)
1. FRIELAND	51	16.5	180.6	38.4	11.5
3. GRONETAL	69	23.9	-	-	1.0
4. NEHIGH	62	17.9	-	-	1.0
6. UPRIVER	197	70.5	-	-	1.0
7. TWENMOND	15	5.5	-	-	1.0
8. TWENTEND	3	1.5	-	-	1.0
10. IJSLAKES	166	65.8	8036.9	323.3	57.4
11. NORHOLL	125	48.2	37.1	22.3	4.5
12. A-R.MOND	105	32.5	-	-	1.0
17. LOPIKWAR	11	3.3	2.0	11.2	1.3
18. VECHT	23	7.9	9.0	1.8	12.1
19. IJSLMOND	33	9.2	-	-	1.0
20. GOUDA	439	135.0	-	-	1.0
21. MIDWEST	4895	1737.5	119.2	22.5	12.8
24. GOR+DOR	50	13.9	-	-	1.0
25. LOWRIVER	283	78.1	880.6	716.3	3.5
26. DLTALAKE	39	11.3	819.2	50.0	38.1
31. SAMGRALI	21	7.1	60.1	216.9	1.4
32. DBOS+BOX	159	51.8	-	-	1.0
33. GERTRUID	57	18.3	-	-	1.0

(a) Values are calculated assuming that 1 ha of glasshouses devoted to vegetables produces 232,000 Dfl/yr, while the same area devoted to flowers produces 485,000 Dfl/yr. This reflects average historical prices. In dry years, the prices may change; see Vol. XIV for discussion.

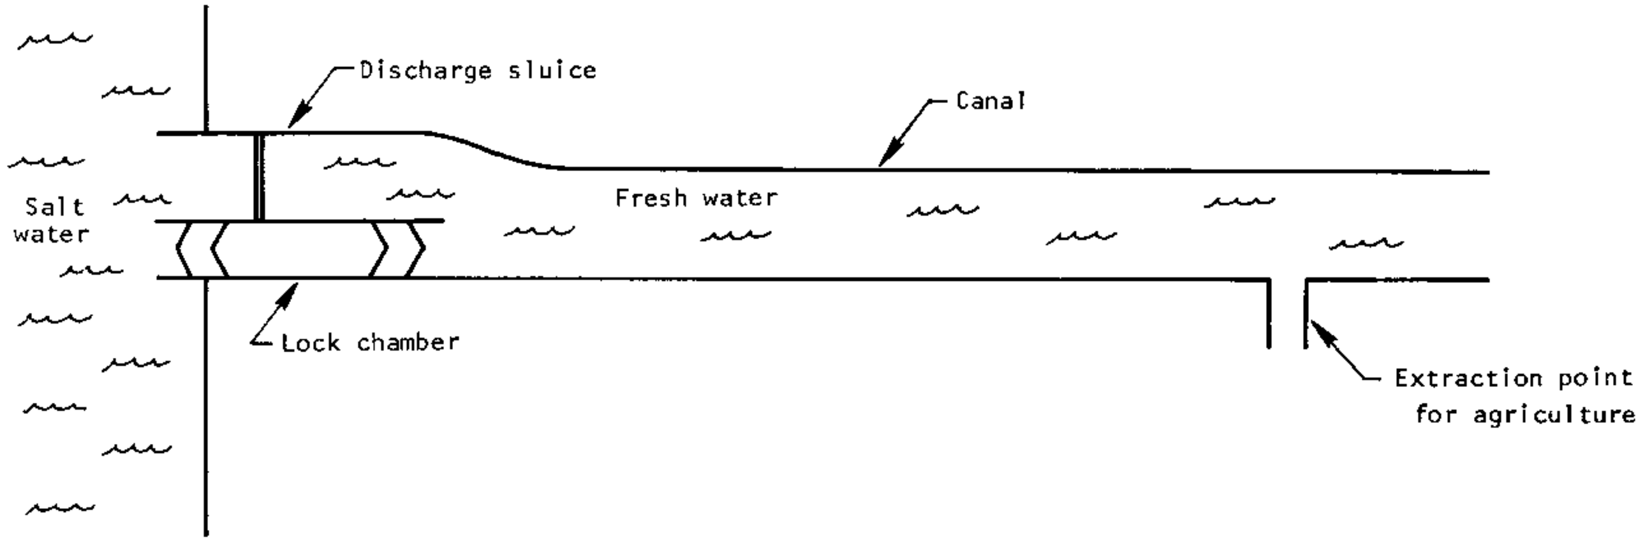

Fig. 4.8--Layout of a lock that may cause local salt damage to agriculture

Function (4.2b) obtained from the lock simulation model. Recall that Function (4.2b) relates the flushing rate and the cost of managerial tactics at the lock to the incremental salt concentration immediately to the freshwater side of the lock.

This very simple relation can be multiplied by the appropriate crop damage coefficient to yield the crop damage as a function of the flushing rate. We add to the result the costs of managerial tactics at the lock, and determine, for each flushing rate, the method of lock operation that results in the smallest total losses due to both lock operating costs and crop damage. The end result is a function that associates a total loss with each flushing rate at the lock.

For each of the four locks, the flushing rate is the rate of flow in one of the links of the network. Thus, these total loss functions appear in MSDM as cost functions associated with the flows in appropriate network links. The functions all have similar shapes, and in Fig. 4.9 we show an example for the lock at Harlingen. The cost is very high at low flushing rates, but drops very rapidly to a minimum cost at quite moderate flushing rates.

4.8. OTHER SALINITY COSTS

Other costs of excess salt in water are borne by industry, commercial activities, and households. Industries use water for two very different sorts of purposes. One is cooling. PAWN has concluded that no improvement in the quality of water used for cooling is likely to yield significant benefits, and we have elected to ignore what slight benefits there may be [4.5].

The other industrial use of water is for process use. In the ways relevant for our study, this use is similar to household, commercial, and municipal use. Dissolved solids in water used for these purposes speed up the deterioration of equipment and furnishings (pipes, faucets, fabrics, appliances, etc.), increasing user's costs by requiring more maintenance and cleaning, and earlier replacement. United States data converted to Dutch equivalents suggest that cost increases for cleaning and for replacement of deteriorated equipment and furnishings amount to about one guilder for each kilogram of dissolved solids. The Dutch measure of salinity is chloride ion, which constitutes about one-third of total dissolved solids in typical river water. Accordingly, the above cost becomes three guilders per kilogram of chloride ion.

This would suggest that industries and households would benefit by three guilders per kilogram of chloride removed from the Rijn. But only a small fraction (less than 0.2 percent) of the Rijn flow is used as industrial process water or in households. It may be worthwhile to remove dissolved substances from the water actually used for these purposes, but it is surely not worthwhile to treat the entire Rijn flow in the same way.

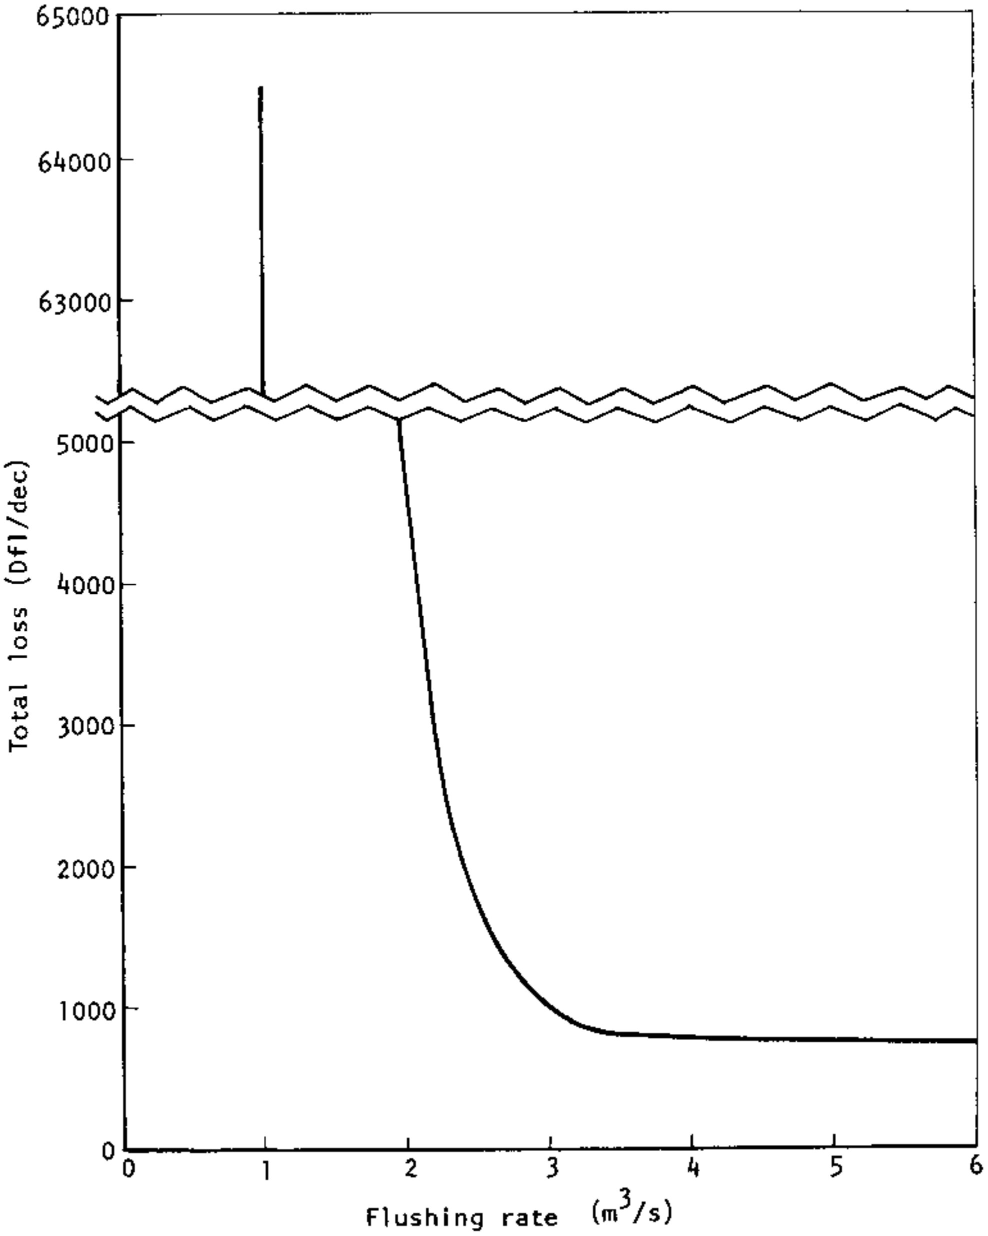

Fig. 4.9--Local crop damage plus managerial tactic costs for the lock at Harlingen

In fact, the costs that salt imposes on industry, commercial activities, and households can be treated entirely separately from MSDM. These costs depend, not on the salinity at any particular time, but on the average salinity over the long term. In addition, reductions in the salinity of Rijn or other water can only be brought about by tactics applied steadily over the long term. Since MSDM is directly concerned only with very-short-term actions, we have elected not to include any costs to industry, commercial activities, or households among the salt-related costs in MSDM.

4.9. WATER QUALITY STANDARDS FOR SALINITY

Excess salt in the water supply also affects public health and the natural environment. Unfortunately, it has proved impossible to determine the costs in guilder terms that salt imposes on either of these areas. Instead of using even a rough estimate of such costs, we impose a water quality standard, a limit beyond which we do not wish the salt concentration to rise. Imposing a standard has the advantage that the standard can be set on the basis of physiological, biological, or ecological data directly. No assumptions need be made about the equivalent economic value of better health or a more diverse ecology. But the level at which the standard is set does have economic consequences, in the sense that in order to meet the standard someone may suffer a loss, or give up a profit.

4.9.1. The Drinking Water Standard

Excess total dissolved solids (which include chloride) are objectionable in drinking water because of possible physiological effects (e.g., high blood pressure), unpalatable tastes, and unappealing appearance and odor.

Few estimates quantifying the effect on taste, appearance, and odor have been prepared using Dutch data. United States estimates suggest that the effect is important to users. Chloride ions have frequently been cited as having a low taste threshold in water. Data from Ref. 4.6 on a taste panel of 53 adults indicated that 61 mg/l NaCl was the median level for detecting a difference from distilled water. At a median concentration of 395 mg/l chloride a salty taste was distinguishable, although the range was from 120 to 1215 mg/l. Lockhart et al. [4.7] evaluated the effect of chlorides on water used for brewing coffee. Threshold concentrations for chloride ranged from 210 mg/l to 310 mg/l, depending on the associated cation. These data indicate that a level of 250 mg/l chlorides (the EPA criterion) is a reasonable maximum level to protect consumers' drinking water.

In the Netherlands, there seems to be nearly complete agreement that the best overall drinking water quality standard for chloride is 200 mg/l. This standard has been adopted, at least provisionally by the International Commission of the Rijn [4.8].

Drinking water standards are not, however, particularly relevant for MSDM. A relatively small amount of water is required to meet drinking water standards. Incidentally, much of the process water used by industry and most of the water used for commercial and household purposes is purchased from drinking water companies, and is required to meet these standards.

A great deal of the drinking water in the Netherlands comes from either groundwater or from the Maas, both of which contain little salt. For example, the city of Rotterdam and the province of Zeeland both have their drinking water piped from the Maas.

Locations that rely solely on highly saline water (usually Rijn water) for their drinking water supply could treat this water, for example by reverse osmosis plants. The annualized cost of such treatment is in the neighborhood of two or three guilders for each kilogram of chloride removed, depending on the size of the plant [4.5].

For the most part, however, treatment is unnecessary. Drinking water is extracted from the rivers only when they are relatively clean. This usually occurs when they have relatively high flows. This water is then stored in reservoirs for later use. Thus, the average, or even the minimum, salinity of the Rijn and other rivers is more relevant for drinking water standards than is the maximum salinity. Certainly, the salinity at any particular instant is of limited interest for assessing the impact of salt on drinking water.

4.9.2. Environmental Standards

The salinity of the water in a particular area will influence the composition of the plant community found there. RIN distinguishes three types of water, each accompanied by its characteristic types of vegetation. Atmotrophic water is rainwater, or water from the unpolluted headwaters of a river. Lithotrophic water is groundwater or unpolluted river water. Hydrotrophic water is seawater or polluted river water. The amount of salt, as measured by chloride, increases from atmotrophic to lithotrophic to hydrotropic water.

Actually, the chloride concentration alone is not sufficient to distinguish the three kinds of water. RIN considers that calcium ion is another important determinant of vegetation type. Probably it is a gross simplification to rely on either ion alone, or even both together; the entire composition of the water is undoubtedly important. The suggestion was made to use the fraction of Rijn water supplied to an area as a measure of water type, but this is so highly correlated with chloride concentration that we deemed it no improvement. In spite of the oversimplification involved, therefore, chloride is our only measure of water type.

Any of the three types of water can support a lush plant life, but different types of plants will grow in the three cases. What is to be avoided is fluctuations in the type of water, since the fluctuations will, if large enough, prevent any plant from becoming well established. Accordingly, RIN has suggested different standards in different PAWN districts. In some PAWN disticts they want chloride not to exceed 20 mg/l, and in others they want it always to exceed 500 mg/l. In many districts, they suggest the same standard as the IMP, namely 200 mg/l.

Figure 4.10 shows RIN's suggested standards [4.9] for each of the PAWN districts. The pattern is quite clear. In the northeastern and southern highlands, where present salinities are low, RIN wishes them to remain low. In areas along the main rivers, where salinities are presently higher, RIN would allow them to remain so, although they would prefer to see reductions in the salinity of the Rijn and adjacent districts. In the lowlands of the midwest and Friesland, they have suggested that 300 mg/l be the standard. These are areas that suffer from considerable salt seepage, and for which there is no prospect of meeting a more stringent standard. In North Holland, also subject to high salt seepage, they prefer that the chloride concentration be maintained <u>above</u> 200 mg/l. There is no prospect for bringing the salinity below this figure. In a few isolated areas--along the Noordzeekanaal, on the shores of the Volkerak and the Oosterschelde, there exist very high salinities at present, and RIN wishes them to be kept that way. In short, RIN has opted for the status quo ante, for less saline water where salinities are now relatively low, and for highly saline water where salinities are now high.

Unfortunately, like the drinking water standards, the environmental standards have little relevance for MSDM. The reason is that they apply to the water supplied directly to the plant communities. This water is a mixture of rain, groundwater, and river water, and is almost always quite different in composition from the water supplied to an MSDM network node. In most cases, the water supplied to a network node is almost entirely river water.

4.9.3. Salinity Standards in MSDM

We have recognized that the water quality standards for salt have little relevance for MSDM. Nonetheless, we have adopted standards for use in some MSDM runs. Most of the time we take 200 mg/l to be our standard everywhere, but in some runs we have considered the location-specific RIN standards. In all runs, we calculate the salt concentrations, so that they may be compared with such standards as may be deemed appropriate.

4.10. THE VALUE OF WATER USED FOR REDUCING SALT CONCENTRATIONS

We distinguish three ways in which water can be used to reduce the damage caused by salt. The first is to increase the flow of water in the Nieuwe Maas (link 9 NIEWMAAS), thereby reducing the

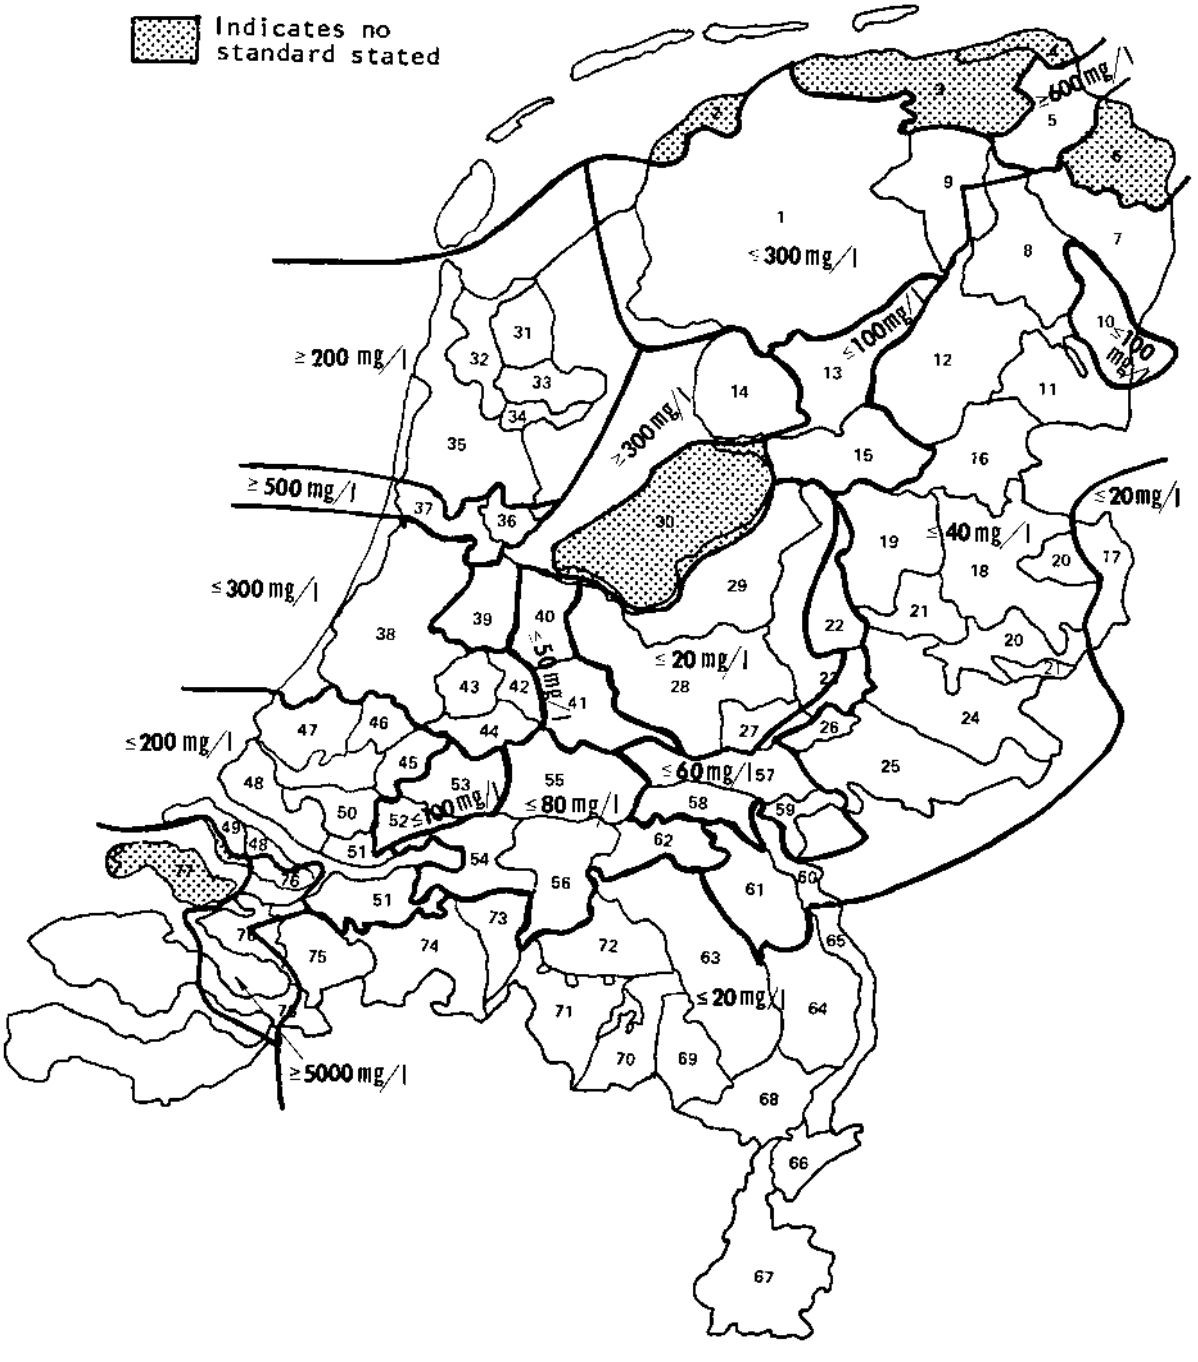

Fig. 4.10--RIN chloride standards by PAWN district

effect of the salt wedge on the chloride concentration at nodes 19 IJSLMOND, 20 GOUDA, and 21 MIDWEST. This will have the effect of reducing the damage to crops grown under glass that are supplied from these nodes.

The second use of water that will reduce salt damage is flushing at locks. This includes both locks at which salt intrusion is confined to a local area, and locks from which intruding salt can spread widely through the network. Water used for this purpose will reduce the salt concentrations at various points in the network, and thereby reduce the amount of damage to crops under glass.

The third method of using water to reduce salt damage is to dilute salty water at nodes supplying glasshouse crops by flushing them with less saline water. This method is potentially effective in lowland areas, where salt seepage acts to elevate chloride concentrations. Elsewhere, the water available for flushing has the same chloride content as the water already available at the node.

4.10.1. Combating the Salt Wedge

The damage due to the Rotterdam salt wedge is almost entirely damage to crops grown under glass. As explained in Sec. 4.7.1, these are the most salt sensitive crops in the PAWN study, as well as being the most valuable per hectare. Furthermore, they are concentrated in the Midwest, where the influence of the salt wedge is felt.

To estimate the damage produced by the salt wedge, we calculate the incremental exposure to chloride due to the salt wedge at the three relevant nodes, 19 IJSLMOND, 20 GOUDA, and 21 MIDWEST. Then we will multiply these incremental exposures by the values of glasshouse crops supplied from each node and by the sensitivities of these crops to salt damage. The final step will be to calculate by how much the damage would be reduced if various key flows in the network were increased. The ratio of the damage reduction to the increase in flow will be the value of water in that particular application.

4.10.1.1. Nontransportable Effect at IJSLMOND. Recall that we have represented the effect of the salt wedge on chloride concentrations as the sum of two components: a transportable component which consists of salt that flows passively through the network downstream from IJSLMOND; and a nontransportable component, which is confined to IJSLMOND. Let us first consider the effect of the nontransportable component. In MSDM, the size of this effect on chloride concentration depends only on the flow in link 9 NIEWMAAS, and may be found in App. D of Vol. VA.

From Table 4.1, we find that a value of only 9.2 Dflm/yr of glasshouse crops are supplied from IJSLMOND. From Fig. 4.7 we find that an increase of one mg/l chloride in the root zone of glasshouse crops for one decade results in a reduction of 17 parts per million in their eventual value. We multiply this figure by two to account for the difference between the concentration in the surface water as compared to the root zone. Thus, if we define F(9) to be the flow in link 9 NIEWMAAS, and if we denote by NTC(F(9)) the nontransportable component of the salt wedge as a function of F(9), measured in milligrams chloride per liter, we can write the damage due to the nontransportable component as:

$$SW_1 = 9.2*2*0.000017*NTC(F(9)) = 0.0003128*NTC(F(9)) \tag{4.3}$$

4.10.1.2. Transportable Effect on IJSLMOND and GOUDA. The effect of the transportable component of the salt wedge is felt at all three nodes, but its effect at IJSLMOND and GOUDA is simpler than its effect at MIDWEST because water is stored at the latter node.

In MSDM, the size of the transportable effect is expressed as an amount of salt discharged into node 19 IJSLMOND. The amount of the discharge depends only on the flow in link 9 NIEWMAAS, as described in App. D of Vol. VA. The resulting incremental chloride concentration due to this discharge can be calculated as the ratio of the discharge to the total flow of water leaving IJSLMOND, which to a very good approximation is F(9) - min{F(61),0}. Here, F(9) is (as before) the flow in link 9 NIEWMAAS, while F(61) is the flow in link 61 HOLIJSEL. The presence of the negative sign is explained by the fact that the flow in link 61 is considered positive when water moves from GOUDA to IJSLMOND, and negative when it moves in the opposite direction.

The value of glasshouse crops supplied from IJSLMOND and GOUDA can be obtained from Table 4.1. Because the transportable component of the salt wedge exposes all of these crops to the same incremental chloride concentration for the same amount of time, we can simply add these values, obtaining a total of 144.2 Dflm/yr. If we denote by TC(F(9)) the transportable component of the salt wedge, measured as kilograms chloride discharged per second, we can write a damage function similar to Eq. (4.3), expressing the damage caused by the transportable component of the salt wedge at IJSLMOND and GOUDA.

$$SW_2 = 144.2*2*0.000017*\frac{1000.0*TC(F(9))}{F(9)-\min\{F(61),0\}} \tag{4.4}$$

$$= 4.9028*\frac{TC(F(9))}{F(9)-\min\{F(61),0\}}$$

4.10.1.3. Transportable Effect on MIDWEST. Finally, we calculate the effect of the transportable component of the salt wedge on glasshouse crops supplied from node 21 MIDWEST. As explained in Sec. 4.7.1, the damage to crop will be proportional to the average effect of the salt wedge on the chloride concentration during the decade, multiplied by the equivalent exposure time from Table 4.1. For MIDWEST, the average effect on chloride concentration will be approximately half of the total effect, and the total effect will be approximately the ratio of the total amount of salt entering MIDWEST during the decade due to the salt wedge, divided by the volume of water stored at MIDWEST.

The total amount of salt entering MIDWEST during the decade is straightforward to calculate. It must be the product of the incremental concentration at GOUDA, which we have discussed in Sec. 4.10.1.2, multiplied by the flow from GOUDA to MIDWEST. We may take this flow to be -F(61). As in the previous two sections, it is possible to write an expression for the damage due to the salt wedge, this time damage to glasshouse crops supplied from MIDWEST.[2]

$$(4.5)\quad SW_3 = \frac{-\min\{F(61),0\}*0.876*12.8*1737.5*2*0.000017*1000*TC(F(9))}{2*119.2*(\ F(9)-\min\{F(61),0\})}$$

$$= 2.7785*\frac{TC(F(9))*(-\min\{F(61),0\})}{F(9)-\min\{F(61),0\}}$$

The factor 0.876 converts the flow F(61) from the units of cubic meters per second to those of millions of cubic meters per decade.

4.10.1.4. The Value of Water Used to Combat the Salt Wedge. We are now in a position to estimate the value of water used to combat the salt wedge. First, we calculate the losses due to the salt wedge as a function of F(9) and F(61), as the sum of Eqs. (4.3-4.5). Because the losses depend on two variables, it is difficult to represent this function graphically. Figure 4.11 is an attempt to do so by showing several constant-loss contours. Note that the contours are much closer together--which means that the loss is much more sensitive to changes in flows--for low flows than for high flows in the Nieuwe Maas, and for high withdrawals than for low withdrawals from the Hollandsche IJssel at Gouda. Note also that the loss is more sensitive to a change of 1 m^3/s in the withdrawals at Gouda than to the same change in the flow in the Nieuwe Maas.

Water can only have an effect on the salt wedge if it alters either the flow in the Nieuwe Maas or the withdrawals at Gouda. But there are many managerial tactics that will accomplish this. Four likely ones are:

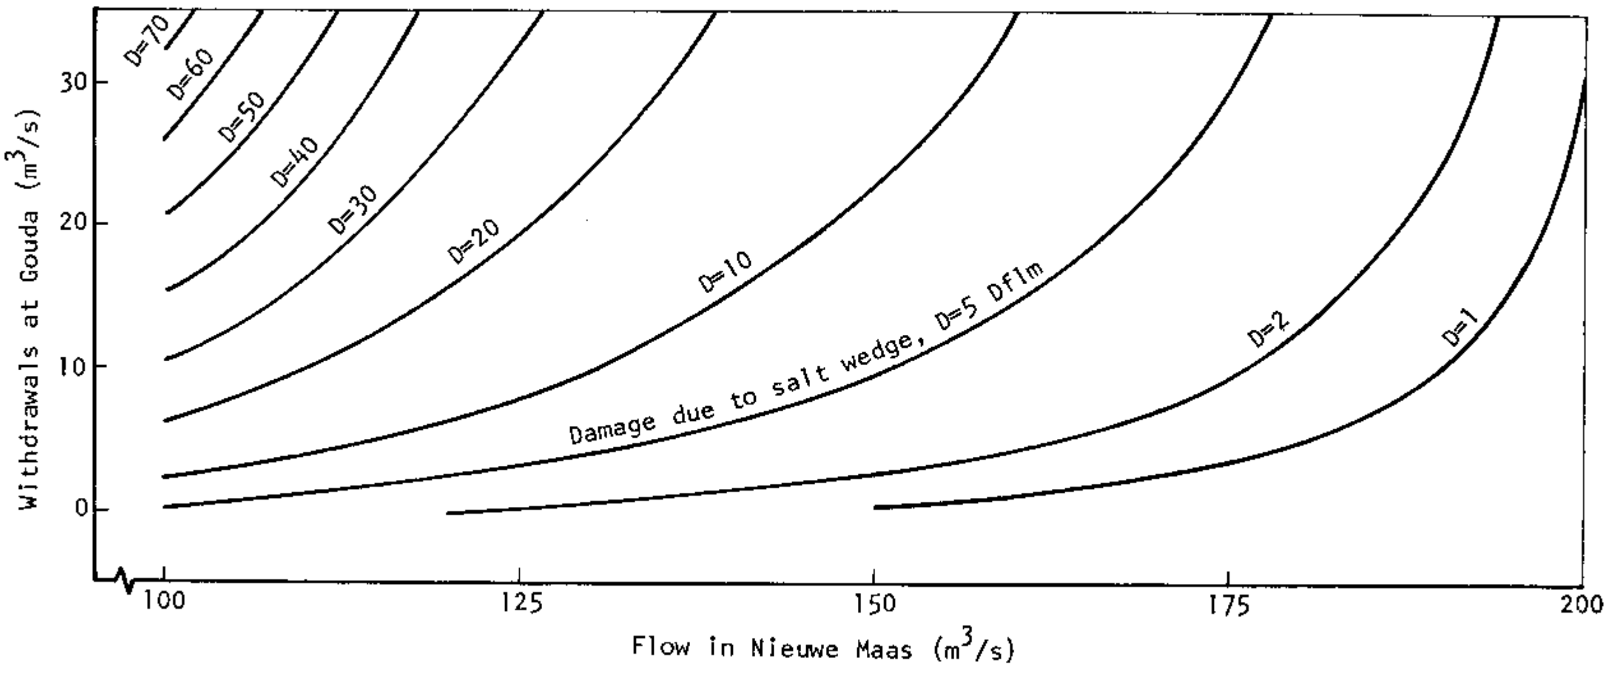

Fig. 4.11--Constant-loss contours of the damage due to the Rotterdam salt wedge

- Reduce the flow in the Amsterdam-Rijnkanaal north from the Lek (link 29), and let it flow west along the Lek instead (link 7). The hydrological equations for the lower rivers part of the network (see App. A of Vol. VA) predict that for every cubic meter per second so diverted, the flow in the Nieuwe Maas (link 9) will increase by 0.864 m^3/s.
- Withdraw water from the Waal at Tiel (along link 28), and let it flow west along the Lek (link 7). The lower rivers hydrological equations predict that each cubic meter per second so diverted will increase the flow in the Nieuwe Maas by 0.503 m^3/s.
- Open the wier at Driel to permit more water to flow west on the Neder-Rijn (link 6). The upper rivers hydrological equations predict that for every 1 m^3/s additional flow in the Neder-Rijn, the flow in the Waal will be reduced by 0.58 m^3/s, and the flow in the IJssel by 0.42 m^3/s. The lower rivers hydrological equations predict that these changes in flow will result in an increase of 0.65462 m^3/s in the Nieuwe Maas.
- Reduce withdrawals at Gouda. The reduction would have a double effect, since it would not only directly reduce withdrawals at Gouda, but would increase the flow in the Nieuwe Maas by 0.864 m^3/s for each 1 m^3/s reduction in withdrawals. Of the four tactics described here, this is the only one that affects withdrawals at Gouda. But this tactic must be accompanied by an increase in the water supply to the midwest from other sources, or a reduction in water use in the midwest.

The value of water used according to any of these managerial tactics can be estimated by comparing the damage due to the salt wedge in two situations that differ only in degree to which the tactic is exercised. For example, we can compare two situations whose only differences are that one more cubic meter per second flows north along the Amsterdam-Rijnkanaal in the first situation than in the second situation, and that there are compensatory differences in the flows in the Lek, the Nieuwe Maas, and other branches of the lower rivers network. The difference in the damage caused by the salt wedge during a decade in the two situations divided by the number of cubic meters that were diverted from the Amsterdam-Rijnkanaal in that decade would be the value of water used according to the first of our managerial tactics.

In Tables 4.2 and 4.3 we present some illustrations of the value of water used to combat the salt wedge by these various tactics under several different circumstances. In Table 4.2 we examine four cases for each of which the withdrawal at Gouda is 19 m^3/s, which is typical of a summer decade. The four cases differ in their Rijn flows, from 1000 m^3/s to 700 m^3/s. A flow of 1000 m^3/s or less in the Rijn occurs in 6.4 percent of all decades, or an average of 2.25 decades per year. A flow of 900 m^3/s or less may be expected 1.23 decades per year; 800 m^3/s, 0.55 decades per year; and 700 m^3/s only 0.11 decades per year. In fact, the Rijn flow has been less than 700 m^3/s in only five decades in the 47 years between 1930 and 1976.

Table 4.2

VALUE OF WATER FOR COMBATING THE SALT WEDGE
(For a Withdrawal at Gouda = 19.0 m^3/s)

Flows (m^3/s):				
Rijn	1000.0	900.0	800.0	700.0
Nieuwe Maas	211.50	185.32	159.35	133.6
Salt from Salt Wedge:				
TC, (kg/s)	0.0	6.19	16.82	34.4
NTC, (mg/l)	277.96	368.0	470.6	572.3
Damage from Salt Wedge (Dflm/dec):				
Eqn. (4.3)	0.0869	0.1151	0.1472	0.1790
Eqn. (4.4)	0.0	0.1487	0.4631	1.1072
Eqn. (4.5)	0.0	1.5994	4.9787	11.9006
Total	0.0869	1.8632	5.5890	13.1868
Value of Water (Dfl/m^3):				
A-R kanaal	0.0007	0.069	0.166	0.352
Tiel	0.0004	0.040	0.096	0.205
Driel	0.0005	0.052	0.125	0.266
Gouda	0.0007	0.135	0.368	0.826

Table 4.3

VALUE OF WATER FOR COMBATING THE SALT WEDGE
(For a Withdrawal at Gouda = 35.0 m^3/s)

Flows (m^3/s):				
Rijn	1000.0	900.0	800.0	700.0
Nieuwe Maas	197.44	171.26	145.29	119.54
Salt from Salt Wedge:				
TC, (kg/s)	2.73	11.5	24.94	48.87
NTC, (mg/l)	320.11	423.5	526.1	627.8
Damage from Salt Wedge (Dflm/dec):				
Eqn. (4.3)	0.1001	0.1325	0.1646	0.1964
Eqn. (4.4)	0.0576	0.2734	0.6782	1.5504
Eqn. (4.5)	1.1422	5.4220	13.4525	30.7525
Total	1.2999	5.8279	14.2953	32.4993
Value of Water (Dfl/m^3):				
A-R kanaal	0.100	0.172	0.308	1.110
Tiel	0.058	0.100	0.179	0.646
Driel	0.076	0.130	0.233	0.841
Gouda	0.124	0.283	0.576	1.697

The flow in the Nieuwe Maas depends only partly on the flow in the Rijn. To obtain the Nieuwe Maas flows in the table, we assumed that the wier at Driel was allowing only 25 m^3/s to flow in the Neder-Rijn, and that the withdrawals at Tiel were the minimum required to maintain the flow in the Lek at 5 m^3/s. Extractions along the Waal, the Neder-Rijn, and the Lek are larger than average, and correspond to a decade with no rain and very high evaporation.

The value of water used according to each of the four tactics is presented in the table. Note that at a Rijn flow of 1000 m^3/s, the salt wedge has hardly any effect, and so there is no benefit from combating it. At lower Rijn flows, however, the benefits of the tactics increase. The value of 1 m^3/s reduction in the withdrawal at Gouda is always considerably greater than the value of exercising any of the other tactics by the same amount.

Table 4.3 shows the value of water for the same four Rijn flows. In this table, however, the withdrawal at Gouda is assumed to be 35 m^3/s, the maximum possible rate given the present infrastructure. The greater withdrawal at Gouda implies a smaller flow in the Nieuwe Maas, so that a Rijn flow of 1000 m^3/s is no longer enough to eliminate the effect of the salt wedge (a flow of nearly 1050 m^3/s is required). In addition, more salt is taken into the midwest by the larger withdrawal, so that the damage done is larger. Otherwise, the pattern is much the same as with the smaller withdrawal. The value of water is higher for lower Rijn flows, and reducing the withdrawal at Gouda is the tactic offering the highest value of water in combating the salt wedge.

4.10.2. Flushing at Locks with Widely Spread Salt Intrusion

In this section, we will estimate the value of water used to flush salt into the sea that intrudes through the locks Kornwerderzand, Den Oever, IJmuiden, Volkerak, Philipsdam, and Kreekrak. These are the locks from which intruding salt has the opportunity to spread widely throughout the MSDM network. Later, in Sec. 4.10.4, we will address the question of the value of water for flushing locks at which salt intrusion is locally confined.

We conclude that flushing these locks does not result in a significant reduction in salt damage to agriculture. (Recall that we have earlier concluded that salt damage to other users must be expressed in terms of water quality standards, or ignored.) Unless it is necessary to achieve the standards, therefore, MSDM will elect not to flush these locks.

4.10.2.1. Method of Estimation. The method we use to estimate the value of flushing water at these locks is essentially the same as the method used to estimate the value of water for combating the salt wedge. First, note that each of these locks is located at a node with storage. We divide the change in the amount of salt that intrudes by twice the volume of water stored at the node, to obtain the effect of flushing on the average salt concentration in storage during the decade. (Dividing by the volume alone would yield the effect on the concentration at the end of the decade.) We then multiply the change in the average concentration by the value of glasshouse crops supplied from the node, and by their salt sensitivity. The result is further multiplied by the equivalent exposure time of salt at the node. Nearly all the data needed to carry out this program can be found in Table 4.1.

4.10.2.2. Den Oever and Kornwerderzand. Both the locks at Den Oever and Kornwerderzand are located at the node IJSLAKES. The greatest effect flushing can have at either of these locks is to reduce salt intrusion by 0.0237 kg for each cubic meter of water used for flushing. This figure assumes that all factors--e.g., traffic intensity, technical and managerial tactics employed, etc.--combine to make flushing as effective as possible. If we apply our formula to this datum, therefore, we will obtain a maximum value for flushing. The result is 0.00000000019 Dfl/m^3, a value certainly too small to be noticed.

IJSLAKES is the primary source of water for various other nodes, in particular NORHOLL, FRIELAND, and GRONETAL. Glasshouse crops supplied from these nodes will also benefit somewhat from an improvement in IJSLAKES quality. However, even if they benefit to the same degree as the glasshouse crops supplied directly from IJSLAKES, the value of flushing water is still only 0.00000000045 Dfl/m^3.

The reason this value is so low is that the salt intruding at Den Oever and Kornwerderzand is assumed to mix with such a large volume of water that the effect on the salt concentration is negligible. In reality, the intruding salt remains near the locks for a substantial time, producing a much larger local effect on concentration. A model that took into account these local concentration gradients might conceivably yield a different result.

4.10.2.3. IJmuiden. Salt intrusion through the locks at IJmuiden affects the salt concentration at the nodes HAL+IJMU and AMSTEDAM on the Noordzeekanaal. No glasshouse crops are supplied from either of these nodes, nor does water flow from either node to any third node. Thus, our formula estimates the effect of flushing at IJmuiden on salt damage to be zero.

This estimate ignores several possible benefits of flushing. First, the brackish water in the Noordzeekanaal contributes to salt seepage in the parts of North Holland bordering the canal. A lower salt concentration in the canal would mean less salt in the seepage water.

Second, water from the Noordzeekanaal sometimes intrudes through the Oranjesluis (link 38 ORANJESL) into the IJSLAKES. If the results of the previous section are accepted, this has an infinitesimal effect on salt damage. In reality, however, it has a large, though very localized, effect on the salt concentration. Whether this local effect is the cause of any agricultural or other damage is a question we cannot address here.

Third, the salt intrusion rate through the locks at Spaarndam depends on the salt concentration in the Noordzeekanaal. We discuss the value of water used to flush the lock at Spaarndam in Sec. 4.10.4 below. Flushing the locks at IJmuiden offers another means to reduce salt damage at Spaarndam, since flushing at IJmuiden will reduce the salt in the Noordzeekanaal available for intrusion at Spaarndam. However, it requires many times more water to achieve a given effect on salt damage at Spaarndam by flushing at IJmuiden than by flushing Spaarndam directly.

4.10.2.4. Volkerak. Until DLTALAKE is made fresh, the Volkerak locks will permit salt to intrude into LOWRIVER. According to the lock simulation model (Vol. IX), the maximum effect flushing can have at the Volkerak locks is to reduce salt intrusion by 0.988 kg per cubic meter of water. Applying our formula yields a value of flushing of only 0.000000006 Dfl/m^3.

As with Den Oever and Kornwerderzand, a different estimate might be obtained if the situation near the Volkerak locks were modeled in greater detail.

4.10.2.5. Philipsdam and Kreekrak. Once DLTALAKE becomes fresh, the Volkerak locks will no longer serve as a route for salt intrusion, but the Philipsdam and Kreekrak locks will. However, the designs and situations of these locks are such that the tactic of changing the flushing rate is not considered reasonable. At Philipsdam, any flushing beyond the minimum lock loss necessary to pass ships through the locks (10 m^3/s) would cause, it is expected, unacceptable damage to the saltwater ecology on the other side. At Kreekrak, the design is such that the salt intrusion rate is well controlled even with minimal flushing (5 m^3/s).

4.10.3. Flushing at Nodes Supplying Glasshouse Crops

Flushing a node is the action of cleaning a pollutant out of a location by letting water flow through it. In PAWN, we define the flushing rate at a node to be the rate of water outflow from the node, where the outflow includes only the water carrying pollutants out of the node. A given flushing rate can, therefore, accompany a range of inflows of water from the remainder of the network, any deficit being made up by rain, and any surplus disposed of by evaporation, in order to maintain a balance between total inflows and outflows. We define the managerial tactic of increasing or decreasing the flushing rate at a node to involve changing the inflow of water from other nodes in the network and the outflow of water from the node in question by the same amount. Thus, flushing, considered as a managerial tactic, does not affect the water balance at the node.

4.10.3.1. Nodes At Which Flushing Reduces Salt Damage. In MSDM, the tactic of changing the flushing rate affects salt damage only at nodes which supply water to crops under glass. A list of these nodes can be found in Table 4.1. Not all of these nodes are much affected by flushing. For example, at nodes which have no storage, flushing generally has negligible effect, because the only water available for flushing is water already flowing into the node, and increasing its flow, where this is possible, will not reduce the chloride concentration at the node. At TWENMOND and TWENTEND, for example, the chloride concentration is that of the Rijn, regardless of the managerial strategy followed.

Of the remaining nodes from Table 4.1, 17 LOPIKWAR has such a small storage volume that very little flushing is needed to maintain its water at Rijn quality, and that little is supplied as a byproduct of providing the city of Utrecht with its water. Node 18 VECHT receives its water from drainage of groundwater. This water is sometimes highly contaminated with other pollutants, but its chloride content is much lower than that of the IJssel lakes, which would be the source of water for additional flushing. Node 31 SAMGRALI, which is located on the Maas, suffers no salt seepage and hence requires no flushing to reduce salt damage.

Thus, the only nodes at which we will consider flushing as a potentially worthwhile tactic for reducing salt damage are the nodes with storage in areas subject to seepage of salt water. These nodes are 1 FRIELAND, 10 IJSLAKES, 11 NORHOLL, 21 MIDWEST, 25 LOWRIVER, and 26 DLTALAKE.

4.10.3.2. Method of Estimation. A rough idea of the effect of flushing at a node with storage can be obtained from the following argument. Suppose flushing is increased by an amount dQ at a given node. The extra inflow of water will bring additional salt into the node, at the concentration that exists in the inflowing water. The extra outflow of water will carry with it salt at the concentration in storage, which we take to be the average concentration in storage during the decade. The average concentration will, of course, depend to some degree on the amount of extra flushing, dQ, so we write:

$$dp_{in} = dQ * (c_{in} - (\bar{c} + d\bar{c}))$$

where dp_{in} = the net change in salt (pollutant) inflow due to the extra flushing, in metric tons per decade;

dQ = the change in the flushing rate, in millions of cubic meters per decade;

c_{in} = the salt concentration in the extra inflowing water, in mg/l;

$\bar{c}$ = the average salt concentration in storage during the decade, with no extra flushing, in mg/l;
$d\bar{c}$ = the change in the average salt concentration in storage during the decade, due to the extra flushing, in mg/l.

As we have mentioned, the average salt concentration is influenced by extra flushing rate. Indeed, if this were not so, there would be no point to flushing. A fair approximation to the change in the average concentration is:

$$d\bar{c} = \frac{dp_{in}}{2V}$$

That is, the net change in salt inflow divided by the volume V will equal the ultimate change in concentration. The factor "2" in the denominator accounts for the fact that the average change in the concentration is only half the ultimate change. Combining the two expressions, we find that:

$$d\bar{c} = \frac{dQ*(c_{in} - \bar{c})}{2V + dQ} \qquad (4.6)$$

From this point, we estimate the value of water for extra flushing in the same way as we estimated the value of water for combating the salt wedge. We multiply the change in the average concentration by the value of glasshouse crops supplied from the node, and by their salt sensitivity. The result is further multiplied by the equivalent exposure time of salt at the node. Nearly all the data needed to carry out this program can be found in Table 4.1. The only missing data are the differences in concentration between the water at the node and the water that might be used for extra flushing.

4.10.3.3. A Test of the Estimation Formula. As a test of our formula for estimating the value of flushing water, we look to runs of the Distribution Model. The only runs that seem directly useful are those used to estimate the expected value of a pipeline to bring low-salinity Maas water to Delfland. The benefits from this pipeline were estimated from four Distribution Model runs, each assuming a different pattern of rainfall, evaporation, and river flows during the various decades in the year. (Such a pattern is called an *external supply scenario*.) Depending on the probabilities one assigns to these external supply scenarios, one can estimate the average annual benefits from the pipeline to be between 20 and 50 Dflm/yr. The same procedure estimates that the average amount of water the pipeline actually supplied was between 1.1 and 1.3 m^3/s, so the average value of the water was 0.57 to 1.19 Dfl/m^3. For further discussion of how the benefits of the pipeline were estimated, see Vol. II.

This high value stems from the following three facts: (a) the open water volume of Delfland is only 20 million cubic meters, so relatively little flushing water could cause a relatively large improvement; (b) the value of glasshouse crops grown in Delfland is extremely large (1173 Dflm/yr); and (c) the difference between the chlorinity of Maas water and the present chlorinity of the boezem and ditch water of Delfland is relatively large (135 mg/l). To apply our formula, we need one other datum, the equivalent exposure time of salt in Delfland. But the average outflow of water from Delfland is 7.2 m^3/s, which according to App. B of Vol. VA implies an equivalent exposure time of 7 decades. Our formula applied to these data yields a value of water for flushing of 0.94 Dfl/m^3. This is quite comparable with the value obtained from the Distribution Model runs.

4.10.3.4. Flushing Node 1 FRIELAND. It is difficult to establish a typical salt concentration in an area as large as that represented by the node FRIELAND. Some waterways represented by this node are near the Waddenzee, and suffer chloride concentration as high as 1500 mg/l. But most of the surface water in FRIELAND contains chloride at concentrations between 150 and 250 mg/l throughout most of the year. Indeed, these concentrations are significantly lower than those found in the IJsselmeer, which is represented by the MSDM node IJSLAKES, and which is the only reasonable source of flushing water for FRIELAND. Accordingly, increases in the flushing rate of FRIELAND should increase, rather than reduce, salt damage.

Our modeling efforts indicate that the waters of FRIELAND will typically have chloride concentrations between 25 and 50 mg/l lower than IJSLAKES. According to our formula, then, each cubic meter of additional flushing should increase salt damage by 0.0005 to 0.0009 Dfl.

4.10.3.5. Flushing node 10 IJSLAKES. It is not possible to flush the entire amount of storage represented by the MSDM node IJSLAKES. This node already receives as much water as it can receive, under most circumstances. But IJSLAKES actually consists of several separate lakes, the largest two of which are the IJsselmeer and the Markermeer. There has been much thought given to the possibility of using IJsselmeer water to flush the Markermeer (Vol. XVI).

According to Ref. 4.10, the water in the Markermeer is more saline than the water in the IJsselmeer only in the first four months of the year. Later, after the IJsselmeer has received the more highly saline summer discharges of Rijn water, the salinity of the IJsselmeer comes to exceed that of the Markermeer, and flushing would be counterproductive. In the early months of the year, however, the difference in salt concentration may be as high as 100 mg/l in a favorable direction. By April, the difference, while still favorable, has dropped to 25 mg/l.

Before we can apply our formula to estimate the value of flushing, we must estimate the volume and equivalent exposure time for the Markermeer alone. Table 4.1 contains these figures only for the entire IJSLAKES. But most of the water discharged from IJSLAKES flows through the Markermeer, so that the reduction in volume is compensated by a proportionate reduction in equivalent exposure time. The result is that flushing the Markermeer may have a value of from 0.0002 Dfl/m^3, for a concentration difference of 25 mg/l, to 0.0008 Dfl/m^3, for a concentration difference of 100 mg/l.

4.10.3.6. Flushing Node 11 NORHOLL. The waterways represented by the node NORHOLL consist of two quite different systems of canals. One of them, the Schermerboezem (a boezem is a main drainage and water supply canal) is only five to ten mg/l more saline than the Markermeer in summer, from which it takes its water. It is from this canal system that all of the glasshouse crops in North Holland are supplied, and if we apply our formula we find the value of flushing water is only 0.0005 to 0.001 Dfl/m^3.

However, the Schermerboezem does not directly supply water to agriculture. Instead, it supplies water to a network of small ditches, which in turn supply agriculture. The ditches supplied from Schermerboezem contain perhaps 50 mg/l more salt than the boezem, and if the full 60 mg/l concentration difference is used, then our formula would estimate the value of additional flushing to be 0.006 Dfl/m^3.

PAWN has not modeled the boezem and ditch systems found in the low-lying areas of the Netherlands in sufficient detail to determine which concentration difference is the appropriate one to use, or whether an intermediate concentration is more nearly correct, or whether the entire approach we have used here is too ridiculously oversimplified to be credible. The correct approach would undoubtedly consider two flushing rates, the rate at which the boezems are flushed with water from an outside source, and the rate at which the ditches are flushed with water from the boezems. It is our opinion, however, that the correct value of water for flushing boezem and ditch systems is quite low, as our crude approach estimates, except perhaps in very special (and probably locally unique) conditions. By their actions, the Dutch appear to agree with this assessment, since flushing is the first use of water that they reduce in times of shortage.

One exception to this conclusion occurs in the Wieringermeer polder in North Holland. This polder is served by a boezem and ditch system quite separate from the Schermerboezem. The polder has by far the highest rate of seepage of salt water in the Netherlands, a fact which explains the chloride concentration of 1000-1500 mg/l in the ditches. At the same time, the chloride concentration in the boezem is hardly above that in the IJsselmeer, about 225 mg/l.

The very high ditch concentration is sufficient to cause salt damage to crops grown in the open air, so that flushing the ditches might have some benefit even though they supply no glasshouse crops. We can calculate the benefits of a reduction in the ditch chloride concentration if we make several assumptions. First, suppose the only crops to suffer damage are those irrigated from surface water. Second, assume that the salt concentration in the root zones of these crops is 1.5 times the ditch concentration. Then, we can calculate that a reduction in ditch chloride by 1 mg/l for one decade would reduce the salt damage by about 1600 Dfl. Other necessary data are the ditch volume (9.725 million cubic meters), and the nominal flushing rate (2.6 m^3/s). These numbers lead to an equivalent exposure time (see App. B of Vol. VA) of about 9.3 decades. If we take the concentration difference to be 1000 mg/l, our formula estimates a value of flushing of 0.77 Dfl/m^3.

Note, however, it is not enough to flush the boezems in order to achieve this substantial benefit. The boezems already have low chloride concentration. It is necessary, rather, to flush the ditches. At Wieringermeer, as at most polders, this is difficult to do, since the works for taking in water are physically colocated with the works for pumping water out. To operate both works simultaneously would only result in circulating the already contaminated ditch water.

4.10.3.7. Flushing Node 21 MIDWEST. The node MIDWEST, from which such a large value of glasshouse crops is supplied with water, would appear at first sight to be a likely candidate for increased flushing. However, the salt seepage in this area is relatively modest, so that the rain is sufficient to maintain the salinity in both the boezems and the ditches below that of the Rijn. Since the Rijn is the sole potential source of water for flushing, additional flushing will not usually be beneficial.

There are occasions, however, when the chloride concentration in the Rijn drops as much as 50 mg/l below that in the midwest boezems and ditches. This occurs when the Rijn flow is high, generally in the first quarter of the year. When such a concentration difference is observed, our formula estimates that the value of additional flushing is 0.16 Dfl/m^3.

4.10.3.8. Flushing Node 25 LOWRIVER. The only way to increase the flushing rate of LOWRIVER is to increase the flow through the Haringvlietsluizen. This diverts Rijn water that would otherwise flow into the North Sea via the Rotterdamse Waterweg. However, the water stored at LOWRIVER is a mixture of Rijn and Maas water, and is less saline than Rijn water alone. Flushing, therefore, is not worthwhile at LOWRIVER.

4.10.3.9. Flushing Node 26 DLTALAKE. The MSDM node DLTALAKE represents the Zoommeer, which is now a saltwater basin but is soon to be dammed off and made fresh. It will also serve to represent the Grevelingenmeer, which is presently a brackish lake, if it were decided to turn it fresh. Once made into a freshwater basin, DLTALAKE will take its supply of fresh water from LOWRIVER, through what is now the Volkerak lock complex. Simulations carried out by WW

District Southwest (unpublished) suggest that the chloride concentration difference between LOWRIVER and DLTALAKE might regularly reach 100 mg/l. But even if the difference were five times as large, that is 500 mg/l, our formula estimates that the value of water used to flush DLTALAKE would only be 0.005 Dfl/m³. This is because very few glasshouse crops would be supplied from this node.

4.10.4. Flushing of Local Locks

In MSDM, discretionary flushing of local locks occurs at Den Helder, Parksluis, Harlingen, and Spaarndam. In each case, the reason for flushing at greater than the required minimum rate is to reduce the local damage to crops due to salt intrusion through the lock. From the results of Sec. 4.7.2, it is possible to estimate the reduction in damage for any increase in the flushing rate. Selected estimates are shown in Table 4.4 below.

The values shown in Table 4.4 assume that no technical or managerial tactics are employed at any of the locks in order to reduce the salt intrusion by means other than flushing. In addition, the methodology in Sec. 4.7.2 uses upper-bound estimates of the crops at risk from local salt intrusion. Thus, the estimated values from Table 4.4 are overestimates of the actual value of water for flushing local locks. Nevertheless, it appears that small increases from the minimum flushing rates at these locks may have significant benefits. However, the value of additional increments quickly declines to essentially zero.

Table 4.4

VALUE OF WATER FOR FLUSHING LOCAL LOCKS

	Flushing Rate Increase (m^3/s for one decade)		Value of Water
Location	From	To	(Dfl/m^3)
Den Helder	2	4	0.042
	4	6	0.0005
	6	8	0.000006
Parksluis	2	4	0.115
	4	5	0.081
	5	5.6	0.002
	5.6	7.5	0.0005
Harlingen	1	2	0.083
	2	4	0.002
	4	5	0.00002
Spaarndam	3	5	0.0092
	5	7.5	0.00066
	7.5	10	0.000037

NOTES

1. Although time and resource limitations prevented it, we would have liked to compare the MSDM approach for estimating damage to crops under glass with the approach used in DISTAG (see Vol. XII). It is desirable that the two models agree, so that results are consistent throughout the PAWN study, and so that strategies developed using MSDM will perform well when simulated by the Distribution Model (which uses DISTAG as a subroutine). Agreement between the models, however, would not demonstrate the correctness of either one. Only agreement with real-world observations would do that, and we were unable to locate the appropriate data.
2. One of our reviewers has drawn our attention to some runs of the Distribution Model (DM) in which the effects of the salt wedge on damage to Midwest agriculture appear to persist considerably longer than the 12.8 decade equivalent exposure time used here. On this basis he suggests that we may be underestimating the damage.

   His suggestion may indeed be true, but the DM runs he refers to don't prove it. For example, if we have underestimated the amount of water stored at node 21 MIDWEST, e.g., by leaving out root zone or subsoil moisture incorrectly, the equivalent exposure time would increase in proportion to the stored volume. Since in Eq. (4.5) we multiply by the one and divide by the other, our estimate of damage would remain the same.

   It is true, however, that Eq. (4.5) is a very crude estimate of damage to Midwest agriculture due to the salt wedge, and that the true damage could be considerably different (although we are not prepared to guess whether higher or lower). To form a more reliable estimate, however, would have required much more time and resources than PAWN could afford to allocate to this part of the analysis.

REFERENCES

4.1. "Kwaliteitsonderzoek in de Rijkswateren," (Water Quality Research in Dutch State Waters), published yearly prior to 1971, quarterly thereafter, by Rijksinstituut voor Zuivering van Afvalwater (State Institute for Wastewater Treatment, RIZA), Lelystad, Netherlands.
4.2. Unpublished information from Rijksinstituut voor Zuivering van Afvalwater (State Institute for Wastewater Treatment, RIZA) concerning salt management of Rijn branches in the Netherlands, May 1975 (PAWN file DW-072).
4.3. MR-120 (unpublished PAWN memorandum), "Some Comments and Suggestions Relating to MW-087, '--Main Management Issues--Salinity,'" February 1978.

4.4. Langfristiges Arbeitsprogramm (LZP)--Gesamtfassung, Internationale Kommission zum Schutze des Rheins Gegen Verunreinigung (Long-Range Work Program--Global Version), International Commission for the Protection of the Rhine from Pollution, 2 volumes, Postfach 309, 5400 Koblenz, Germany.
4.5. MR-132 (unpublished PAWN memorandum), "Costs of Water Quality Changes (TDS and TH) for Municipal and Industrial Users," March 1978.
4.6. Richter, C. O., and A. MacLean, "Salt Taste Threshold of Humans," Am. J. Physiol., 126:1, 1939.
4.7. Lockhart, E. E. et al., "The Effect of Water Impurities on the Flavor of Brewed Coffee," Food Research, 20:598, 1955.
4.8. De Bestrijding van de Verontreiniging van het Oppervlaktewater: Indicatief Meerjarenprogramma 1975-1979, Ministerie van Veerkeer en Waterstaat (The Combat Against Surface Water Pollution in the Netherlands: Prospective Multi-annual Program 1975-1979), Ministry for Transport and Public Works, The Hague, 1975.
4.9. Unpublished information from the Rijksinstituut voor Naturbeheer (State Institute for Nature Management, RIN) reviewing water quality standards for the protection of the natural environment (PAWN file DW-472).
4.10. Unpublished information from Zuiderzeewerken (ZZW), consisting of salinity contour maps of the IJssel lakes, February 1978 (PAWN file DW-327).

Chapter 5

WATER MANAGEMENT AND THERMAL POLLUTION

5.1. INTRODUCTION

5.1.1. Definition and Reasons for Concern

Thermal pollution occurs when heat is discharged into a waterway. This increases the temperature of the water above the point it would have reached in the absense of heat discharges, the so-called natural or background temperature. The measure of thermal pollution used by the Dutch is the excess temperature, which is the rise in temperature of a water body above its natural temperature.

There are several reasons to be concerned about thermal pollution. For example, increases in temperature will decrease the amount of oxygen available in the water by reducing the solubility of oxygen, and by accelerating the decay of organic matter, a process that uses oxygen. Organic matter is often present due to discharges of sewage into the water. A low oxygen level may kill fish and other organisms.

A second reason is that at higher temperatures, many aquatic organisms are more sensitive to toxic materials than at lower temperatures. Since Rijn water, at least, contains significant concentrations of toxic materials, this consideration is relevant to the Netherlands.

A third reason is that the metabolism, growth rate, and reproduction of aquatic organisms are greatly influenced by temperature. Especially the larvae of certain species may be extremely sensitive to temperature changes. To species well adapted to an environment, the changes we are considering are well within their tolerance; but such a change can cause just enough additional mortality among the larvae to eliminate a poorly adapted species from a habitat.

5.1.2. PAWN's Approach to Modeling Thermal Pollution

In PAWN, we have constructed two related models for dealing with thermal pollution. In each of them, we consider that electric power generating plants are the sole sources of heat discharges in the Netherlands. In both models, heat discharged into a waterway raises the excess temperature at the point of discharge. In addition, the heat is carried downstream with the flowing water, which loses the heat at some rate along the way. Thus a discharge of heat raises the temperature at points downstream of the discharge, although to a lesser extent than at the discharge point itself.

In both models, we can reduce the excess temperature at any point by reducing the heat discharges at or upstream of that point. This is done by reducing the amount of power generated by plants discharging their waste heat there. But such a reduction in the power generated at one point must be compensated by an increase in the power generated elsewhere, if the demand for electric power is to be satisfied. How this redistribution is carried out is reported in detail in Vol. XV, which describes one of our models called EPRAC (Electric Power Reallocation and Cost). EPRAC carries out this objective assuming that the amounts of water available to cool the power plants are precalculated by the Distribution Model (DM) Vol. XI.

In MSDM, we use a similar model to accomplish a similar, but more general objective. In MSDM, we relax the assumption that the amounts of cooling water are given and try to calculate what might be done to lower the cost of power generation by changing the water distribution as well as the power generation schedule.

5.1.4. Organization of This Chapter

In MSDM the thermal pollution problem has six elements: the calculation of excess temperatures at nodes of the MSDM network; thermal standards; heat discharges into MSDM nodes, largely by electric power generating plants; the demand for electric power; the capacities of the various parts of the electric power generating and distributing system in the Netherlands; and the objective of the problem, to minimize the cost of satisfying the demand. Sections 5.2 through 5.7 discuss each of these elements in turn. Section 5.8 assembles the elements into the overall MSDM thermal polluton problem.

Section 5.9 brings out various assumptions and complicating factors involved in the formulation of the MSDM thermal pollution problem, but not mentioned in earlier sections.

Sections 5.10 through 5.15 comprise an analysis of the relation between water management and thermal pollution.

5.2. CALCULATING EXCESS TEMPERATURES

As explained in detail in App. B of Vol. VA, our simple description of pollutant transport through the network can be expressed mathematically as a system of simultaneous linear equations. The unknowns in the equations are the pollutant concentrations at each node, which in the case of thermal pollution are excess temperatures. The coefficients of these unknown excess temperatures are calculated from the flows of water in the network links, and the amounts of water stored at the nodes. As explained in Chap. 8, whenever we make this calculation, we will know, at least provisionally, the flows of water in every link of the network. The right-hand sides of the equations are the discharges of the pollutant, which for thermal pollution is heat, into each node of the network. Mathematically, then, the calculation of excess temperatures in MSDM can be expressed as:

$$(5.1) \qquad \text{MATRIX} * \text{TEMP} = \text{HEAT}$$

In Eq. (5.1), HEAT is a vector of heat discharges expressed in megacalories per second (Mcal/s), one element of the vector corresponding to each node. TEMP is a vector of excess temperatures in degrees Celsius, likewise having one element for each network node. MATRIX is the matrix of coefficients in the simultaneous linear equations that relate the excess temperatures to the heat discharges. As mentioned above, the individual entries in the matrix depend on the water flows in network links, and water storage at nodes. Details of how these entries are determined can be found in App. B of Vol. VA.

5.3. THERMAL STANDARDS

The Dutch presently apply a standard of three degrees Celsius to the excess temperature in their major river branches (see Vol. XV, and App. E of Vol. VA for a discussion of the reasons behind the standard). These branches include, for example, the Maas, the Waal, the Neder-Rijn, the Lek, and the many river branches in the Lower Rivers area. Also included are the IJssel River, and the IJsselmeer, Markermeer, and Randmeren (border lakes). However, all canals and all non-state waters are excluded. For example, no thermal standard is applied to the Amsterdam-Rijnkanaal or the Noordzeekanaal.

In PAWN, we have adopted as our nominal assumption a thermal standard of three degrees (Celsius) excess temperature in all waterways. Even though the present law applies this standard only to the main rivers, we have investigated the effect of applying it as well to the large canals. As an excursion case, we have also investigated the effect of a less stringent standard of seven degrees on canals. For the sake of comparison, we consider the case where no thermal standards are applied anywhere.

Mathematically, we can express the requirement that the excess temperature conform to the standards as:

$$(5.2) \qquad \text{TEMP} \leq \text{STD}$$

In Inequality (5.2), STD is a vector of thermal standards having one element for each node of the MSDM network. Note that the standard can be different at different nodes.

5.4. HEAT DISCHARGES

The heat discharged into each node of the MSDM network may come from power plants generating electricity in the Netherlands. Or it may come from the Rijn or the Maas, which have been heated above their natural temperatures by heat discharges in other countries.

5.4.1. Excess Temperatures in the Rijn and Maas at the Borders

Power plants discharge heat into the Rijn and the Maas all along their lengths, before they enter the Netherlands. It is not known by how much their temperatures are raised above natural levels by the time they reach the borders, so we have been forced to make assumptions regarding their excess temperatures.

For the Rijn, we have assumed that for flows smaller than 1500 cubic meters per second at Lobith, the excess temperature at the border is three degrees Celsius. For Rijn flows greater than 1500 m^3/s, we assume a constant excess heat flow of 4500 Mcal/s across the German border, enough heat to raise the temperature of 1500 m^3/s of water by three degrees. For the Maas, we have adopted the assumption of M. Hofstra [5.1]. Both the excess temperatures and the corresponding excess heat loads in the Rijn and Maas at their respective borders are shown in Figs. 5.1a and 5.1b. These assumptions are discussed in Vol. XV. Their influence on our results is discussed in App. E of Vol. VA.

In PAWN, we are most interested in periods of low flow in both the Rijn and Maas, well below the maximum flows at which the rivers have excess temperatures of three degrees at their borders. On the Rijn, even low flows are high enough that the water has little time to shed its excess heat before reaching nodes on main branches of the Rijn (the Waal and the IJssel). Thus, assuming that the excess temperature of the Rijn at the border is three degrees eliminates most of the capacity of the Rijn to absorb waste heat from Dutch sources.

On the Maas, the problem is not so severe, since at low flows the many weirs on the Maas impound the water for rather long periods. The water thus has an opportunity to shed virtually all of the excess heat in the water at the border.

5.4.2. Heat from Power Plants

We have assembled inventories of Dutch power plants for the years 1976 and 1985 (see App. E of Vol. VA). Each power plant contains one or more generating unit, which serve as the basic units of our analysis. Each generating unit has a gross capacity, expressed as the maximum amount of power (in megawatts, abbreviated MW) it is capable of generating. A certain amount of the capacity must be reserved for in-plant uses, however, and an additional five percent must be allocated to serve as a spinning reserve. We call the remainder the effective capacity of the unit.

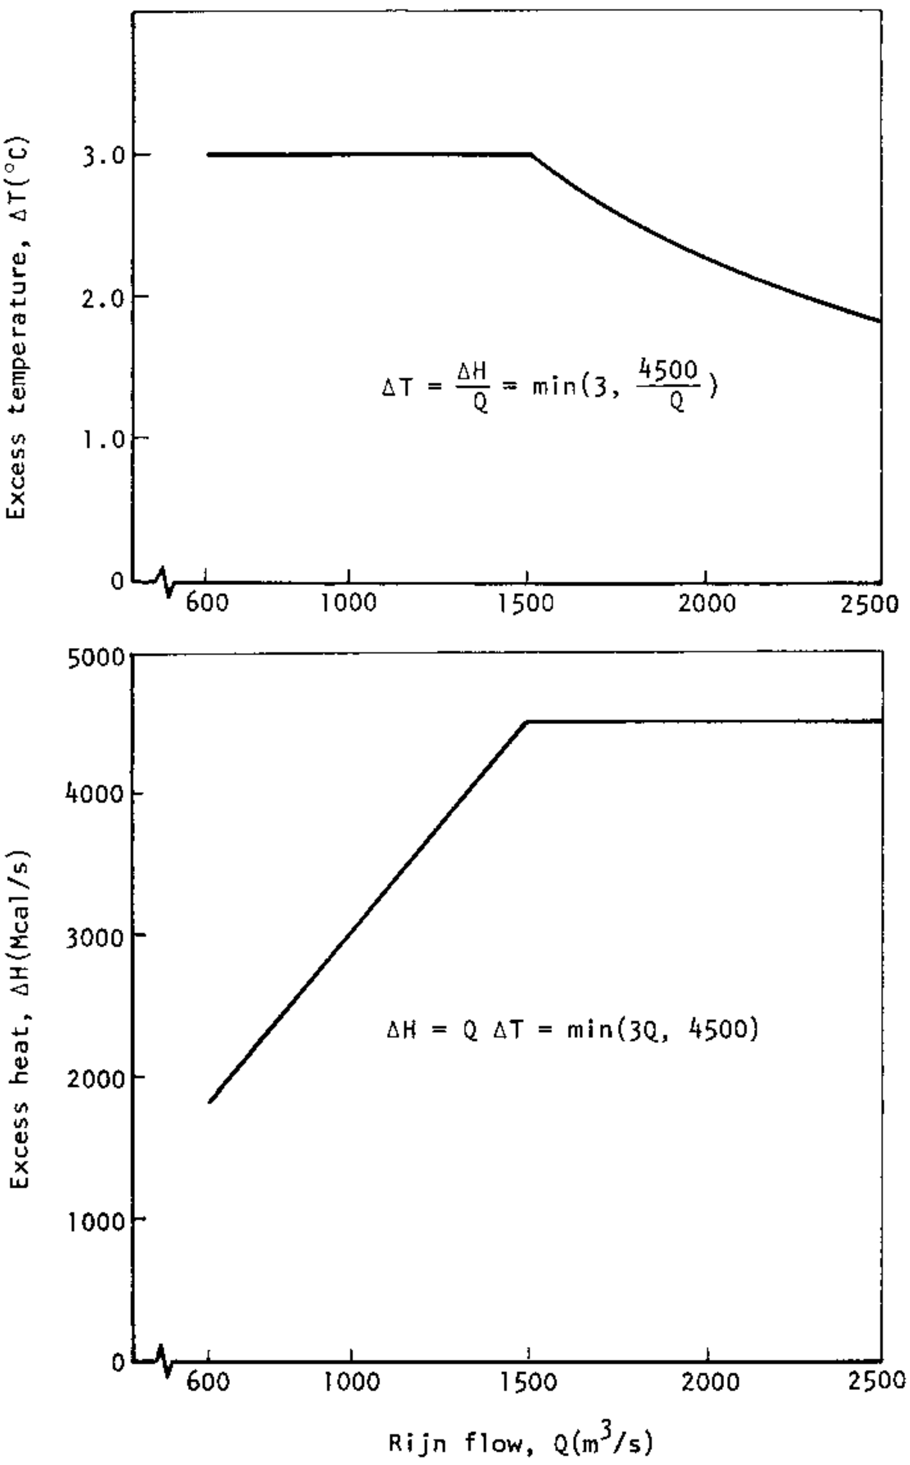

Fig. 5.1a--Excess heat and temperature in the Rijn at Lobith

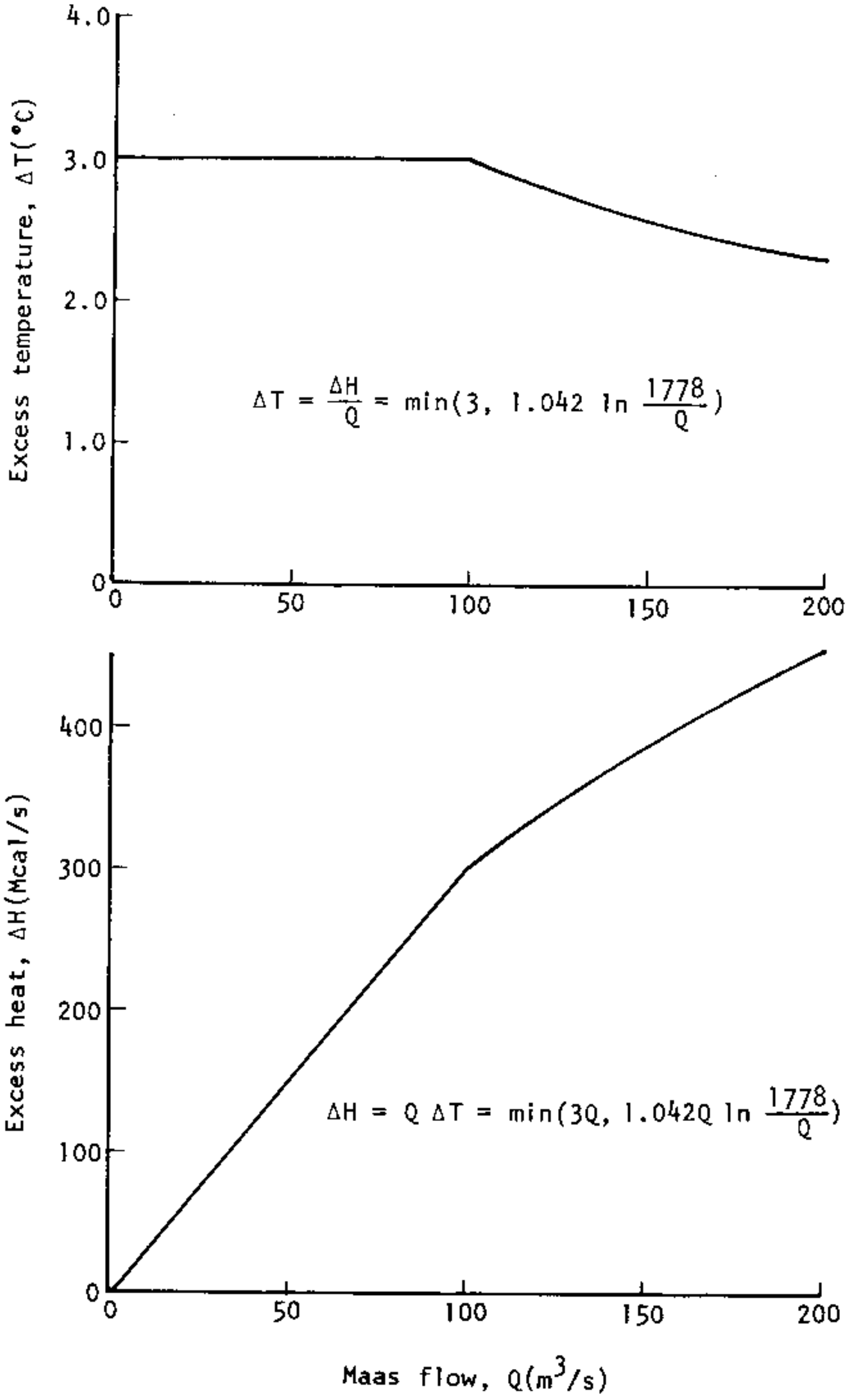

$$\Delta T = \frac{\Delta H}{Q} = \min(3,\ 1.042\ \ln \frac{1778}{Q})$$

$$\Delta H = Q\ \Delta T = \min(3Q,\ 1.042Q\ \ln \frac{1778}{Q})$$

Fig. 5.1b--Excess heat and temperature in the Maas at Eijsden (from M. Hofstra)

Generating units of different designs have different efficiencies. This means that to generate a given amount of power, different units will require different amounts of fuel, and will discharge different amounts of waste heat. Some of the generating units do not affect the excess temperatures at MSDM nodes because they discharge their waste heat to the air, or into the North Sea, or into waterways not represented in the MSDM network. But other generating units do affect the excess tmperatures we calculate, and for each of these, we have assigned a node into which its waste heat is discharged. In addition, there are a few generating units in the 1985 power plant inventory at which the use of cooling towers is optional. These units may or may not discharge their waste heat into the network.

5.4.3. Calculation of Heat Discharges

Now we define:

DISO = a vector of heat discharges having one element for each generating unit that discharges heat into a MSDM node, <u>not</u> including units with optional cooling towers. It is expressed in Mcal/s.

DIST = a vector of heat discharges, in Mcal/s, having one element for each unit with an optional cooling tower.

LOCO = a matrix with one row for each node and one column for each generating unit that discharges heat into a node, not including units with optional cooling towers. The column corresponding to a generating unit will consist of all zero entries, except for an entry of '1' in the row corresponding to the node at which the generating unit discharges its waste heat. In other words, this matrix specifies the LOCation of each generating unit's heat discharge.

LOCT = a matrix of the same kind as LOCO, but whose columns correspond to units with optional cooling towers.

EXOG = a vector of exogenously supplied heat discharges in Mcal/s, having one element for each node. In our model, all but two of the elements are zero. The two nonzero elements correspond to the nodes at which the Rijn and Maas enter the Netherlands, and the values of those elements are the amounts of heat imported from other countries by those rivers.

With these definitions, we can express the heat discharges at the various nodes as:

(5.3) HEAT = LOCO*DISO + LOCT*DIST + EXOG

5.4.4. Units with Optional Cooling Towers

In the 1976 inventory of power plants, all power plants with cooling towers must use them, as they have no alternative means of cooling. They can never, therefore, discharge their waste heat into the network. In the 1985 inventory, however, there are units at two power plants that can be operated either with or without towers. The units are found at power plants Amer (at node 33 GERTRUID) and Maasbracht (at node 29 LINNE). Without the cooling tower, such a unit discharges its waste heat into a network node. Using the cooling tower eliminates heat discharges into the network, but it also reduces the efficiency of the unit by about eight percent, and raises its fuel consumption by the same percentage.

In MSDM these units are represented as two variables, whose sum is required not to exceed one. The first variable represents operating the unit without cooling towers. When this variable equals one, the unit is producing power at full capacity, and discharging heat into the network at the maximum rate. The second variable represents operating the units with cooling towers. When the second variable equals one, the unit is again producing power at full capacity, but all the waste heat is being discharged through the cooling tower to the air. No heat is discharged into the network. When the two variables sum to less than one, the deficit may be thought of as leaving the unit partly idle. For example, when the sum is zero, no power is generated, no cost is incurred, and no heat is discharged.

These two variables can be averaged in any proportions. Thus, it is possible to operate the unit at one-half capacity without cooling towers by setting variable one to 1/2, and variable two to zero. In the model it is also possible to operate the units at full capacity, but cool them using partly the towers and partly Maas water. In real life, this is evidently not possible, but we do not trouble to rule out this possibility from the model.

Mathematically we can describe our representation of these units as follows. Let:

DMAX = a diagonal matrix of maximum heat discharges, in Mcal/s, having one row and column for each unit with an optional cooling tower. On the diagonal of the matrix are placed the heat discharges that would be observed if the units were operated at full capacity without cooling towers. All off-diagonal elements are zero.

F1, F2 = two vectors, each having one element for each unit with an optional cooling tower. Each element is a

> fraction between zero and one. Elements of F1 correspond to the first, and elements of F2 the second, of the two variables describing units of this kind.

Although we will not require F2 until a later section, it is convenient to introduce it now. Given these definitions, we can write:

$$\text{DIST} = \text{DMAX}*\text{F1} \tag{5.4}$$

5.5. THE DEMAND FOR ELECTRICITY

We have partitioned the Netherlands into five power regions for electricity (Vol. XV), and we have estimated both the supply and the demand by region in both 1976 and 1985. MSDM requires that the demand in each region be satisfied from one of four sources: (1) from generating units located in the region that discharge heat into a node, not including units with optional cooling towers; (2) from units located in the region with optional cooling towers; (3) from units in the region that do not discharge heat into the MSDM network; and (4) from units located outside the region, whose output must be transmitted into the region. In the fourth case, the power is transmitted through a grid, an operation entailing transmission losses that we assume to be one percent of the amount transmitted.

5.5.1. Power Regions

Each of the five power regions is composed of one or more provinces. Each region has its own demand for power. In addition, each unit is located in one of the regions, and hence can supply power to that region without transmission. That is, we assume that power generated anywhere within a region can be transmitted to any other point in the same region at no cost and with no restrictions. However, we do consider that there will be a loss of one percent of the power transmitted between regions.

In Table 5.1, we identify the provinces making up each region, and the supply and demand by region for both 1976 and 1985. Note that for both 1976 and 1985, the generating capacity in each region exceeds the demand by a factor of nearly two. This does not mean that no power will be transmitted from one region to another. Some of the apparent excess capacity in each region consists of units that are relatively expensive to operate. It is worthwhile to transmit power from some regions to others so that these more expensive units may stand idle, while cheaper units in other regions produce the needed electricity.

Table 5.1

SUPPLY AND DEMAND BY POWER REGION
(Megawatts)

		1976 Scenario		1985 Scenario	
Region	Provinces	Supply Capacity	Demand	Supply Capacity	Demand
I.	Friesland Groningen Drenthe Overijssel	2367	1056	2833	1627
II.	Noord-Holland	1940	958	2133	1654
III.	Zuid-Holland	3027	1308	3722	1600
IV.	Limburg Noord-Brabant Zeeland	3456	1935	4650	2668
V.	Utrecht Gelderland IJsselmeer Polders	2114	884	2415	1766
	Netherlands	12902	6141	15753	9315

5.5.2. Relation Between Heat Discharges and Power Generation

If we choose in our models to operate a generating unit at full capacity, we have data to determine the heat discharge that will occur. Because we have no data on part-load characteristics of Dutch power plants, we have assumed that if we operate a unit at part load, the ratio of heat discharged to power generated will be constant. For example, at 60 percent of capacity, the waste heat discharge will be 60 percent of its full-capacity value. This assumption is discussed in App. E of Vol. VA.

An important consequence of this assumption is that there is a constant ratio between the heat discharged from a unit and the amount of power it generates. This is important because we have been expressing the level of operation for units that discharge heat into the network (not including units with optional cooling towers) in terms of the amounts of heat that they discharge, rather than in the more natural terms of the amount of power produced. But this assumption says that we can obtain the second by merely multiplying the first by a constant.

5.5.3. The Demand Constraints

Given the above assumptions, it is possible to write a set of linear equations, one for each power region, that express the requirements that the demand in each region be satisfied. We define:

DEMAND = a vector of demands for electricity, having one element for each region. It is expressed in megawatts (MW).

EPERD = a matrix with one row for each power region and one column for each generating unit that discharges heat into a node, not including units with optional cooling towers. In the column corresponding to a generating unit, the only nonzero element will be the element in the row corresponding to the power region in which the generating unit is located. The value of this element is the amount of energy generated per unit of heat discharged into the network; hence the name of the matrix, Energy PER Discharge. Its units are (MW-s)/Mcal.

EMAX = a matrix with one row for each power region and one column for each generating unit with an optional cooling tower. In the column corresponding to a generating unit, the only nonzero element will be the element in the row corresponding to the power region in which the generating unit is located. The value of this element is the amount of energy generated, in MW, when the unit is operated at full effective capacity. This matrix does for units with optional cooling towers what EPERD does for units without cooling towers, that discharge heat into the network.

ENOD = a vector of amounts of energy in MW generated by units that do not discharge heat into the network, having one element for each power region; hence the name, Energy with NO Discharge.

GRID = a matrix describing the transmission grid. It has one row for each power region, and one column for each link in the grid. A column contains only two nonzero elements, a "-1" in the row corresponding to the region <u>from</u> which the link transmits power, and a "0.99" in the row corresponding to the region <u>to</u> which the link transmits power. The links in the model are considered to be unidirectional. Thus, if power can be transmitted in both directions between two regions (although not simultaneously), we must include two links in the model, one transmitting in each direction.

TRAN = a vector of amounts of power transmitted in MW, having one element for each link of the transmission grid.

Given these definitions, we can write:

(5.5) EPERD*DISCH + EMAX*(F1 + F2) + ENOD + GRID*TRAN = DEMAND

In Eq. (5.5), the four sources of power for each region correspond to the four terms to the left of the equals sign. The matrix product EPERD*DISCH is a vector with one element per region, showing the amount of power generated in that region by units discharging heat into the network, not including units with optional cooling towers. The product EMAX*(F1 + F2) is a vector showing the power generated in each region by units with optional cooling towers. They may operate completely without their cooling towers ($F1>0$ and $F2=0$), or completely with them ($F1=0$ and $F2>0$), or partly with and partly without them ($F1>0$ and $F2>0$). ENOD, by definition, is a vector showing, for each region, the amount of power generated by units in that region that do not discharge heat into the network. Finally, the matrix product GRID*TRAN is a vector showing the net amount of power transmitted into that region, calculated as the amount transmitted into the region, after transmisson losses, minus the amount transmitted out, before losses.

5.6. CAPACITIES

As mentioned earlier, each generating unit has a capacity. In addition, each link in the transmission grid has a capacity. The capacities of the units that discharge heat into the MSDM network, excluding those with optional cooling towers, are expressed as the amounts of heat discharged when they are generating their maximum amounts of power, because these units are represented in MSDM by the vector DISO of heat discharges. The capacities of units with optional cooling towers are embodied in the matrix EMAX above; their maximum heat discharges appear in the matrix DMAX. To ensure that these maxima are not exceeded, it is only necessary that F1+F2 not exceed one. The capacities of units that do not discharge heat into the network, and the capacities of the transmission links, are expressed in the more natural units of megawatts. The actual capacities we use in MSDM can be found in App. E of Vol. VA. We define:

DCAP = a vector of heat discharge capacities, in Mcal/s, having one element for each generating unit that discharges heat into the network.

ECAP = a vector of effective generating capacities, in MW, having one element for each generating unit that does not discharge heat into the network.

TCAP = a vector of transmisson capacities, in MW, having one element for each link in the transmission grid.

Then we can write:

$$\text{DISO} \leq \text{DCAP} \tag{5.6}$$

$$\text{F1} + \text{F2} \leq 1 \tag{5.7}$$

$$\text{ENOD} \leq \text{ECAP} \tag{5.8}$$

$$\text{TRAN} \leq \text{TCAP} \tag{5.9}$$

5.7. THE OBJECTIVE: MINIMIZE GENERATING COST

The thermal pollution problem in MSDM has the objective of minimizing the cost of generating electricity and delivering it to the regions that demand it. There are only three sources of generating cost: (1) the cost of generating power at units which discharge heat into the network, excluding units with optional cooling towers; (2) the cost of generating power at units with optional cooling towers; and (3) the cost of generating power at units which do not discharge heat into the network. The first of these terms is a function of the elements of the vector DISCH, the second a function of the elements of the vectors F1 and F2, and the third a function of the elements of the vector ENOD.

We assume that the fuel cost is the only cost that varies with power output; all other costs (e.g., maintenance, repair, personnel) must be paid regardless of how much power is generated by the unit. Thus, the marginal cost of power generation is taken to be equal to the fuel cost (for discussion, see Vol. XV and App. E of Vol. VA).

If we know the price of the fuel, we can calculate the fuel cost of operating a unit at its effective capacity. Because we have no data on part-load characteristics of Dutch power plants, we have assumed that if we operate a unit at part load, for example at 60 percent of capacity, the cost of fuel will be 60 percent of its full-capacity value. The latter assumption is equivalent to saying that the cost of generating power is a linear function of the levels at which the various generating units are operated. The former assumption offers a means for calculating the coefficients of the operating levels of the different generating units in the linear cost function.

Finally, we assume the cost of transmission does not depend on the amount of power transmitted--i.e., the marginal cost of transmission is zero--but the fact that one percent of transmitted power is lost means that extra power must be generated--at some cost--to make up the losses. Thus, transmission is not without an indirect effect on cost. Now define:

CPD = a vector of costs per unit heat discharge, in (Dflm/dec)/(Mcal/s), having one element for each generating unit that discharges heat into the network, excluding units with optional cooling towers.

CNOCT = a vector of costs, in Dflm/dec, having one element for each unit with an optional cooling tower. Each element is the cost of operating the corresponding unit at full effective capacity without using the cooling tower; hence the name, Cost with NO Cooling Tower.

CWCT = a vector of costs, in Dflm/dec, having one element for each unit with an optional cooling tower. Each element is the cost of operating the corresponding unit at full effective capacity, this time using the cooling tower; hence the name, Cost With Cooling Tower.

CPE = a vector of costs per unit of electric power generated, in Dflm/(MW-dec), having one element for each generating unit that does not discharge heat into the network.

Then we can write the cost as:

$$\text{Cost} = \text{CPD*DISO} + \text{CNOCT*F1} + \text{CWCT*F2} + \text{CPE*ENOD} \tag{5.10}$$

5.8. THE THERMAL POLLUTION PROBLEM

We are now in a position to state the thermal pollution problem in MSDM mathematically. It takes the form of a linear program, with Eqs. (5.1) through (5.9) serving as constraints, and with the objective of minimizing the cost function (5.10) discussed in the previous section. We write the thermal problem below as Problem (5.11), in which we have combined Eqs. (5.1), (5.3) and (5.4) in order to eliminate the intermediate variables HEAT and DIST.

$$(5.11)\ \begin{cases}
\text{Minimize:}\ \ \text{CPD*DISO} + \text{CNOCT*F1} + \text{CWCT*F2} + \text{CPE*ENOD} & \\
\\
\text{Subject to:} & \\
\quad \text{MATRIX*TEMP} - \text{LOCO*DISO} - (\text{LOCT*DMAX})\text{*F1} & = \text{EXOG} \\
\quad \text{TEMP} & \leq \text{STD} \\
\quad \text{EPERD*DISO} + \text{EMAX*(F1+F2)} + \text{ENOD} + \text{GRID*TRAN} & = \text{DEMAND} \\
\quad \text{DISO} & \leq \text{DCAP} \\
\quad \text{F1} + \text{F2} & \leq 1 \\
\quad \text{ENOD} & \leq \text{ECAP} \\
\quad \text{TRAN} & \leq \text{TCAP}
\end{cases}$$

Problem (5.11) is a standard linear program (e.g., see Ref. 5.2). The variables are TEMP, DISO, F1, F2, ENOD, and TRAN. The variables, other than TEMP, represent activities that cannot be operated in reverse. Thus, generating units cannot be operated to absorb power and produce fuel, and power cannot be transmitted backwards to yield as output 101 percent of the amount of power input. Hence, they are restricted to be nonnegative. (TEMP, the vector of excess temperatures, will naturally be nonnegative if the coefficient matrix MATRIX is properly constructed (see App. B of Vol. VA)).

5.9. ASSUMPTIONS AND COMPLICATIONS

In formulating Problem (5.11), we have made a number of simplifying assumptions in addition to those mentioned above. In addition, there are some complicating factors that we have taken into account but which, to avoid confusion, we have not mentioned heretofore. In this section, we will briefly mention each new assumption and each complication, if only to reassure the reader that they have not been neglected. Further discussion can be found in App. E of Vol. VA.

5.9.1. The Heat Transfer Rate

The rate of heat transfer from water to air is used in calculating the matrix of coefficients, MATRIX (see Sec. 5.2), that relate heat discharges to excess temperatures. But this rate is extremely variable. Depending on the wind velocity, the surface temperature of the water, the relative humidity, and the size and shape of the water body, it can easily vary by an order of magnitude [5.3]. Our Dutch colleagues have advised us to adopt a heat transfer rate in the middle of the range, which we have done for most of our cases. We have, however, investigated the sensitivity of our results to this parameter by making some calculations with an extremely low, pessimistic value. These sensitivity results can be found in App. E of Vol. VA.

5.9.2. Time variation

The demand for electricity varies by time of day. Since electricity is inconvenient to store, it must be generated at the moment of consumption, and therefore the generation rate must also vary by time of day. In MSDM, we have assumed that each generating unit is assigned a constant fraction of the total demand during each hour of the day, instead of determining a new generating schedule for each demand level. This permits us to calculate only one generating schedule instead of many. (How much error this might introduce into our results is discussed in App. E of Vol. VA.)

If demand varies by time of day, then so must the heat discharge. The thermal standard applies to the peak instantaneous excess temperature. The straightforward application of Eq. (5.1), using the average heat discharges from each power plant as the right-hand sides, would calculate the average excess temperatures during the decade. To calculate the actual peak excess temperatures would require us to expand the method to calculate the hour-by-hour excess temperatures at each MSDM network node, taking into account the delay time between the discharge of heat from a power plant and its arrival at a downstream plant. Instead, we use the peak heat discharge rates in place of the average rates. The price we pay for this simplification is that the peak excess temperatures are overestimated. Conversely, under this assumption the thermal standards have the maximum possible effect on power plant operation.

However, considering that the average demand--and hence average heat discharge--are approximately 80 percent of the peak values (see App. E of Vol. VA), we do not think this simplification introduces an unduly large error.

5.9.3. Unusual Cooling Effects

For most power plants, one calculates the peak heat discharge rate into a MSDM network node by multiplying the power generated, expressed as a fraction of effective capacity, by the waste heat discharged at capacity. This calculation is not correct for all power plants, however. Power plants at which unusual cooling effects occur are located at nodes 33 GERTRUID, 2 B'GUMMER, 3 GRONETAL, 14 DIEMEN, 15 AMSTEDAM, and 16 HAL+IJMU.

The Amer power plant, which discharges its waste heat into node 33 GERTRUID, is affected by tidal action. Thus, when the flow in the Maas is low, the water flows back and forth past the plant instead of remaining nearly stagnant, so that the waste heat is dispersed within a rather large volume of water. In effect, tidal action provides the Amer plant with a large cooling pond.

Generating units discharging heat at 2 B'GUMMER, 3 GRONETAL, 14 DIEMEN, 15 AMSTEDAM, and 16 HAL+IJMU, divert their cooling water from one of the links of the MSDM network, discharge their heat into a diversion channel called a cooling circuit, and then return the heated water to the main waterway. As it flows through the cooling circuit, the water sheds some of its excess heat, so that the discharge of heat into the network is less than the discharge of heat out of the generating unit.

At nodes 33 GERTRUID, 2 B'GUMMER, and 3 GRONETAL, we have represented these phenomena by including links in the MSDM network that loop from each of these nodes back to the same nodes. These links are the links AMER REC, 46 B'GM REC, and 50 H-H REC (see Table 2.2). At the other nodes, 14 DIEMEN, 15 ANSTEDAM, and 16 HAL+IJMU, the cooling circuits appeared to us to have relatively minor effects, and hence were ignored.

5.10. OPTIMAL THERMALLY UNCONSTRAINED GENERATING SCHEDULES

The cost of generating power to meet the demand will be lowest if the thermal standards are ignored. That is, the optimal thermally unconstrained generating schedule (which is the solution to Problem (5.11) with the vector of standards, STD, made very large) will provide a lower bound on the cost of power generation. If the standards are enforced, the cost of meeting the demand for electric power will rise, because some inexpensive units will have to be shut down to avoid violating the standards at some nodes, and their output will have to be replaced by more expensive power generated elsewhere. In the case that no change in managerial tactics occurs, the rise in cost due to enforcing the standards will be a maximum. Costs between these two extremes may be attained if the appropriate managerial tactics are employed. The reductions in cost thus attained can serve as measures of the value of water used according to the tactics.

In this section and in the several that follow it, we investigate the relation between managerial tactics and the cost of power generation. This section determines the optimal thermally unconstrained generating schedule, and identifies the nodes at which this schedule leads to violations of the three-degree standard. Later sections investigate what may be done at each of these "hot" nodes to meet the standards, and what is the cost of doing so.

5.10.1. The 1976 Optimal Unconstrained Schedule

As mentioned above, we have considered two different inventories of power generating units in MSDM, one for 1976 and one for 1985 (see App. E of Vol. VA). In Fig. 5.2 we summarize some features of the 1976 inventory. In 1976, the Netherlands had a total effective generating capacity of approximately 13000 megawatts (MW), distributed among over 30 power plants. The largest of these plants, the Amer plant (at node 33 GERTRUID), had an effective capacity of over 1700 MW, and could discharge nearly 500 Mcal/s of waste heat into the Maas when operating at full effective capacity. Another large plant, Velsen (at node 16 HAL+IJMU), had an effective capacity of 990 MW, and discharged 342 Mcal/s into the Noordzeekanaal when operating at effective capacity. Most of the remaining large plants were located where their waste heat was discharged into the Rotterdam harbor, or directly into the North Sea.

To construct this figure, we calculated the marginal fuel cost for each generating unit in the 1976 inventory, expressed in Dutch cents per kilowatt-hour. We ordered the units according to their marginal costs, from the lowest cost to the highest, and accumulated the effective capacities. The result was the lower curve in Fig. 5.2. Next, we integrated the lower curve to obtain the total fuel cost. To express the total cost as millions of guilders per decade, we assumed units could be operated at just over 80 percent of effective capacity (the ratio of average to peak hour demand), for 243.333 hours per decade (one thirty-sixth of the hours in a year). This results in the upper curve of Fig. 5.2.

From the figure, we see that it will cost 46.65 Dflm/decade to supply the fuel needed to generate enough electricity to supply the summer, 1976 demand. This result ignores thermal standards; the least expensive units are used without regard for where or how large their heat discharges may be. It also ignores the possible need to transmit power from one part of the Netherlands to another. As discussed earlier, we have included transmission effects in the thermal problem (5.11); as demonstrated in App. E of Vol. VA, it has very little effect on our results.

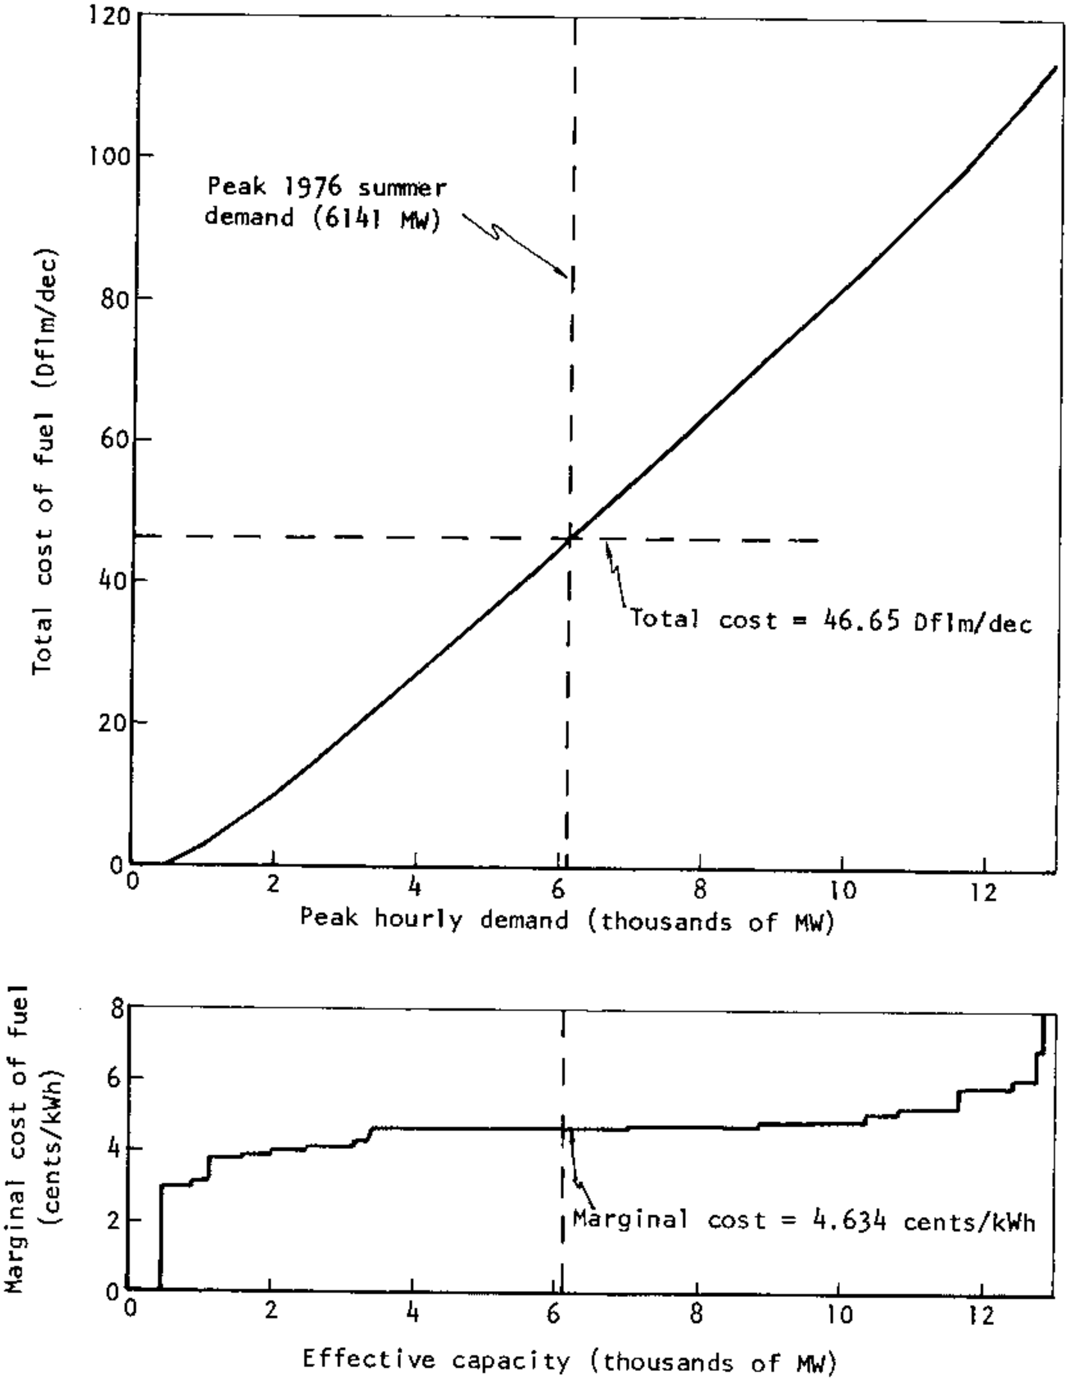

Fig. 5.2--Fuel cost for power generation (1976 inventory of power plants)

The most interesting feature of the lower curve is that, except for about 500 MW capacity of nuclear plants (assumed to have zero fuel cost), the curve is remarkably flat. In fact, if we look at the distribution of demand over the hours of a typical summer day, we see that between the minimum hourly demand (3471 MW) and the maximum (6141 MW), the marginal fuel cost varies only from 4.601 to 4.634 cents/kWh. This implies that the upper curve is almost perfectly linear between the minimum and maximum hourly demands.

5.10.2. The 1985 Unconstrained Optimal Schedule

By 1985, the Netherlands expects to have expanded their generating capacity to about 16000 MW. At the time of this writing, much of the planned expansion had already taken place. For example, in 1979 the Amer plant's capacity had been expanded to an effective capacity of 1900 MW, by a combination of installing new, more efficient units and retiring old, inefficient units. The potential for discharging waste heat from Amer was only marginally increased to 525 Mcal/s by this expansion. In addition, in 1977 and 1978 a plant at Maasbracht (at node 29 LINNE) was brought into production, at an effective capacity of 1203 MW. This plant discharges 344 Mcal/s waste heat when operating at effective capacity. The Velsen plant, however, is scheduled to remain the same in 1985 as it was in 1976. In Fig. 5.3, we summarize some features of the projected 1985 power plant inventory. We constructed this figure in the same way as Fig. 5.2. Clearly, the same remarks can be made about Fig. 5.3 as about the earlier Fig. 5.2. In 1985, because the demand for electricity has increased over the 1976 demand, the total fuel cost for meeting the demand has risen to 71.53 Dflm/decade. As before, this ignores transmission requirements, but as before, these have small impact.

5.10.3. Nodes Likely to Violate Thermal Standards

In PAWN, we considered a three-degree limit on excess temperature as the nominal thermal standard. In Tables 5.2 and 5.3 we show the excess temperature at each of the MSDM nodes at which power plants can discharge waste heat. Table 5.2 deals with the 1976 inventory of power plants and demand, while Table 5.3 deals with the 1985 inventory and demand. In the two tables, we assume the peak heat discharges at each node that characterize the optimal generating schedules discussed in Secs. 5.10.1 and 5.10.2, respectively. To calculate the excess temperatures, we have assumed a water distribution characteristic of a rather dry decade, with a Rijn flow of 951 m^3/s at Lobith, and a Maas flow of 34.2 m^3/s at Monsin (Belgium), of which 22.4 are reserved for Belgium. The flows in the links of the network reflect the managerial policies presently being used by the Dutch, as described in Vol. XI. Excess temperatures are separated into a component caused by local heat discharges directly into the node, and a component due to remote heat discharges, originally discharged into nodes farther upstream and carried down to the node in question before all the heat is lost to

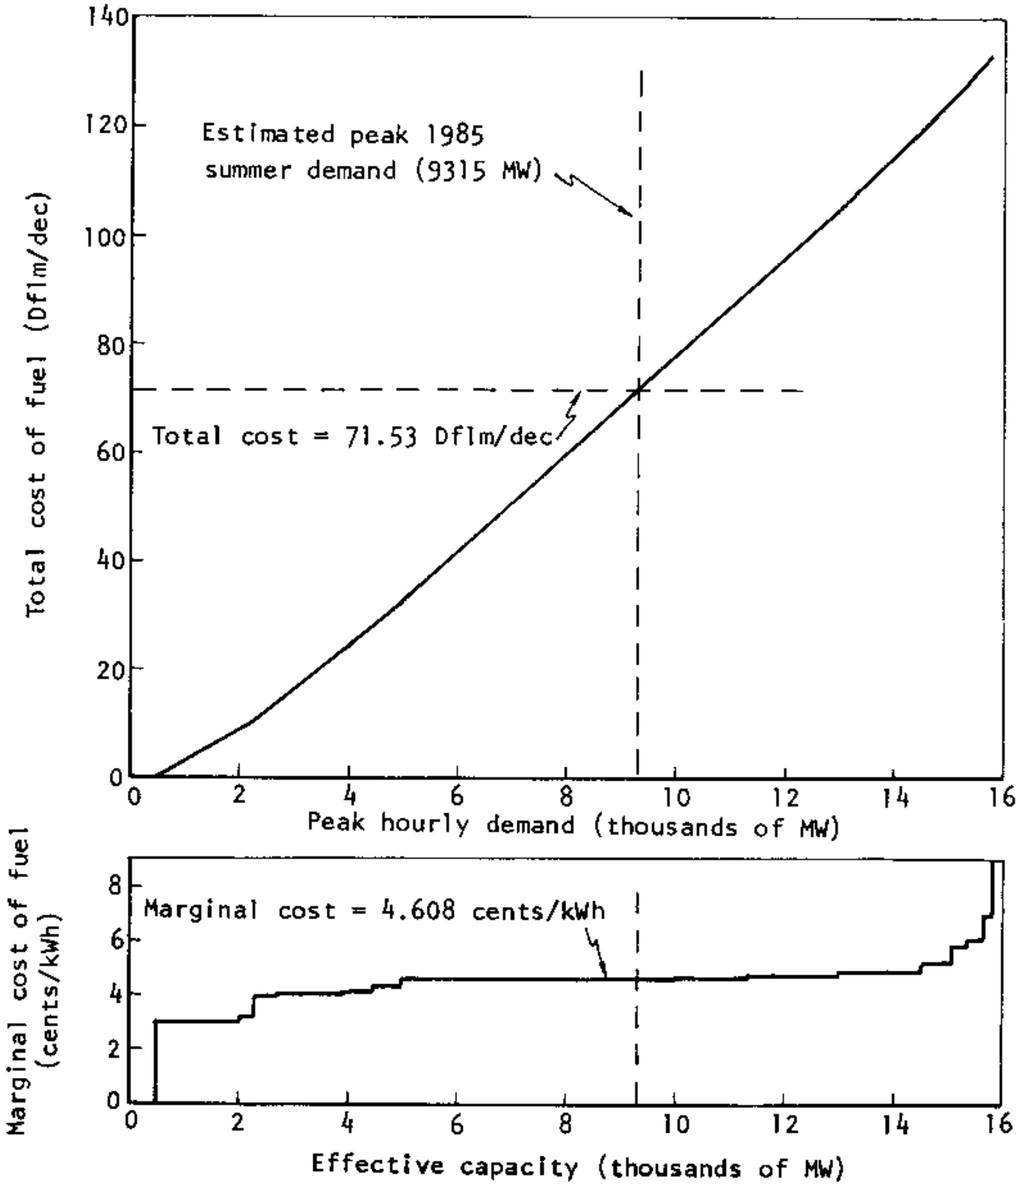

Fig. 5.3--Fuel cost for power generation (1985 inventory of power plants)

Table 5.2

EXCESS TEMPERATURES AT MSDM NETWORK NODES DURING A DRY DECADE
(1976 Scenario)

				<- Excess Temperatures -> (Degrees Celsius)		
MSDM Network Node	Peak Heat Discharge (Mcal/s) (a)	Flow (m^3/s) (b)	Cooling Factor (c)	Locally Caused (d)	Remotely Caused (e)	Total
GRONETAL	0.	3.9	0.55	0.	0.	0.
B'GUMMER	171.73	13.2	0.38	4.94	0.	4.94
IJSLEND	55.008	187.0	1.	0.29	1.6	1.89
NIJMEGEN	133.122	737.0	1.	0.19	2.7	2.89
IJSLAKES	137.915	-	0.	0.	0.	0.
UTR+MAAR	18.233	7.7	1.	2.37	0.1	2.47
HAL+IJMU	342.015	33.3	1.	10.26	0.51	10.77
DIEMEN	0.	11.7	1.	0.	0.11	0.11
AMSTEDAM	92.032	21.6	1.	4.26	0.	4.26
GOR+DOR	0.	587.0	1.	0.	1.38	1.38
GERTRUID	407.083	35.0	0.65	7.49	0.	7.49
ROER+BEL	30.591	34.0	0.43	0.39	0.	0.39
LINNE	0.	15.1	0.41	0.	0.03	0.03

(a) This is the peak heat discharge into each node under the optimal generating schedule with no thermal standard, as discussed in Sec. 5.10.1.

(b) Flows are total flows leaving the node via all links. The water distribution from which the flows were taken was generated by the Distribution Model, using rather low flows in the Rijn (951 m^3/s) and Maas (34.2 m^3/s), and using the RWS managerial strategy. This managerial strategy is discussed in Vol. XI; it is intended to approximate the present Dutch practice.

(c) The cooling factor is the factor by which the heat discharge must be multiplied to account for cooling circuits, tidal cooling at GERTRUID, or the surface area at a node with storage (see App. E of Vol. VA).

(d) The locally caused excess temperature is the product of the peak heat discharge and the cooling factor, divided by the flow. Performing this calculation using the numbers in the table may not yield the excess temperature given there due to rounding.

(e) The remotely caused excess temperature is due to heat discharges upstream of the node in question that have not had time to decay.

Table 5.3

EXCESS TEMPERATURES AT MSDM NETWORK NODES DURING A DRY DECADE (1985 Scenario)

				<- Excess Temperatures -> (Degrees Celsius)		
MSDM Network Node	Peak Heat Discharge (Mcal/s) (a)	Flow (m^3/s) (b)	Cooling Factor (c)	Locally Caused (d)	Remotely Caused (e)	Total
GRONETAL	0.	3.9	0.55	0.	0.	0.
B'GUMMER	171.73	13.2	0.38	4.94	0.	4.94
IJSLEND	152.461	187.0	1.	0.82	1.6	2.42
NIJMEGEN	251.001	737.0	1.	0.35	2.7	3.05
IJSLAKES	137.915	-	0.	0.	0.	0.
UTR+MAAR	101.362	7.7	1.	13.13	0.1	13.23
HAL+IJMU	342.015	33.3	1.	10.26	0.80	11.06
DIEMEN	0.	11.7	1.	0.	0.15	0.15
AMSTEDAM	144.691	21.6	1.	6.70	0.	6.70
GOR+DOR	159.6	587.0	1.	0.27	1.45	1.72
GERTRUID	322.902	35.0	0.65	6.00	0.	6.00
ROER+BEL	0.	34.0	0.43	0.	1.44	1.44
LINNE	343.972	15.1	0.41	9.36	0.03	9.39

(a) This is the peak heat discharge into each node under the optimal generating schedule with no thermal standard, as discussed in Sec. 5.10.2.

(b) Flows are total flows leaving the node via all links. The water distribution from which the flows were taken was generated by the Distribution Model, using rather low flows in the Rijn (951 m^3/s) and Maas (34.2 m^3/s), and using the RWS managerial strategy. This managerial strategy is discussed in Vol. XI; it is intended to approximate the present Dutch practice.

(c) The cooling factor is the factor by which the heat discharge must be multiplied to account for cooling circuits, tidal cooling at GERTRUID, or the surface area at a node with storage (see App. E of Vol. VA).

(d) The locally caused excess temperature is the product of the peak heat discharge and the cooling factor, divided by the flow. Performing this calculation using the numbers in the table may not yield the excess temperature given there due to rounding.

(e) The remotely caused excess temperature is due to heat discharges upstream of the node in question that have not had time to decay.

the air. The unusual cooling effects mentioned in Sec. 5.9.3 are taken into account in the "cooling factors" found in these tables. The heat discharge from the power plant, which is reported in "peak heat discharge" column, must be multiplied by the cooling factor for the appropriate node in order to obtain the heat discharge into the node.

There are two conclusions to draw from these tables. First, there are only six nodes (2 B'GUMMER, 13 UTR+MAAR, 15 AMSTEDAM, 16 HAL+IJMU, 29 LINNE, and 33 GERTRUID) at which the thermal standard of three degrees is violated in either 1976 and/or 1985. (We will ignore the trivial violation at 22 NIJMEGEN.) The locations of these nodes are shown in Fig. 5.4. Second, at each of these nodes, the violation is almost entirely due to local heat discharges. Heat brought to these nodes from remote, upstream locations does not contribute much to the excess temperature. Power plants at all other nodes have ample water for cooling under all circumstances, or they are so inefficient that they are not used in the optimal generating schedules.

It is important to note that of these six nodes, only two (GERTRUID and LINNE) are presently subject to a three-degree thermal standard. The other four (B'GUMMER, UTR+MAAR, AMSTEDAM, and HAL+IJMU) are not subject to any formal standard at all, although the excess temperatures of the water discharged from power plants at these nodes cannot be permitted to rise too far for reasons of efficiency.

The flows shown in Tables 5.2 and 5.3 are not unusual. The Maas frequently has low flows, resulting in thermal problems at LINNE and GERTRUID. The flows at nodes B'GUMMER, UTR+MAAR, AMSTEDAM, and HAL+IJMU are all under the control of the water manager, since they occur on canals. The flows shown for these nodes are typical of summer decades with little or no rain. Thus, the excess temperatures shown in Tables 5.2 and 5.3 should not be considered rare occurrences.

5.11. POTENTIAL MEASURES TO COMPLY WITH THE THERMAL STANDARDS

There are two methods for controlling the heat discharges from a power plant into the water. The first is to reduce the power generated at a plant. In order to meet the demand for electricity, such a reduction must be compensated by an increase at a plant in some other location, where the heat discharge is more acceptable. This shifting of load between power plants is not a popular pastime among power plant operators. While it can be done, it must be preplanned and executed carefully, and it involves start-up and shutdown costs.

In addition, in PAWN we have assumed that the power companies try to generate their electricity using the more efficient units, and that the less efficient units stand idle. As described above, under the 1976 scenario, the most costly unit it proved necessary to use in the optimal generating schedule has a marginal fuel cost of 4.634

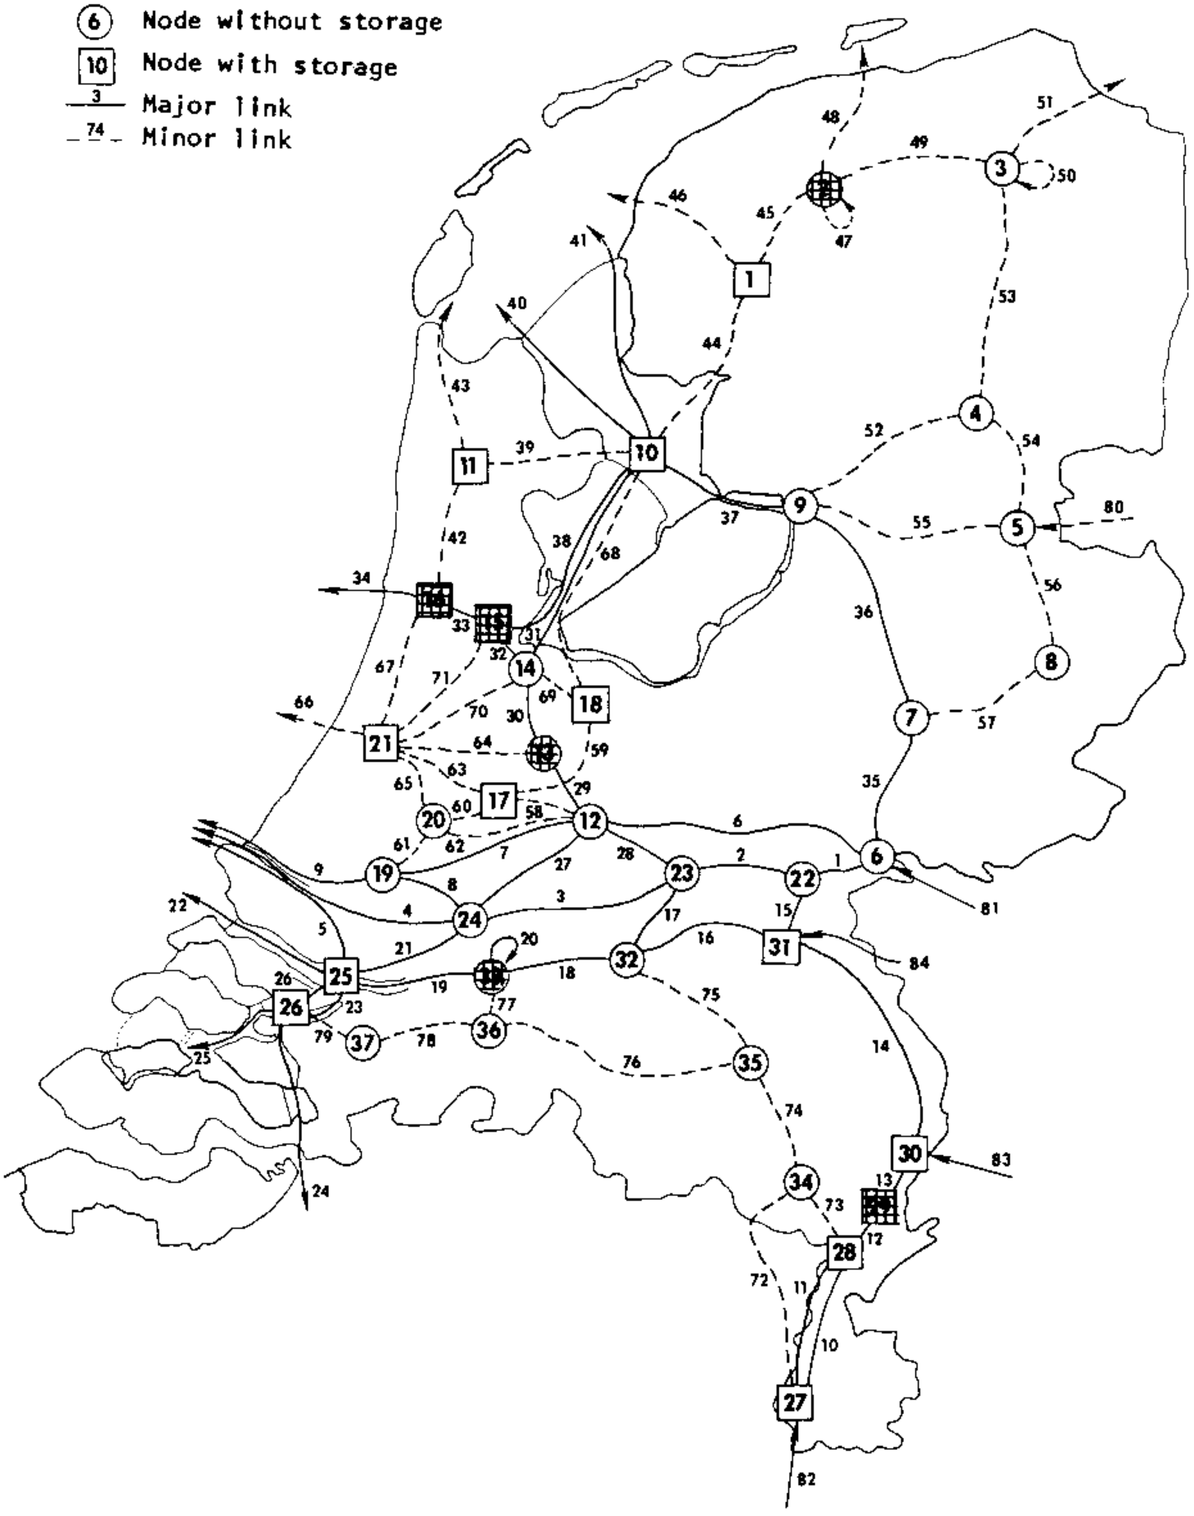

Fig. 5.4--Nodes violating the three-degree standard under the optimal generating schedule

cents/kWh. For 1985, the highest cost required is 4.608 cents/kWh. If transmission costs are included, the highest cost required in either scenario is 4.634 cents/kWh (see App. E of Vol. VA).[1] There are many units with essentially this efficiency, not all of which need be used to meet the demand. Thus, when a load is shifted from a unit used in the optimal schedule, it can always be replaced at a marginal fuel cost of 4.634 cents/kWh. However, the unit from which the load is shifted will generally have a lower marginal fuel cost (never a higher fuel cost). Thus, shifting to other units will generally involve an increase in the cost of generating electric power, in addition to start-up and shutdown costs.

The second method to reduce heat discharge into the water is to discharge the waste heat to the air instead. This is done by means of cooling towers. The power plants Maasbracht and Amer have cooling towers in the 1985 inventory, although the Amer plant does not have enough tower capacity to discharge all the waste heat generated. Towers could be installed at other plants. This method is no more popular with power plant owners and operators than the first method because cooling towers cost money to build and operate.

Instead, the power companies would prefer the water managers to guarantee that enough water will flow past each power plant to dilute the waste heat to an acceptable temperature. This could mean either diverting water, or relaxing the thermal standards at some times on some waterways. Generally, diverting water will not have a significant direct cost, but it may mean depriving another user of water and thus imposing the cost of cooling power plants on other sectors.

5.12. MEASURES AT GERTRUID

At the power plants that discharge heat at GERTRUID, the Amer and Donge plants, the three measures mentioned in Sec. 5.11 may all be employed to meet the three-degree standard. Units can be shut down, and their output replaced by other units elsewhere. The cooling towers at Amer can be used to discharge some of the waste heat into the air instead of the water. (This is an option only with the 1985 scenario, since the cooling towers did not exist in 1976.) Using the cooling tower raises the cost by about 0.12 cents/kWh. Finally, the weir ponds (stuwpanden) in the Maas can be partly drained of water to augment the flow past the Amer plant. The third measure can only work for one or two decades before the water levels in the weir ponds drop to a point that shipping finds critical, and cannot in any case provide very much cooling capacity. Thus, at present there is little that managerial tactics can do to help cool the plants on the Maas when nature decides not to provide the water. However, PAWN has considered a technical tactic, the St. Andries Connection, by which Waal water could be brought to the Maas and could help cool the node GERTRUID.

5.12.1. Optimal Unconstrained Heat Discharge at GERTRUID

The Amer and Donge plants discharge their waste heat at node 33 GERTRUID. In the optimal generating schedule for 1976 (see Sec. 5.10.1), every unit at Amer and Donge with a marginal fuel cost less than or equal to 4.634 cents/kWh is used to full capacity during the hour of peak demand. This results in a peak generating rate at these two plants of 1434 MW, and a peak heat discharge of 407 Mcal/s.

In the optimal generating schedule for 1985 (see Sec. 5.10.2), the units at Amer and Donge with a marginal fuel cost less than or equal to 4.608 cents/kWh are used to generate 1189 MW. In the process they discharge 323 Mcal/s of waste heat. This generation rate is 728 MW below the capacity of the units at GERTRUID with marginal fuel costs in the range mentioned.

In the 1985 schedule, we made an arbitrary choice to use units at Amer at part load, and to use units at Maasbracht at full capacity. The Maasbracht units are in the same region (region IV) as the Amer and Donge units, and have the same marginal fuel cost (4.608 cents/kWh) as some of the Amer units. The capacity of the Maasbracht units is 1203 MW. We could as easily have used the Amer units at their full capacity of 787 MW (an increase of 728 MW), and only used 475 MW of the capacity of the Maasbracht units (a decrease of 728 MW).

As long as no thermal standard is enforced, it makes no difference (in our models) which of these units we use at full capacity, and which we operate at part load. When the three-degree standard is applied, however, it may be possible to shift some load from Maasbracht to Amer at no cost instead of shifting it somewhere else to a unit with a higher marginal fuel cost (after 4.608 cents/kWh, the next highest marginal fuel cost is 4.634 cents/kWh). In MSDM we do consider this possibility, but in the discussion below we do not. Therefore, what follows will slightly overstate the cost of the thermal standard. However, the cost difference between 4.608 and 4.634 cents/kWh is so small that our error can hardly be significant.

5.12.2. Reduction of Heat Discharges

In Fig. 5.5, we summarize the effect of low Maas flows at GERTRUID on the power plants discharging into that node in the 1976 inventory of generating units. At these low flows, the heat discharge from the Amer and Donge plants must be reduced if the three-degree standard is to be met at GERTRUID. Figure 5.5a shows the optimal amount of power generated there consistent with the three-degree standard, as a functon of the Maas flow. At Maas flows higher than 127 m^3/s, all units at these plants with a marginal fuel cost no larger than 4.634 cents/kWh can be used at full effective capacity. At lower Maas flows, less capacity can be used, and units elsewhere must pick up the load at a higher cost.

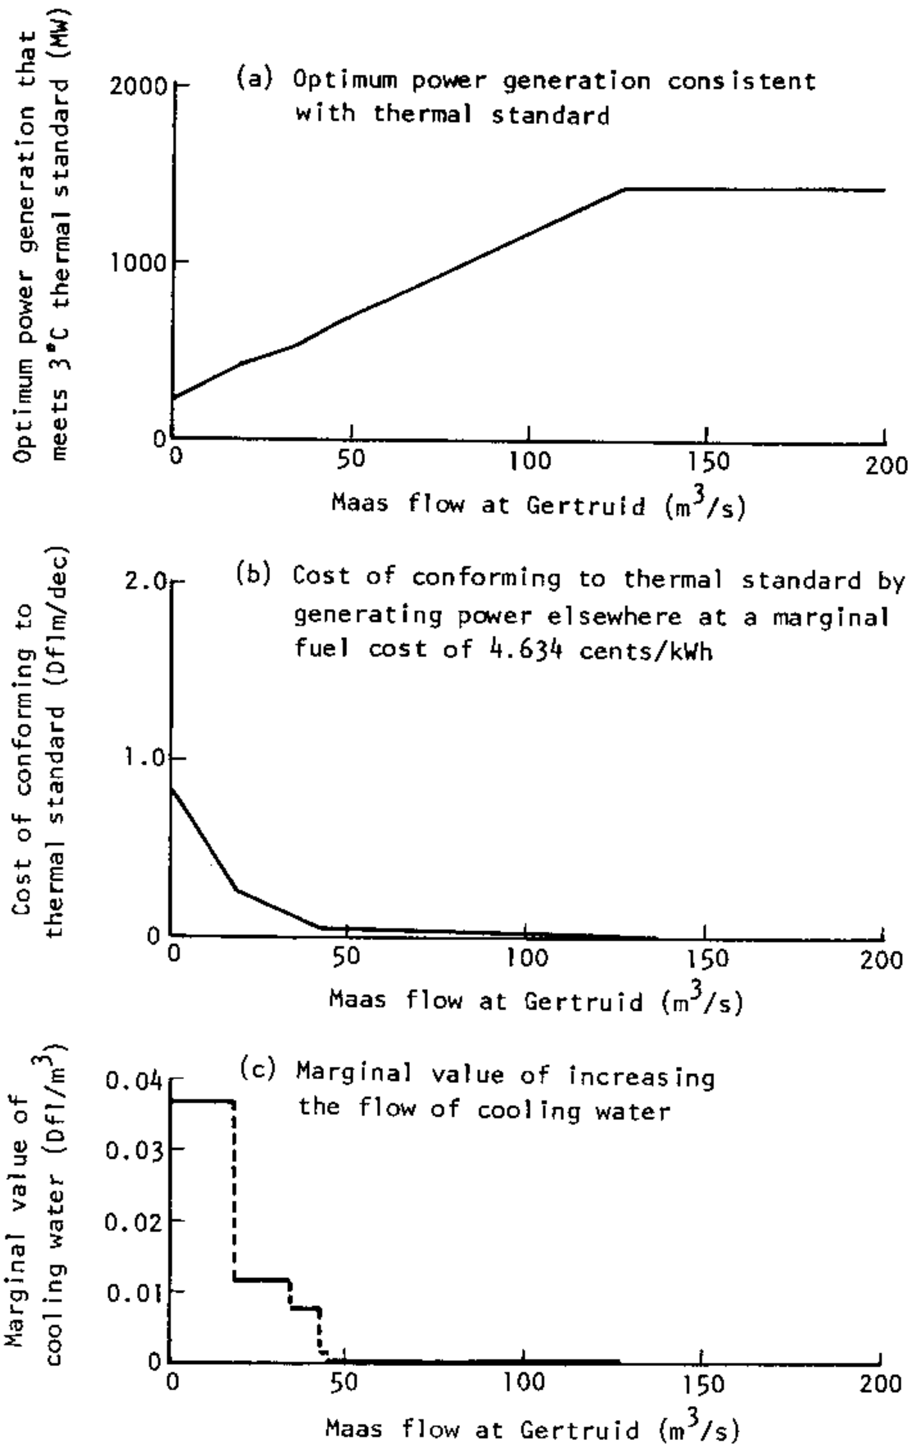

Fig. 5.5--Effect of the 3-deg C thermal standard on the Amer and Donge power plants versus Maas flow at GERTRUID (1976 scenario)

In Fig. 5.5b, we show the effect on total cost of shifting the load to other units. Since no shifting occurs at flows above 127 m^3/s, the cost is zero. If the Maas flow drops to a point between 45 m^3/s and 127 m^3/s, power generation by units with a marginal fuel cost of 4.608 cents/kWh (the most costly units at GERTRUID that are used in the optimal unconstrained schedule) must be reduced, and their capacity replaced by units elsewhere. The cost of power from replacement units need not be larger than 4.634 cents/KHW, so the replacement has little effect on the total cost. Further reductions in Maas flow, however, cause more efficient units to be idled, and hence cost considerably more.

In Fig. 5.5c, we show the value of additional cooling water at GERTRUID. This curve is essentially the derivative of Fig. 5.5b, scaled to read in units of Dfl/m^3. For flows larger than 127 m^3/s, additional water has no value. For flows between 45 m^3/s and 127 m^3/s, each additional cubic meter of water is worth only 0.00055 Dfl. At the extreme, when the Maas flow at GERTRUID is less than 18 m^3/s, the value of additional cooling water is 0.0366 Dfl/m^3.

In Fig. 5.6, we show the same results for the 1985 inventory of generating units. Fig. 5.6a differs from Fig. 5.5a mainly because of the option to use cooling towers at the Amer plant. This is reflected in the flat section of the curve between Maas flows of 27 m^3/s and 90 m^3/s. The use of cooling towers costs about 0.12 cents/kWh, which is reflected in both Figs. 5.6b and 5.6c.

5.12.3. Cooling GERTRUID With Water from Weir Ponds

Under present circumstances, there is only one possibility for increasing the flow of cooling water at GERTRUID. Water can be released from behind the several weirs on the Maas to augment the flow at GERTRUID. In the MSDM network, these weir ponds are represented by nodes 27 MAASLOO, 28 BORN+PAN, 29 LINNE, 30 ROER+BEL, and 31 SAMGRALI. Figure 5.7 shows these nodes with the surface areas of the associated storage basins, as well at their location relative to GERTRUID (node number 33).

The total surface area of the water impounded behind the weirs is 24.77 sq km, so reducing the level by 10 cm will increase the Maas flow downstream of the weirs by about 2.83 m^3/s for an entire decade. The maximum possible reduction in water levels is 130 cm, starting from the levels at which the Dutch prefer to maintain the weir ponds. Further reductions result in water too shallow to accommodate shipping. Thus, by taking full advantage of the water stored in the weir ponds, it is possible to increase the flow in the Maas at GERTRUID by as much as 36.76 m^3/s for one decade.

While the available water is not large, neither are the savings necessarily trivial. For according to Figs. 5.5 and 5.6, each cubic meter of water released from the weir ponds might save as much as 0.0366 Dfl. Emptying the weir ponds just once, therefore, could

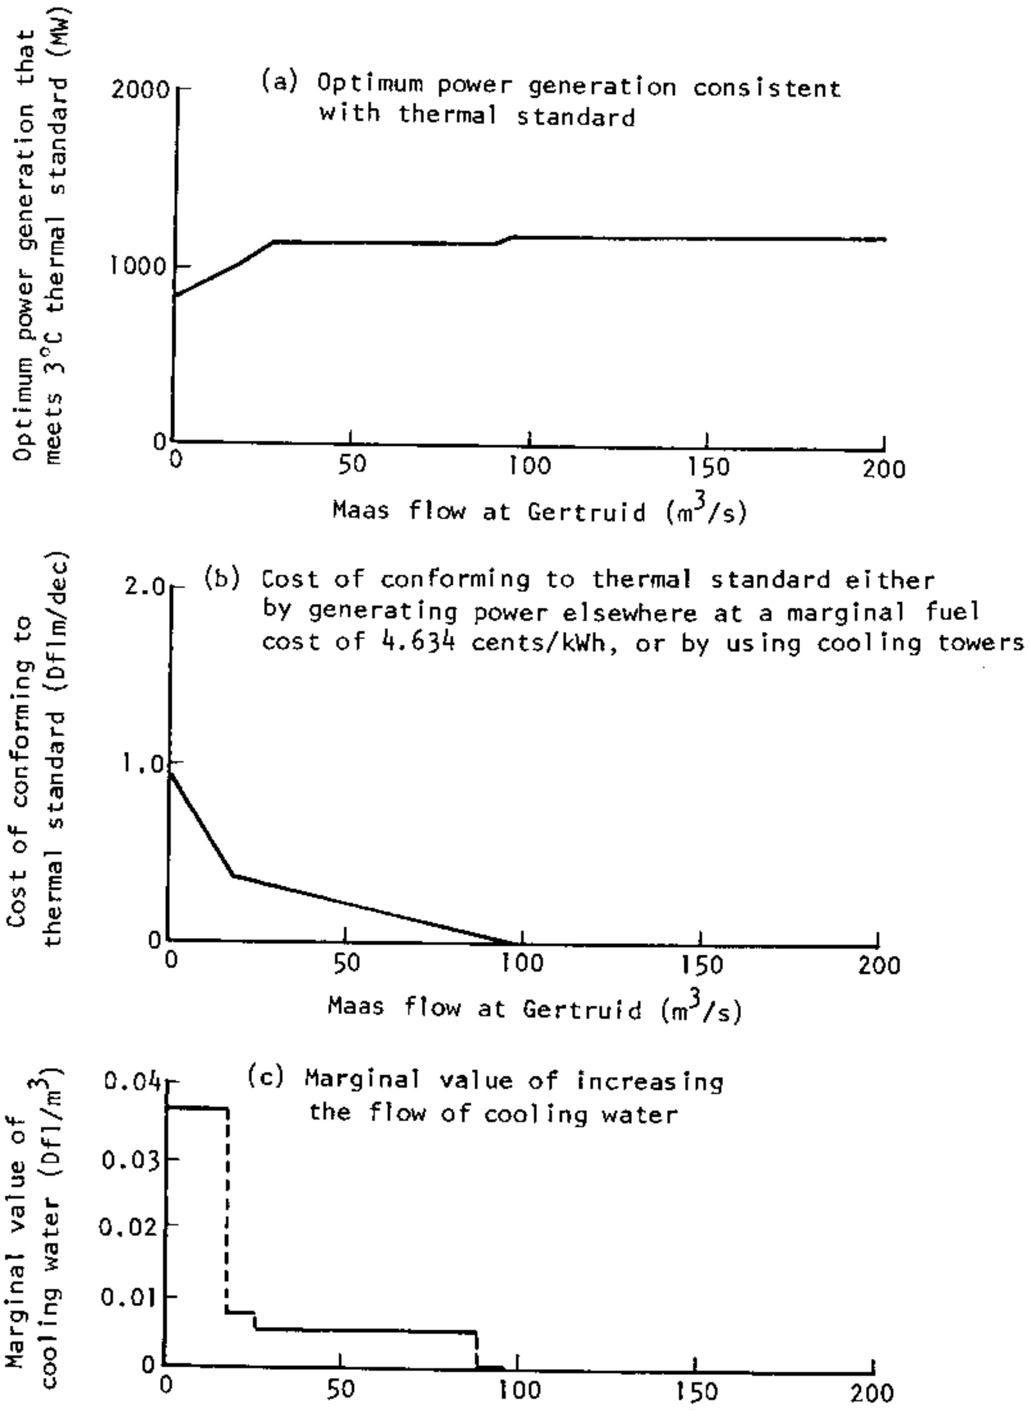
(a) Optimum power generation consistent with thermal standard

(b) Cost of conforming to thermal standard either by generating power elsewhere at a marginal fuel cost of 4.634 cents/kWh, or by using cooling towers

(c) Marginal value of increasing the flow of cooling water

Fig. 5.6--Effect of the 3-deg C thermal standard on the Amer and Donge power plants versus Maas flow at GERTRUID (1985 scenario)

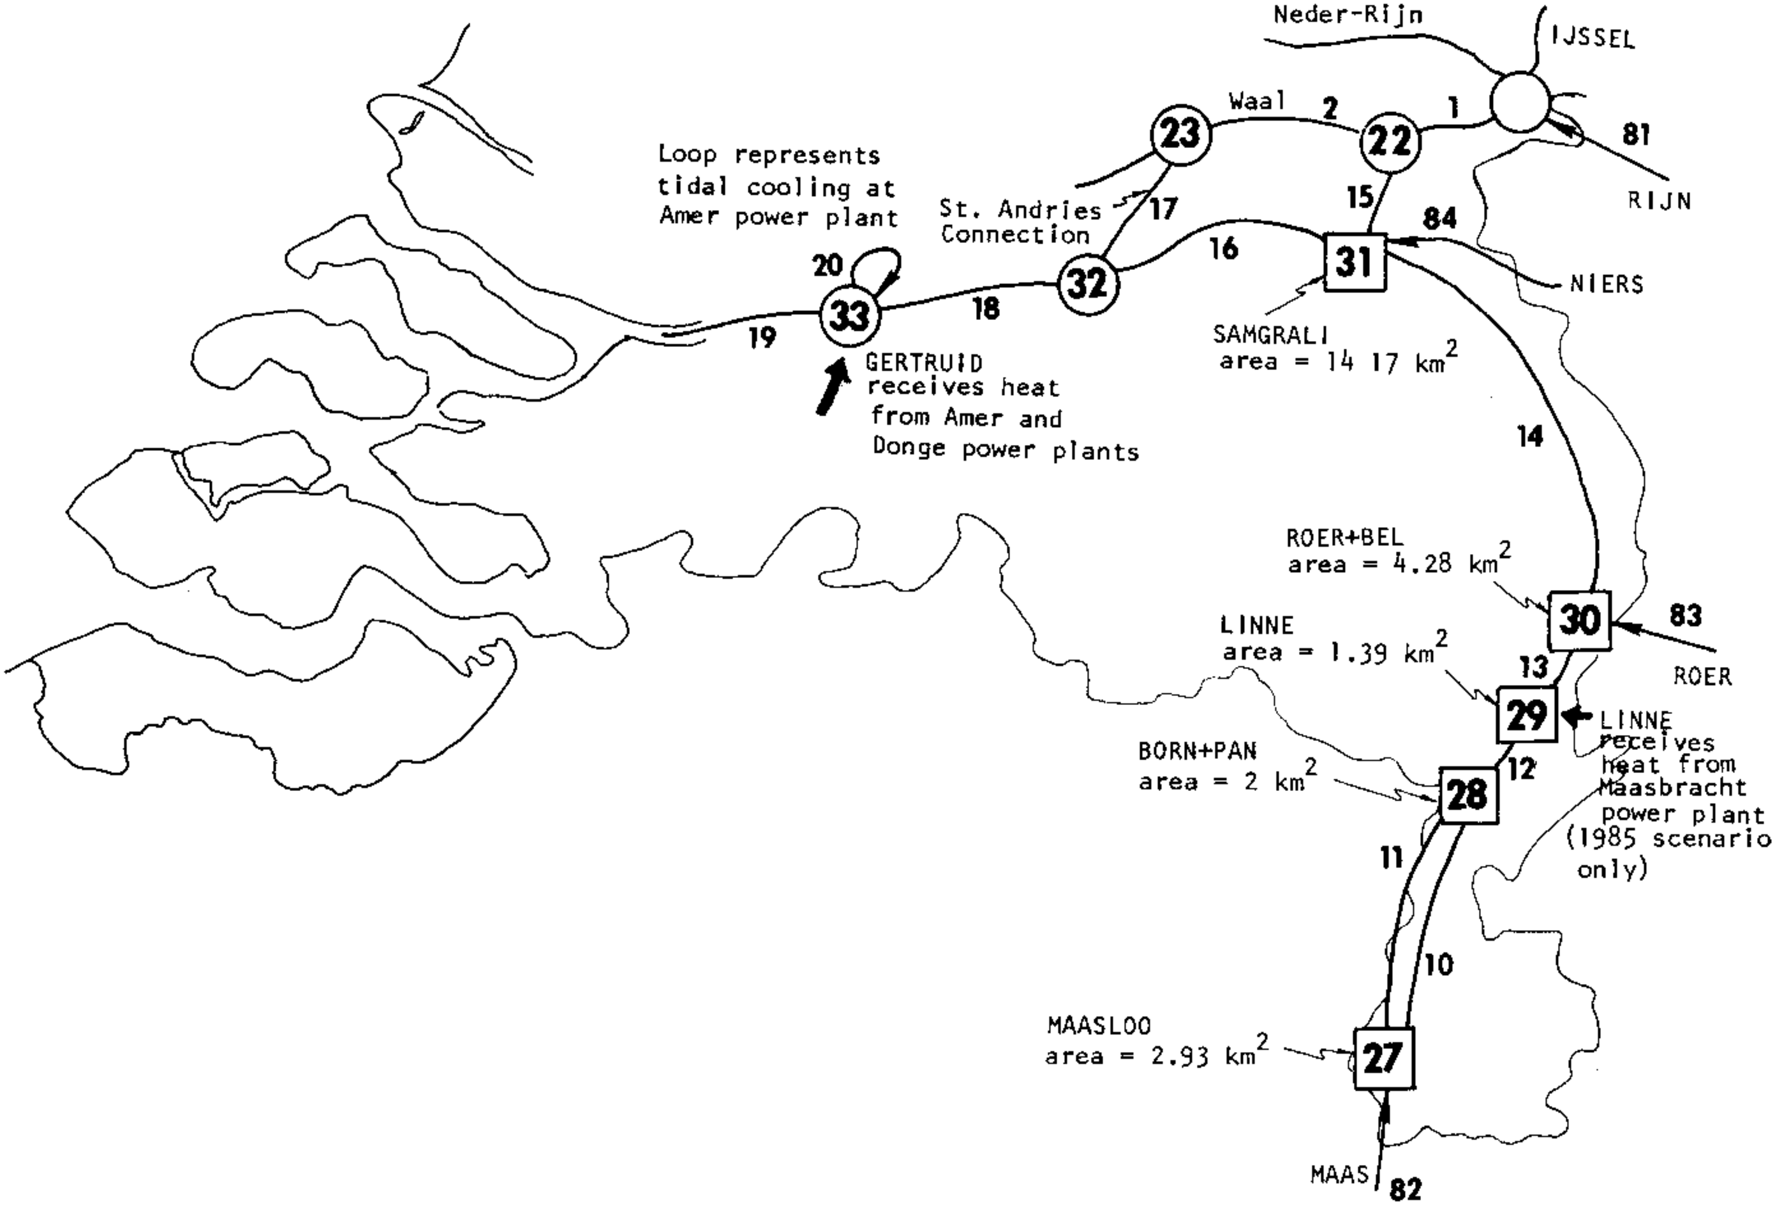

Fig. 5.7--Power plants and weir ponds on the Maas

conceivably save nearly 1.2 million Dfl in power generation costs. A better idea of the possible savings can be formed if one considers the historical frequency distribution of low flows at GERTRUID, and uses it and the cost functions from Figs. 5.5b and 5.6b to compute the average annual increase in power generation costs attributable to applying the three-degree standard at GERTRUID. The frequency distribution of the Maas flow at Lith appears as Fig. 5.8. (We assume the flow at GERTRUID to be essentially the same as the flow at Lith.) This distribution is based on data from the 44 years from 1930 through 1973, plus the year 1976, or 45 years in all. Using this distribution, we calculate that the average annual cost of enforcing the three-degree thermal standard at GERTRUID is 0.304 Dflm/yr for the 1976 scenario, and 1.024 Dflm/yr for the 1985 scenario.

These annual costs represent upper bounds on the savings possible from using the managerial tactic of depleting the weir ponds. As long as enough water is in the weir ponds to cool GERTRUID, no cutback in heat discharge need be made at the power plants Amer and Donge. However, if all the water is released in one decade, none will be available to augment the Maas flow in the next decade. In fact, none will be available until a decade occurs in which the Maas flow is high enough to refill, at least partly, the weir ponds. We must expect, therefore, that only a fraction of the above annual cost can be saved by this managerial tactic. How large a fraction will depend on whether decades with low Maas flows are interspersed among decades with higher flows, or whether low-flow decades tend to occur in bunches.

In order to investigate this question, we carried out a simulation using 45 years of data on the Maas flow at Lith, which we assumed was essentially the same as the flow at GERTRUID. We also assumed that during these 45 years (1930 through 1973, and 1976), the weir ponds were always at their maximum levels. This is certainly not true; we know that in the summer of 1976 the weir ponds were depleted throughout much of July and August. However, this assumption provides the most liberal possible estimate of the amount of water available to augment the Maas flow during these years.

The rules of the simulation were as follows. We specified minimum and maximum water levels for the weir ponds, beginning each simulation with the levels at their maxima. We also selected a target flow for the Maas. If the Maas flow in a decade were lower than the target, we would deplete the weir ponds enough to raise the flow to the target, or until the water levels reached their minima, whichever occurred first. Conversely, if the Maas flow were larger than the target, we would reduce the flow to the target, or enough to refill the weir ponds, whichever was the smaller reduction. The resulting water levels in the weir ponds are taken to be the starting levels in the next decade. The cost for power generation for the decade is calculated from the resulting Maas flow using Fig. 5.5b (for the 1976 scenario) or Fig. 5.6b (for the 1985 scenario).

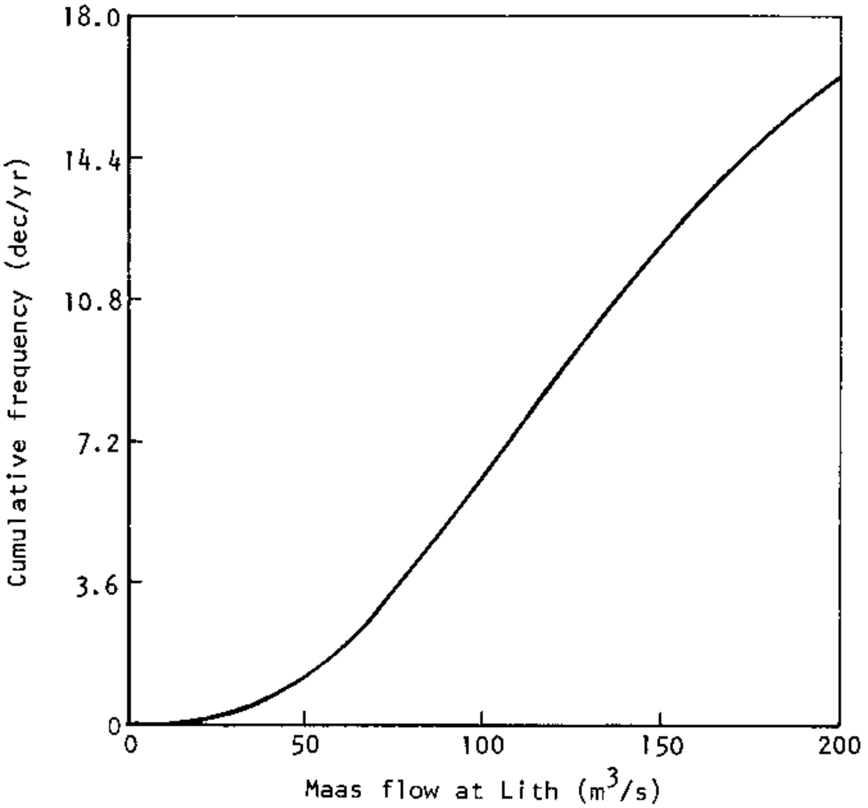

Fig. 5.8--Cumulative frequency of decades with low Maas flow at Lith

The results are shown in Fig. 5.9. For the 1976 scenario, the average yearly cost at GERTRUID of the three-degree standard is 0.304 Dflm/yr, if the target flow is zero. (A zero target flow corresponds to a policy of never depleting the weir ponds.) If the target is raised, the average yearly cost fluctuates, but the policy of depleting the weir ponds to provide cooling water never saves more than 50,000 Dfl/yr. For the 1985 scenario, the cost of the standard is 1.024 Dflm/yr if the weir ponds are never depleted. No depletion policy saves more than 125000 Dfl/yr.

5.12.4. Cooling GERTRUID with Waal Water

One of the technical tactics considered by PAWN is to establish a connection between the Waal and the Maas at St. Andries, which is just upstream of GERTRUID (see Fig. 5.7). The annualized investment cost of this tactic is 0.75 Dflm/yr (Vol. XVI). In addition the withdrawal of water from the Waal at St. Andries causes sedimentation to occur, which obstructs shipping (Vol. XVIII). The sum of these costs proved to be larger than the possible savings in power generation costs under either the 1976 scenario (0.304 Dflm/yr) or the 1985 scenario (1.024 Dflm/yr), and no other benefits were discovered to arise from this tactic (Vol. II). Accordingly, PAWN found the tactic of expanding the St. Andries Connection not to be worthwhile.

5.13. MEASURES AT LINNE

Power plant Maasbracht discharges its waste heat at the node LINNE. This plant did not exist in 1976, so in this section we will only consider the 1985 scenario. As mentioned in Sec. 5.12.1, we have arbitrarily chosen to use Maasbracht at full capacity in the 1985 optimal generating schedule. We could have reduced the load there by as much as 728 MW, shifting it instead to the Amer plant. In this section, as before, we ignore this possibility and estimate the cost of a three-degree standard at LINNE to be the additional cost of generating power elsewhere at a marginal fuel cost of 4.634 cents/kWh. The marginal fuel cost of generating power at Maasbracht is 4.608 cents/kWh.

Figure 5.10 shows the effect of the thermal standard on the power plant Maasbracht. Given the present infrastructure, there is only one measure it is economical to adopt to meet the thermal standard at LINNE. That measure is to reduce the power generated at Maasbracht. It is true that Maasbracht has cooling towers, but even without using them the marginal fuel cost is 4.608 cents/kWh, very close to the highest cost unit it is ever necessary to employ (4.634 cents/kWh). With the cooling towers operating, the cost rises by eight percent to 4.976 cents/kWh.

Note that the cost of meeting the thermal standard at Maasbracht (Fig. 5.10) is much smaller than the comparable costs for the Amer and Donge

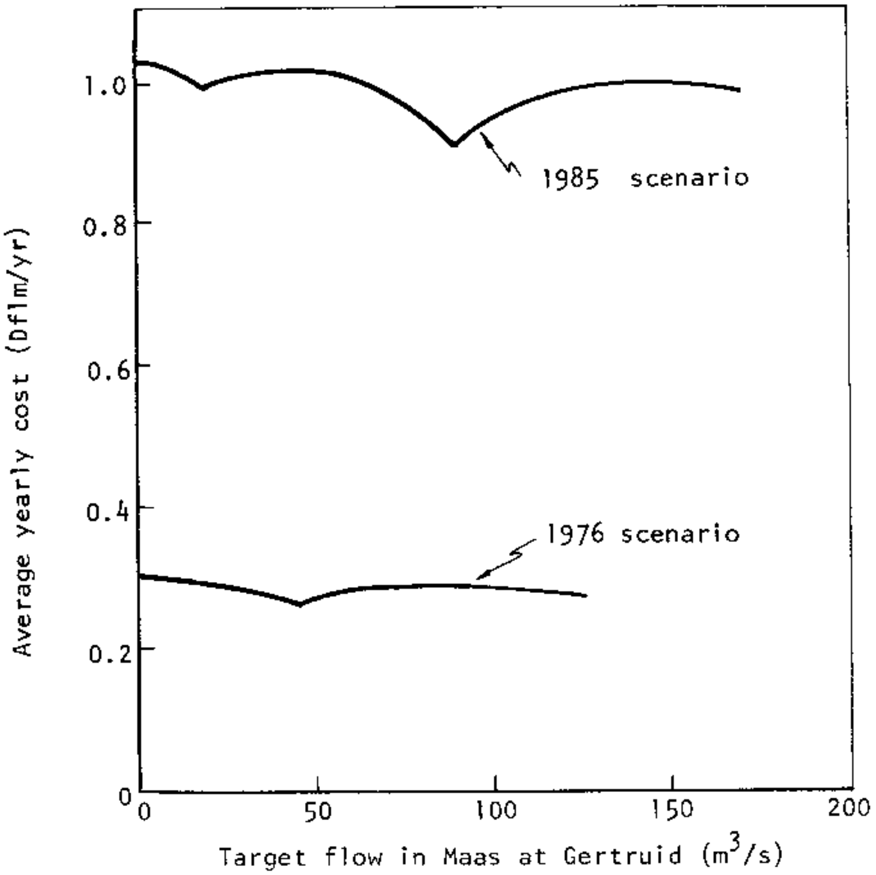

Fig. 5.9--Effect of cooling GERTRUID with water stored in weir ponds on fuel cost of power generation

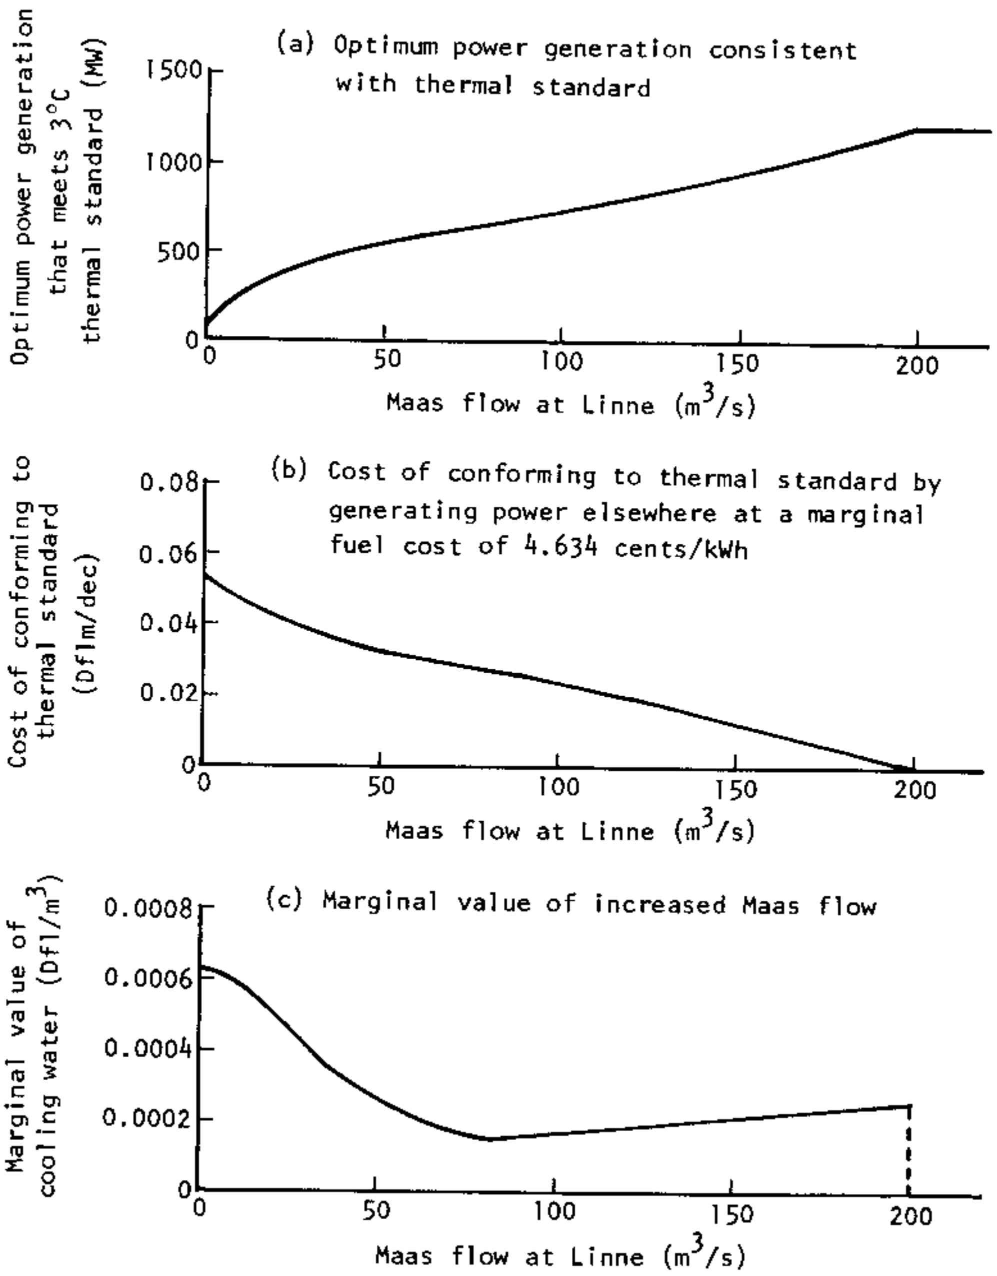

Fig. 5.10--Effect of the 3-deg C thermal standard on the Maasbracht power plant versus Maas flow at LINNE (1985 scenario)

plants (Figs. 5.5 and 5.6). In the earlier figures, the cost of conforming to the thermal standard (Figs. 5.5b and 5.6b) reached nearly one million guilders per decade, for a Maas flow of zero. Similarly, the value of additional cooling water (Figs. 5.5c and 5.6c) nearly reached 4 cents/m^3. In Fig. 5.10, by contrast, the cost of conforming to the three-degree standard at LINNE never reaches even 50000 guilders per decade, and the value of additional cooling water is more than 50 times smaller than the largest values seen in the earlier figures.

The value of increasing the flow in the Maas is shown in Fig. 5.10c. This value is of some interest because of the possibility that Belgium might increase--or decrease--the amount of Maas water they permit to flow into the Netherlands. Even though all units at Maasbracht have the same marginal fuel cost, the value of additional water is not constant. The reason is that excess heat in the Maas at the border is not entirely dissipated by the time the water reaches LINNE. The fraction dissipated is lower for higher flows, which explains the decline in the value of increases in the Maas flow. Where the value begins increasing again (at a Maas flow of about 80 m^3/s) is the point at which we have assumed that the excess temperature at the border begins to drop below three degrees Celsius.[1]

If, on the other hand, additional cooling water could be brought to LINNE with an excess temperature of zero, its value for reducing power generation costs would be 0.061 cents/m^3. This is not to say that there is any way at present to increase the flow of cooling water past LINNE at present. This node lies upstream of most of the water impounded by the weirs on the Maas (see Fig. 5.7). However, PAWN has considered some tactics for pumping water upstream either along the Maas from Roermond--or even from the St. Andries Connection--or along the Zuid-Willemsvaart from Den Bosch (Vol. XVI). In the former case, water would be pumped upstream to BORN+PAN, passing LINNE along the way. This tactic would actually reduce cooling at LINNE unless the amount of water pumped exceeded the nominal Maas flow at LINNE. In the latter case, water would be pumped up the canal Zuid-Willemsvaart to BORN+PAN, and part of it could be allowed to flow downstream along the Maas past LINNE. This tactic would increase the cooling at LINNE. However, the savings in power generation costs would be very small (only 0.061 cents/m^3), and could never be the sole justification for investing in these technical tactics.

5.14. MEASURES AT B'GUMMER

The power plant Bergum discharges its waste heat at node 2 B'GUMMER. In both the 1976 and 1985 optimal generating schedules, it is desirable to use the efficient Bergum units with a marginal fuel cost of 4.061 cents/kWh at their full effective capacity of 613 MW. In the process, they discharge heat at a rate of 172 Mcal/s. Because a cooling circuit exists at Bergum, not all of this heat actually enters the network. The relation between the waste heat discharge and the excess temperature at B'GUMMER is discussed in App. E of Vol. VA.

Bergum has no cooling towers. Therefore, if the three-degree standard is to be met at B'GUMMER, either some of the load must be shifted from Bergum to other units elsewhere, or sufficient cooling water must be brought to B'GUMMER. Both options are available. As always, units with a marginal fuel cost of 4.634 cents/kWh are available elsewhere. And, much more than at nodes on the Maas, a managerial tactic exists to bring cooling water to B'GUMMER.

Cooling water can be supplied to B'GUMMER by the following route. Extract water from the IJsselmeer (node 10 IJSLAKES), and send it north and east along the Prinses Margrietkanaal (link 44 MARGKAN1 to node 1 FRIELAND, and then link 45 MARGKAN2 to node 2 B'GUMMER). The capacity of the Van Starkenborghkanaal (link 49 STABOKAN) is limited by the inability to move water past the lock at Gaarkeuken, but additional water in any desired amounts can be sent north from B'GUMMER along the Nieuwevaart (link 48 B'GMSINK). These routes for cooling water are shown in Fig. 5.11.

The only disadvantage of sending water north from B'GUMMER is that the water must be discharged into the salt Waddenzee, rather than used for a second purpose. This is in contrast to water that can be sent to Groningen province (node 3 GRONETAL) via link 49 STABOKAN, and there used for level control, or irrigation, or even flushing (improving water quality). PAWN has considered a technical tactic that would expand the capacity of link 49 STABOKAN by constructing a bypass around the lock at Gaarkeuken, and possibly constructing one or more pumping stations at various locations on the Van Starkenborghkanaal (Vol. XVI). This tactic could increase the capacity of link 49 STABOKAN from its present 16 m^3/s to as much as 40 m^3/s.

Figure 5.12 shows the effect of enforcing the three-degree thermal standard on the Bergum power plant. In constructing this figure, we have used the MSDM approximation to the actual Bergum cooling circuit that we described in App. E of Vol. VA. It should be pointed out that during dry decades, the present managerial strategy is to send 16 m^3/s east along link 49 STABOKAN, and 4.5 m^3/s north along link 48 B'GMSINK, for a total flow of cooling water of 20.5 m^3/s. According to Fig. 5.12b, this would result in an incremental power generation cost of 0.23 Dflm/dec if the three-degree standard were enforced at B'GUMMER (which is not presently the case). From Fig. 5.12c, we find that each additional cubic meter of cooling water, up to a total of 39.4 m^3/s, should reduce the cost of generating power by approximately 1.4 cents, again if the three-degree standard were applied.

5.15. MEASURES AT UTR+MAAR, AMSTEDAM, AND HAL+IJMU

5.15.1. Heat Discharges With the Optimal Generating Schedule

Numerous power plants discharge waste heat at the nodes 13 UTR+MAAR, 15 AMSTEDAM, and 16 HAL+IJMU. In the optimal generating schedule for

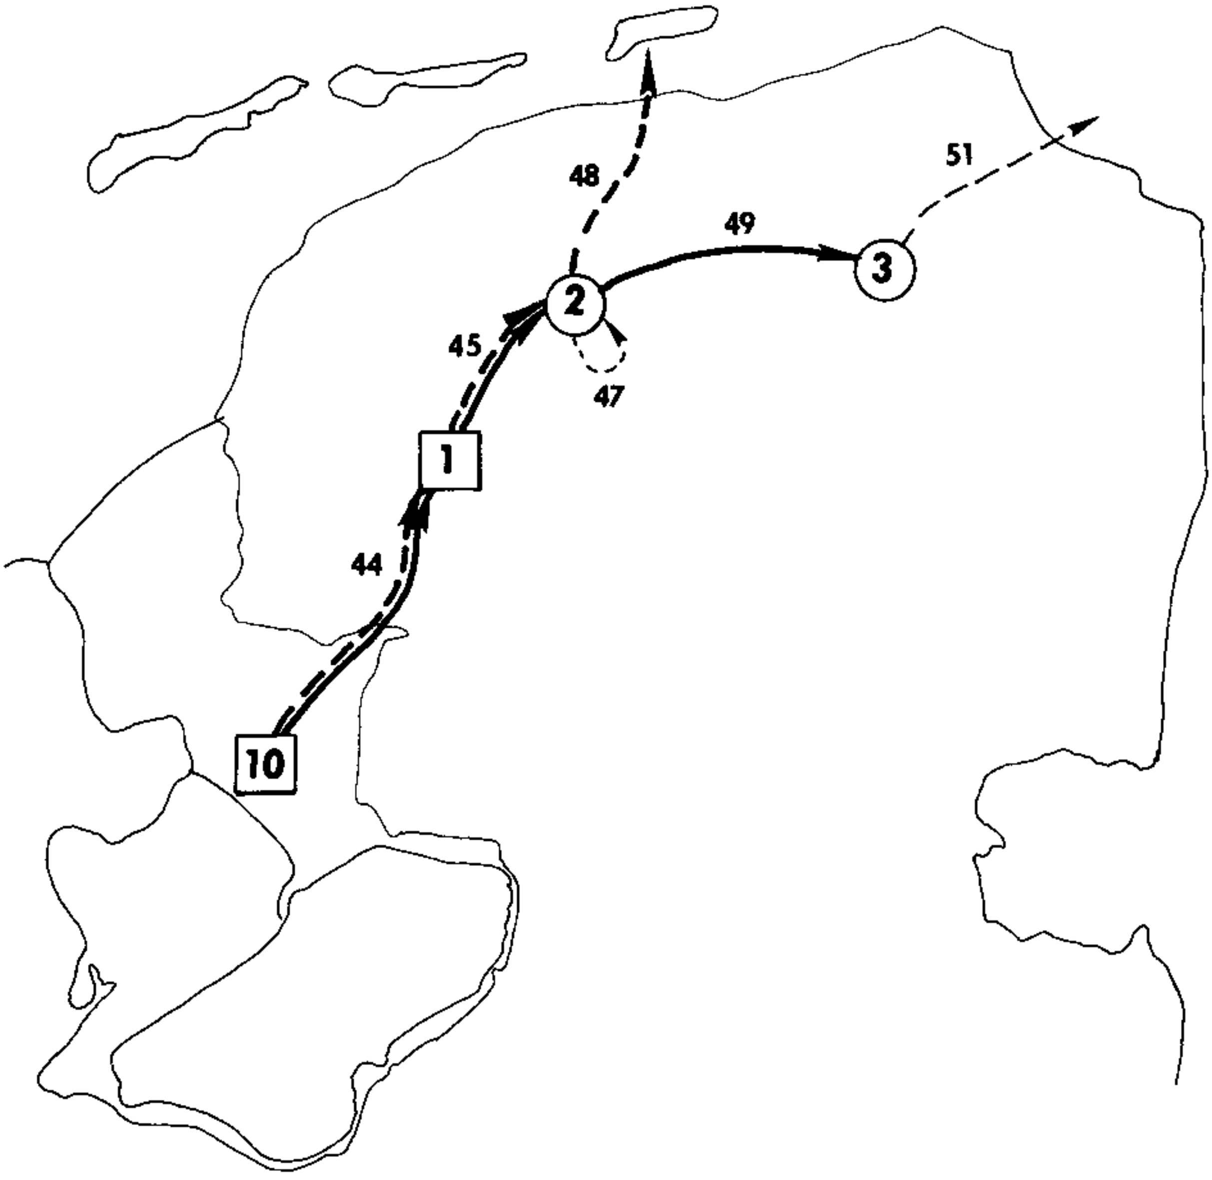

Fig. 5.11--Routes for delivering cooling water to node 2 B'GUMMER

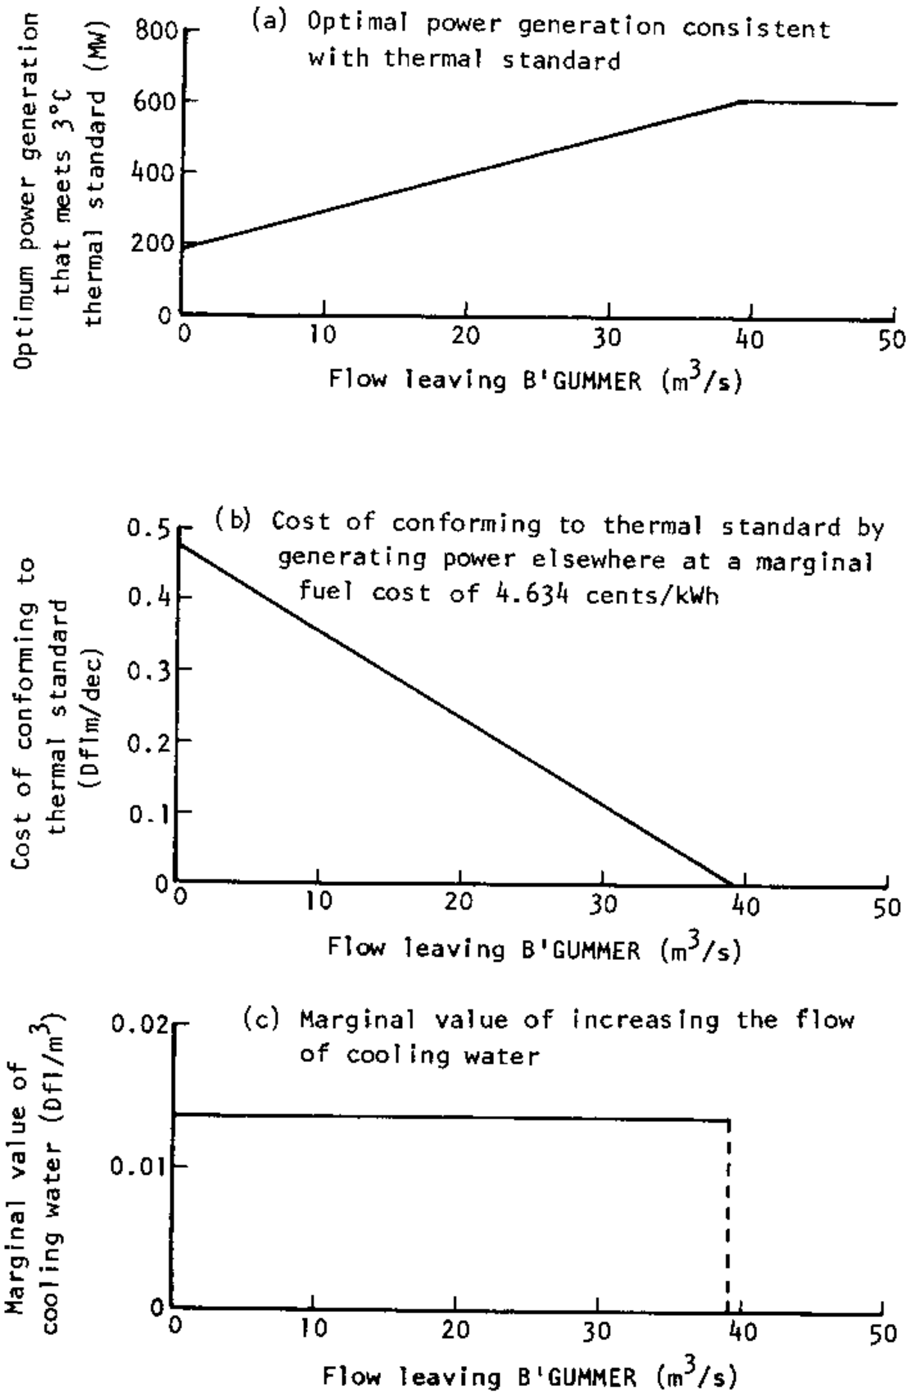

Fig. 5.12--Effect of 3-deg C thermal standard on the Bergum power plant versus flow from B'GUMMER (assumes MSDM representation of cooling circuit at Bergum; see Sec. 5.3.5.3)

1976, we have used 60 MW of capacity at UTR+MAAR, resulting in a peak heat discharge of 18 Mcal/s. The full 200 MW capacity of economical units (i.e., units with a marginal fuel cost less than or equal to 4.634 cents/kWh) at AMSTEDAM are used, resulting in a peak heat discharge of 92 Mcal/s. Finally, the full capacity of economical units at HAL+IJMU is used, totaling to 990 MW of power generated, and 342 Mcal/s of heat discharged. We point out that the unit discharging heat at UTR+MAAR has a marginal fuel cost of 4.634 cents/kWh, and its power can therefore be replaced at no cost, or at most one percent transmission loss.

In the optimal generating schedule for 1985, we use the full capacity of economical units at UTR+MAAR, which totals to 331 MW of power generated, and 101 Mcal/s of heat discharged at UTR+MAAR. Again, 13 MW of this capacity consists of a unit with a marginal fuel cost of 4.634 cents/kWh, and its power can be replaced at trivial cost. If this is done, units at UTR+MAAR deliver only 88 MW of power and discharge only 27 Mcal/s of heat. A unit at AMSTEDAM delivers 481 MW of power and discharges 145 Mcal/s of heat. The optimal 1985 schedule for HAL+IJMU is the same as for 1976.

5.15.2. Measures Available for Meeting the Standard

None of the power plants at these three nodes possess cooling towers, so to comply with the three-degree standard one must either reduce the heat discharges at the nodes, or increase the flow of cooling water. The first option always exists, since there are many idle units with a marginal fuel cost of 4.634 cents/kWh in the optimal generating schedules for both 1976 and 1985. At these nodes, the second option, that of supplying more cooling water, also exists. To cool UTR+MAAR, water can be taken from the Lek at node 12 A-R.MOND, sent north along the Amsterdam-Rijnkanaal (links 29 ARKANAL3, 30 ARKANL4B, and 32 ARKANAL5), and west along the Noordzeekanaal (links 33 NZKANAL1 and 34 NZKANLSL) to the North Sea. The water needed at A-R.MOND can be taken from the Waal to node 23 TIEL and sent along link 28 ARKANAL1, or it can be taken directly from the Rijn by increasing the flow on link 6 NEDRIJN2 (by opening the weir at Driel). Figure 5.13 shows these possible routes for cooling water.

Note that cooling water supplied to UTR+MAAR in this way will also flow past AMSTEDAM and HAL+IJMU. As the water flows beyond UTR+MAAR, it will naturally begin to cool off. However, unless the flow is very slow, it will retain some of the heat from UTR+MAAR by the time it reaches AMSTEDAM. Any heat added at AMSTEDAM will likewise be retained to some extent by the water reaching HAL+IJMU. This demonstrates why it is necessary to consider these three nodes together, instead of separately.

Figure 5.13 also shows a second route for cooling water. This route takes water from the IJsselmeer (node 10 IJSLAKES), and directs it through the Oranjesluis (link 38 ORANJESL) into the Noordzeekanaal at AMSTEDAM. From there it follows the Noordzeekanaal to HAL+IJMU,

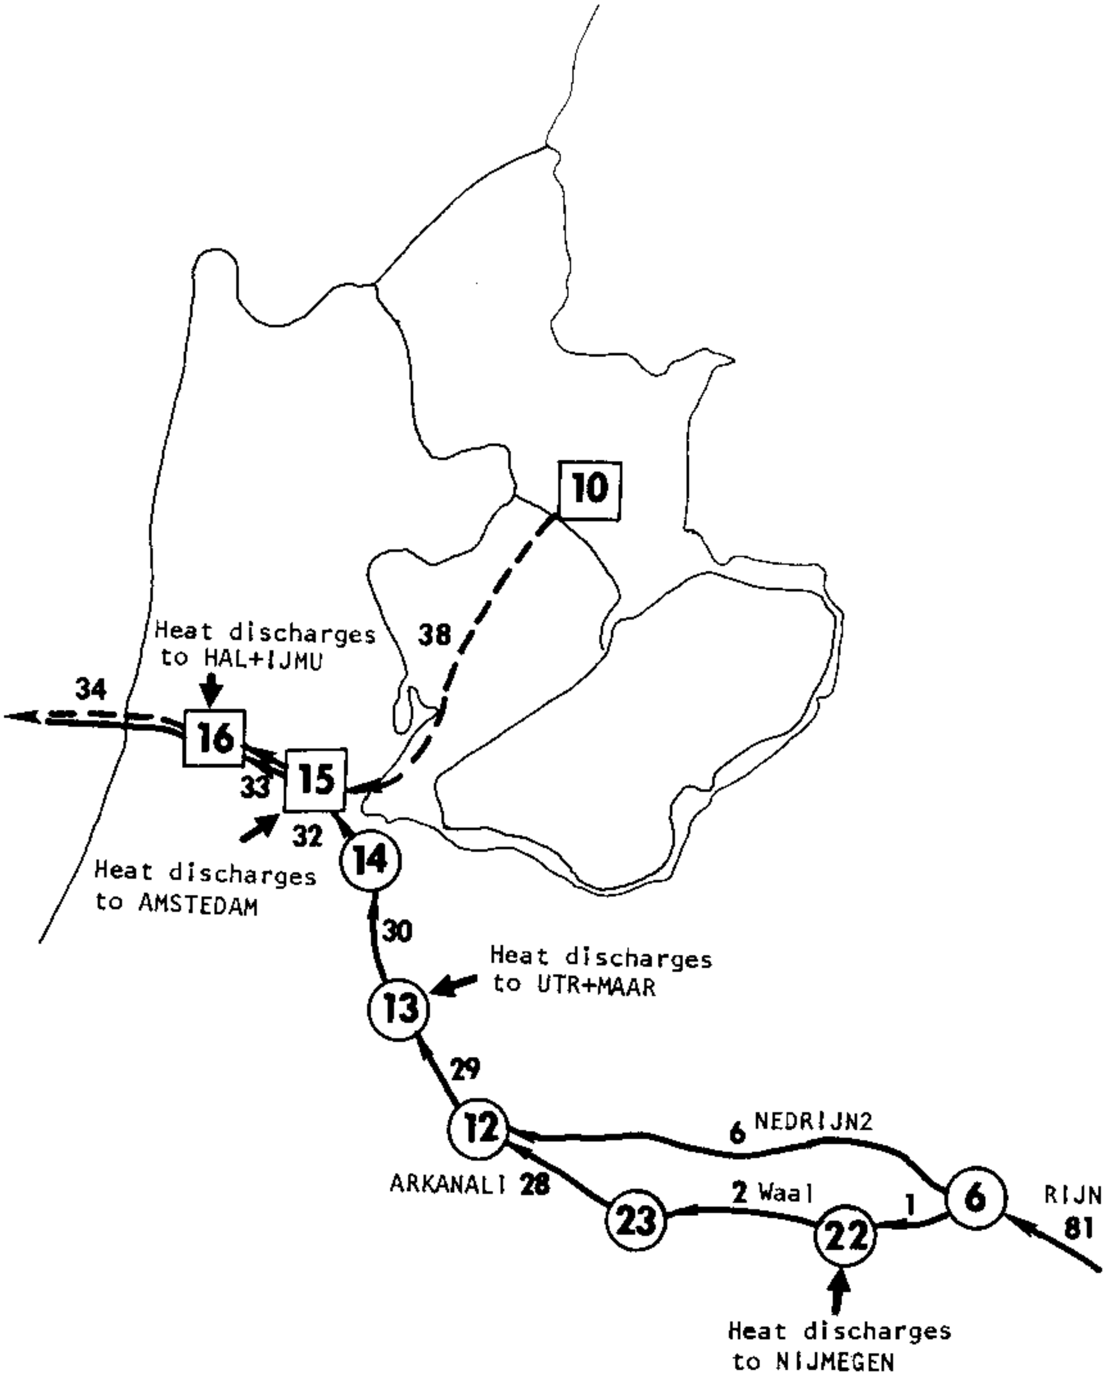

Fig. 5.13--Routes for delivering cooling water to nodes 13 UTR+MAAR, 15 AMSTEDAM, and 16 HAL+IJMU

and is dumped into the North Sea. Water taken via this route cools AMSTEDAM and HAL+IJMU, but not UTR+MAAR.

5.15.3. Effects of Providing Cooling Water

Figures 5.14-5.17 show the effect of providing cooling water via these two routes on power generation at the three nodes. Figures 5.14 and 5.15 were constructed using the 1976 inventory of power plants, while Figs. 5.16 and 5.17 assumed the 1985 inventory. Figures 5.14 and 5.16 show the results if cooling water is taken from the Lek at A-R.MOND and brought north along the Amsterdam-Rijnkanaal, while Figs. 5.15 and 5.17 show results if cooling water is taken from the IJsselmeer.

For all of the figures, we take as our nominal situation flows of 7.7 m^3/s in the Amsterdam-Rijnkanaal at UTR+MAAR, 21.6 m^3/s in the Noordzeekanaal at AMSTEDAM, and 33.3 m^3/s in the Noordzeekanaal at HAL+IJMU. These flows were taken from a run of the Distribution Model picked at random; they have no significance beyond the fact that they are essentially as small values of these flows as one would expect to encounter. In reality a minimum flow of 40 m^3/s is ordinarily maintained in the Noordzeekanaal at IJmuiden, but not all distribution model (DM) runs observed this minimum. In the figures dealing with the effect of cooling water from the Lek, the flow of cooling water is measured as the flow in the Amsterdam-Rijnkanaal at UTR+MAAR. Increases in this flow from its nominal value will increase the flows in the Noordzeekanaal at AMSTEDAM and HAL+IJMU by a like amount. Thus, results in Fig. 5.14 or 5.16 for a flow of, say, 25 m^3/s correspond to an increment over the nominal flow in the Amsterdam-Rijnkanaal at UTR+MAAR of $25 - 7.7 = 17.3$ m^3/s; this implies a flow of $21.6 + 17.3 = 38.9$ m^3/s in the Noordzeekanaal at AMSTEDAM and $33.3 + 17.3 = 50.6$ m^3/s in the Noordzeekanaal at HAL+IJMU.

Similarly, in the figures dealing with the effect of cooling water from the IJsselmeer, the flow of cooling water is measured as the flow in the Noordzeekanaal at HAL+IJMU, whose nominal value is 33.3 m^3/s. Increases from this value will also be added to the flow at AMSTEDAM, but not to the flow at UTR+MAAR. Thus, results in Fig. 5.15 or 5.17 for a flow of, say, 50 m^3/s correspond to an increment over nominal of $50 - 33.3 = 16.7$ m^3/s; this implies a flow of $21.7 + 16.7 = 38.4$ m^3/s at AMSTEDAM. But the flow at UTR+MAAR remains the nominal flow of 7.7 m^3/s.

5.15.3.1. The 1976 Scenario. If we look first at Figs. 5.14 and 5.15, which give results for the 1976 scenario, we find that for flows up to 10 m^3/s larger than nominal--corresponding to flows of 17.7 m^3/s in Fig. 5.14 and 43.3 m^3/s in Fig. 5.15--the marginal value of additional cooling water is between 0.02 and 0.025 Dfl/m^3. For larger flows, the marginal value drops to about 0.01 Dfl/m^3. The reason for this abrupt change is that some of the generating units at AMSTEDAM and HAL+IJMU are constrained by the thermal standard to be used at part load for flows less than 10 m^3/s in

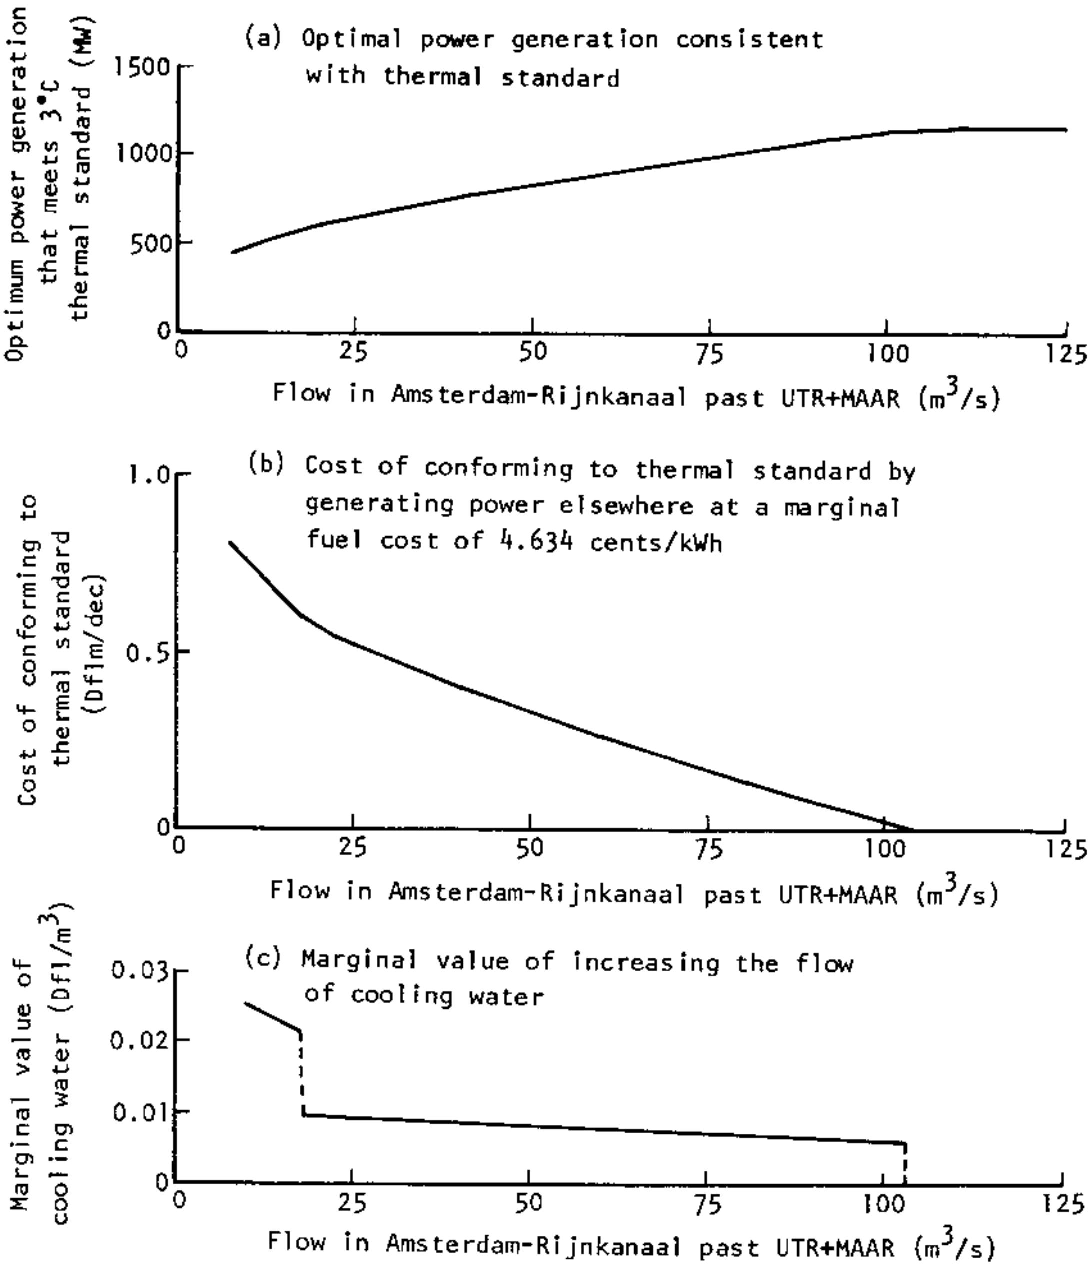

Fig. 5.14--Effect of providing cooling water from the Lek on power generation at UTR+MAAR, AMSTEDAM, and HAL+IJMU (1976 scenario)

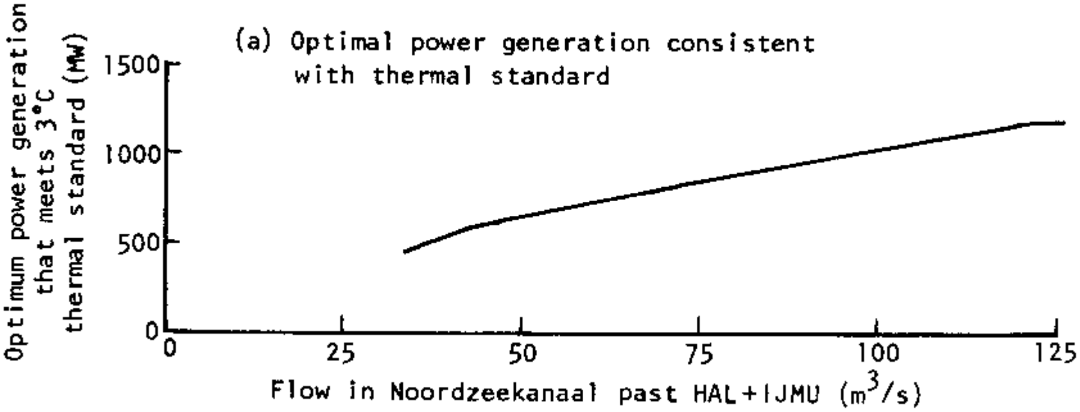
(a) Optimal power generation consistent with thermal standard

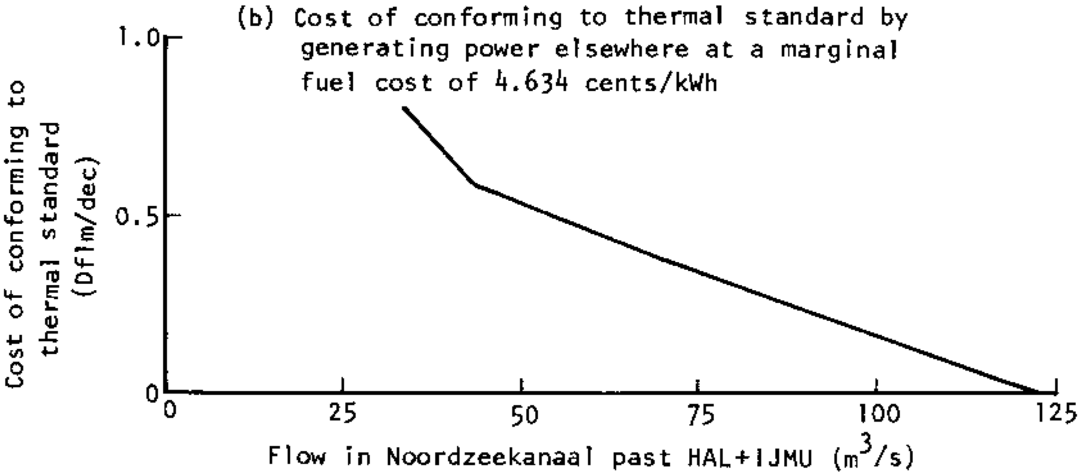
(b) Cost of conforming to thermal standard by generating power elsewhere at a marginal fuel cost of 4.634 cents/kWh

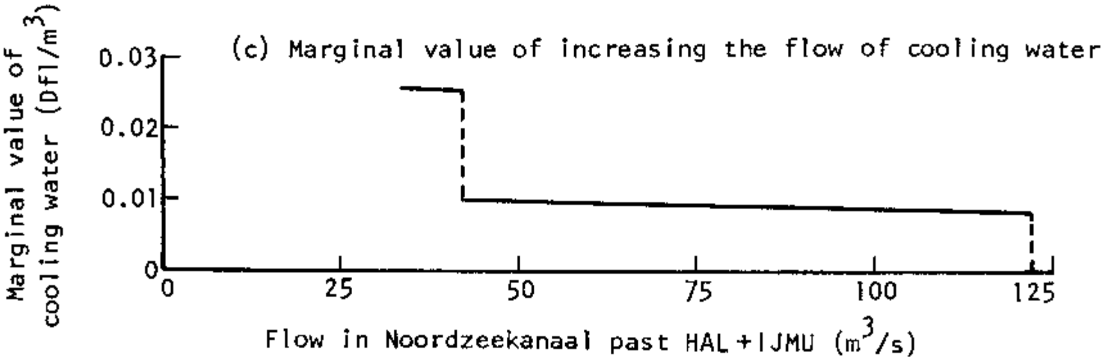
(c) Marginal value of increasing the flow of cooling water

Fig. 5.15--Effect of providing cooling water from the IJsselmeer on power generation at UTR+MAAR, AMSTEDAM, and HAL+IJMU (1976 scenario)

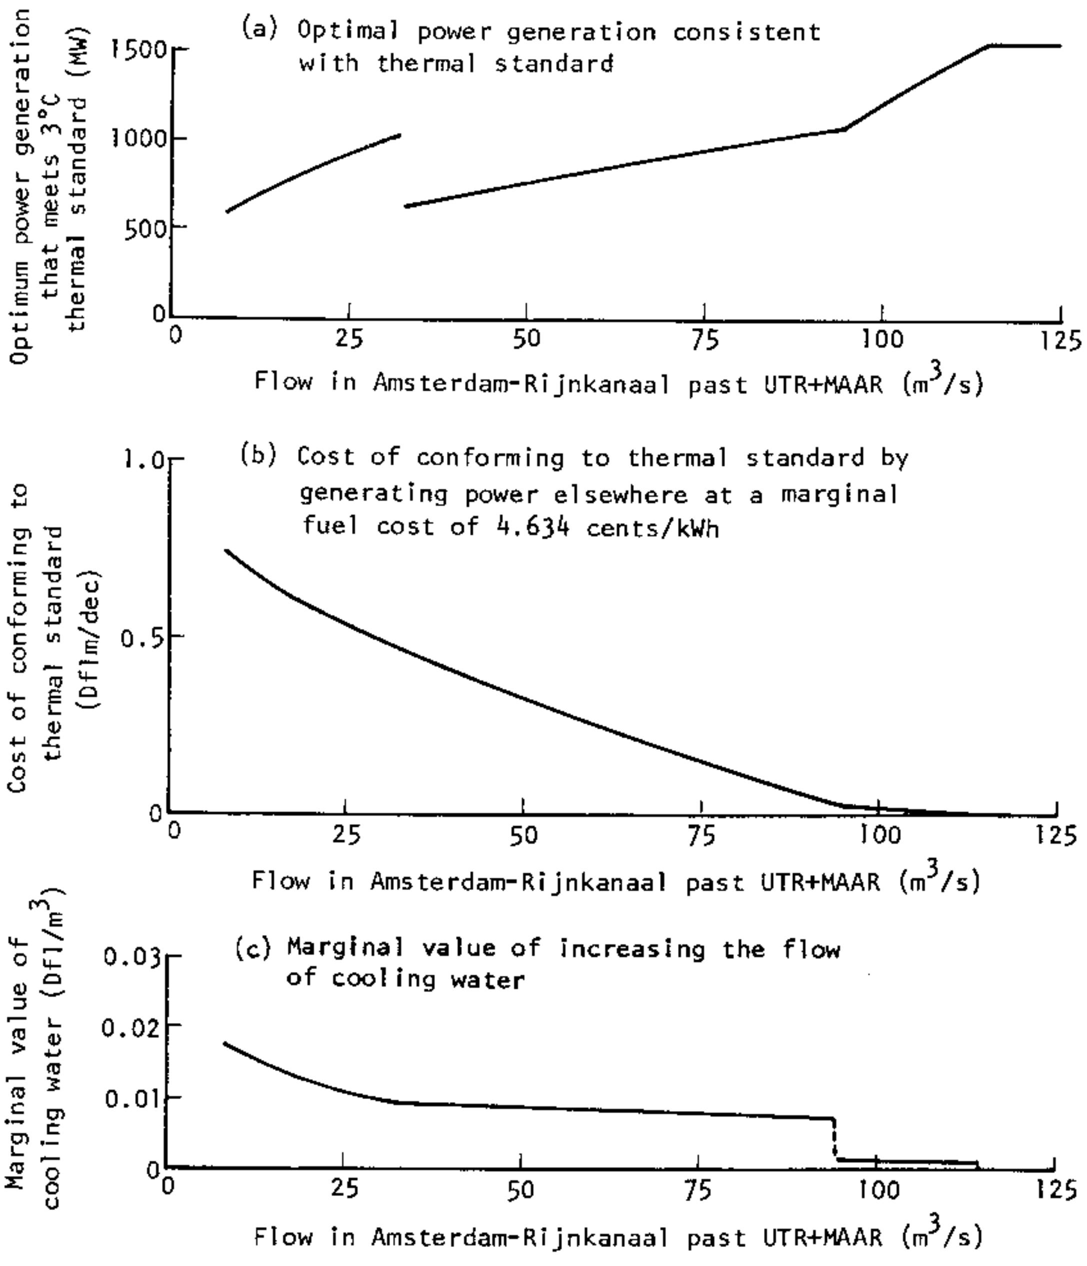

Fig. 5.16--Effect of providing cooling water from the Lek on power generation at UTR+MAAR, AMSTEDAM, and HAL+IJMU (1985 scenario)

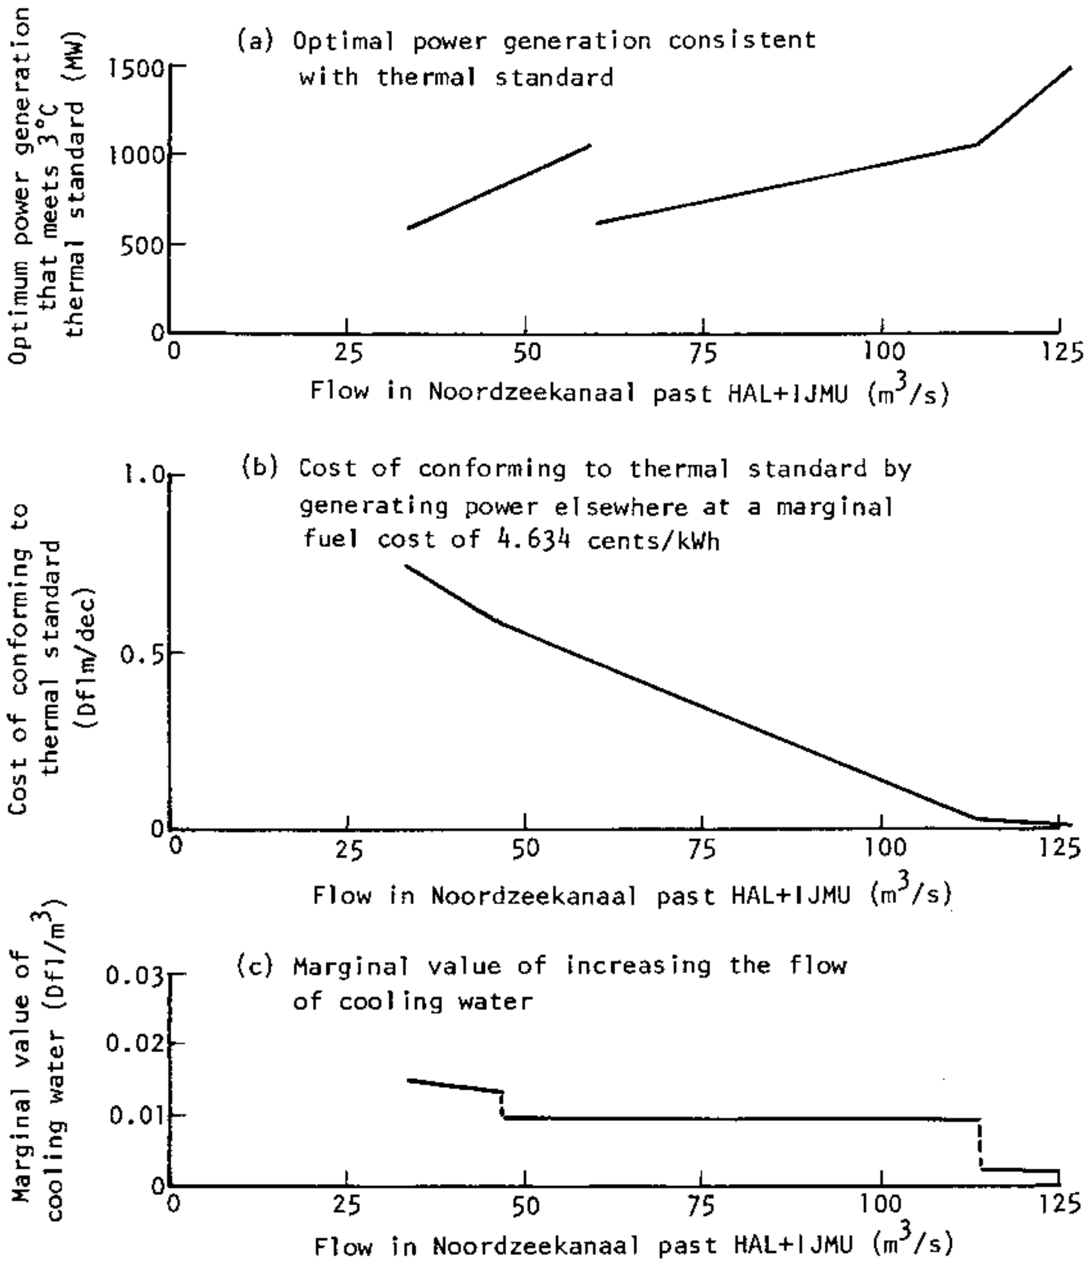

Fig. 5.17--Effect of providing cooling water from the IJsselmeer on power generation at UTR+MAAR, AMSTEDAM, and HAL+IJMU (1985 scenario)

excess of the nominal. At an increment of about 10 m^3/s, however, these units can be used at their full capacity, so further increases in cooling water do not affect them. Since these units have the lowest marginal fuel costs of any at the three nodes, failure to use them at full capacity carries the largest penalty, and implies the largest value for additional cooling water.

Probably the best way to compare the results for the two different sources of cooling water is to compare them at equal increments over the nominal flows. Additional water from the Lek has the advantage that it cools UTR+MAAR as well as AMSTEDAM and HAL+IJMU, but the disadvantage that it is Rijn water, and hence may retain some of the excess heat it contains when it crosses the German border. Additional water from the IJsselmeer does not cool UTR+MAAR, but its excess temperature is zero. Comparing Figs. 5.14 and 5.15 at equal incremental flows reveals little difference. Evidently, both sources of cooling water are equally effective in reducing the costs of complying with the three-degree thermal standard.

5.15.3.2. The 1985 Scenario. If we now examine Figs. 5.16 and 5.17, which present results for the 1985 scenario, we find one immediately interesting feature. The optimal amount of power generated at the three nodes has a discontinuity at an increment of about 26 m^3/s over the nominal, corresponding to flows of 33.7 m^3/s in Fig. 5.16 and 59.3 m^3/s in Fig. 5.17. This discontinuity arises because as the flow on the Noordzeekanaal increases, a larger fraction of the heat added at AMSTEDAM remains in the water when it reaches HAL+IJMU. Thus, at larger flows, a unit of heat discharged at AMSTEDAM requires a larger cutback in the heat discharged at HAL+IJMU. If the required cutback at HAL+IJMU is large enough, it may be more economical to reduce the heat discharged at AMSTEDAM instead.

This could never be the case if the savings per unit of heat discharge were the same at both nodes, since an increase of one unit of heat at AMSTEDAM will require much less than one unit decrease at HAL+IJMU. But it happens that the unit discharging heat at AMSTEDAM has a marginal fuel cost of 4.608 cents/kWh, while the units at HAL+IJMU that must be cut back have a marginal fuel cost of 4.122 cents/kWh. Replacing power from AMSTEDAM by power generated elsewhere costs 4.634 - 4.608 = 0.026 cent/kWh, while replacing power from HAL+IJMU costs 4.634 - 4.122 = 0.512 cents/kWh, or nearly 20 times as much. Above an incremantal flow of 26 m^3/s, more than 19/20 of the heat discharged at AMSTEDAM will reach HAL+IJMU; below that incremental flow, less.

Of course, if the flow becomes large enough, the units at HAL+IJMU will be operating at their full capacity. Further increases in the flow of cooling water will enable the unit at AMSTEDAM to be operated, at least at part load, without requiring cutbacks at HAL+IJMU. In fact, this occurs at an incremental flow of about 86 m^3/s, where the slopes of Figs. 5.16a and 5.17a turn abruptly upward.

The discontinuity, however, is the only peculiar feature to be found in Figs. 5.16 and 5.17. The value of additional cooling water is essentially the same for equal incremental flows, regardless of the source of cooling water. For increments below about 86 m^3/s, the value of additional water is approximately 0.01 Dfl/m^3, while for higher increments it is essentially zero.

5.15.4. Effect of the North-South Connection

Of all the technical tactics considered in PAWN, only the North-South Connection (Vol. XVI) would affect thermal pollution at UTR+MAAR, AMSTEDAM, or HAL+IJMU. In fact, only UTR+MAAR would be directly affected, although the others might be indirectly affected because of a change in the heat transport from UTR+MAAR. The North-South Connection involves building the appropriate pumping stations, etc., to enable water to be taken either from the Lek and transported north into the IJmeer, or from the IJmeer to be transported south to the Lek or the Waal. The route involved is that of the Amsterdam-Rijnkanaal.

At present, water moves north on this route, but if the North-South Connection were used to send water north, there would be two changes in the flow. First, the flow would be increased. In addition, water sent north would enter the IJmeer, rather than flow along the Noordzeekanaal past AMSTEDAM and HAL+IJMU.

When transporting water south, the flows would be entirely different from the present. UTR+MAAR would be cooled entirely by water from the IJmeer, instead of water from the Lek (which is warmer), and the flows would probably be higher in absolute value than at present. Thus, power plants at UTR+MAAR would be less constrained thermally. However, even under the present circumstances these plants can be used nearly to their economical capacities, so relaxing the constraints would be of small value.

In addition, all of the water flowing past AMSTEDAM and HAL+IJMU would have to come from the IJmeer, and none from the Lek. Again, because of the lower temperature of IJmeer water, this would be detrimental. However, since so little of the excess heat from the Lek remains in the water by the time it reaches AMSTEDAM, the reduction in cost would be insignificant.

NOTES

1. In the 1985 case, the marginal fuel cost increases when transmission costs are considered, but in the 1976 case, it remains the same. This is because, for a higher cost unit to be attractive when transmission is considered, it can be no more than one percent more costly than the most costly unit used when

transmission is not considered (assuming that the capacities of the transmission links are never constraining, which is true in this model). This means that in the 1985 case with transmission costs, we will never use units costing in excess of one percent more than 4.608 cents/KWH. Units costing 4.634 cents/KWH are sufficiently economical, and indeed we use some, thus raising the marginal cost. By contrast, in the 1976 case with transmission costs, we will never use units costing in excess of one percent more than 4.634 cents/KWH. This rules out all units in the 1976 inventory, except for those costing 4.634 cents/KWH or less, and so the marginal cost cannot change.

2. According to Fig. 5.1b, the excess temperature of the Maas water at the Belgian border begins to decrease at 100 m^3/s, not at approximately 80 m^3/s. But the Maas flow referred to n Fig. 5.1b is the flow at the border, while the flow referred to in Fig. 5.10 is the flow past the node LINNE. Between the border and LINNE, the flow in the Maas is reduced by both Belgian and Dutch extractions.

REFERENCES

5.1. MR-286a (unpublished PAWN memorandum), "Preliminary Specifications for the Calibration Case, the Base Case, and Various External Supply Scenarios," June 1979.

5.2. Dantzig, G. B., Linear Programming and Extensions, Princeton University Press, Princeton, New Jersey, 1963.

5.3. Sweers, H. E., "A Nomogram to Estimate the Heat Exchange Coefficient at the Air-Water Interface as a Function of Wind Speed and Temperature; A Critical Survey of Some Literature", J. Hydrology, 30:375-401, 1976.

Chapter 6

USES OF WATER: AGRICULTURE

6.1. INTRODUCTION

6.1.1. Water Consumption by Agriculture

Agriculture is the single most important consumer of water in the Netherlands. To get an idea of how much water agriculture consumes, consider that cash crops are grown on more than 20,000 square kilometers of Dutch land. In an average year, approximately 500 mm of water will be transpired from each point of cultivated area, for a total water loss of more than 10,000 million cubic meters of water. In a year warmer than average, such as 1976, this water loss can increase by 50 percent.

We can get an impression of how large this amount of water is by comparing it to other large volumes. For example, the IJsselmeer and Markermeer together have a surface area slightly larger than 2000 square kilometers, and an average depth of four meters. Thus it holds 8000 million cubic meters of water, somewhat less than the annual agricultural consumption. Even this is misleading, since the difference between the minimum and maximum allowable water levels in the IJsselmeer and Markermeer is presently only 20 cm. Thus the usable volume of these lakes is only 400 million cubic meters, or only four percent of annual agricultural consumption.

As another example, consider that the average discharge of the Rijn River during the six months from April through September (approximately the growing season for agricultural products) is just over 2000 m^3/s, and drops to 1400 m^3/s in one year out of six. In 1976, the average Rijn flow during these six months was 1064 m^3/s. By comparison, if the 10000 million cubic meters of water lost by agriculture is spread uniformly over six months, it amounts to 634 m^3/s.

6.1.2. PAWN's Interest in Agriculture

The mere fact that agriculture consumes vast amonts of water makes it important to consider agriculture in the PAWN study. At the very least, the consumption must be estimated and introduced into our models as a net loss of water. But agriculture is important also because some of its consumption can be controlled. If farmers have irrigation equipment, they can determine how much water they apply to their crops, and therefore how much water their crops can transpire. Of course, should they supply less than the optimum amount of water, their crops will suffer damage, and the national income from agriculture will be reduced.

Thus, there are three questions about agriculture that PAWN must answer.

- What is the minimum amount of water that agriculture will consume, as a function of the weather (rain, etc.), geographical location, crop type, kind of soil, and other factors?
- How much additional water might agriculture consume beneficially in areas with irrigation equipment, as a function of the same quantities?
- What fractions of this additional potential benefit would agriculture realize if it were to receive different fractions of this additional water?

6.1.3. Two Agricultural Models

6.1.3.1. DISTAG. In PAWN, we have developed two different agriculture models to help answer these questions. One, a very detailed model called DISTAG (the District Hydrologic and Agriculture Model), is described fully in Vol. XII. At the heart of DISTAG are three PLOT models, which simulate the water flows, salt flows, and crop damage on a single, homogeneous plot of ground throughout the year. For each decade in the growing season, the PLOT models calculate the water lost by transpiration (evaporation from crops), the damage to crops from both drought and excess salinity, and the amount and cost of irrigation. These quantities are functions of the amount of water available in the soil at the start of the decade, the supply of water from rain, and the potential water consumption by crops as determined by the weather. The PLOT models are applied to each of more than 1200 separate plots of ground, each with different location, crop type, kind of soil, mode of irrigation (using surface water, using groundwater, or not irrigated), or other characteristics.

DISTAG and the PLOT models are simulation models. They can be used to answer all three of the questions posed above, but only the first two can be answered conveniently. The first question can be answered by prohibiting all irrigation; the resulting water use is the minimum possible consumption by agriculture. Assuming (as we have--see Vol. XII) that water levels must be held constant in the ditches and other open water bodies accessible to the plot, no action or policy by the water manager can reduce agricultural water consumption further.[1] The second question can be answered by permitting as much irrigation as the farmers desire. This will allow all irrigated crops to transpire as much moisture as they potentially could, according to the prevailing weather conditions. Unless irrigation equipment were installed on more cultivated area, no greater amount of water will be consumed by agriculture.

The third question could be answered by repeatedly simulating irrigated plots, using a different limit on the amount of irrigation water for each simulation. For any plot, a curve might be constructed relating the damage to the crop from drought or excess salt to the amount of irrigation allowed. But the damage suffered in any decade will depend not only on the water applied in that decade, but on how much water is stored in the soil at the start of the decade. The moisture content of the soil is itself a result of the weather and the amount of irrigation in prior decades. Thus, a myriad different curves can be drawn.

6.1.3.2. MSDM Agriculture Model. To avoid this problem, we constructed the MSDM agriculture model. In most ways, it is an aggregated and simplified version of the PLOT models. It is simplified in that the water consumption and damage for many of the 1200 plots considered by DISTAG cannot be influenced by managerial tactics, and hence can be precalculated and provided to MSDM as an exogenously specified input stream. This is true, for example, of all unirrigated plots. It is aggregated in that, of the remaining plots, those that differ only in crop type are combined. Certain types of soils are also combined. But the calculation of crop damage and water consumption is made somewhat more complicated by the need to explicitly calculate the value of different water rations to agriculture.

6.1.4. Organization of This Chapter

This chapter discusses both models, and the relation between them. The discussion here is relatively brief; more detail about DISTAG appears in Vol. XII, while the MSDM agriculture model and the relation between the two models is treated more fully in App. F of Vol. VA. We begin with Sec. 6.2, a discussion of how the cultivated area of the Netherlands is represented geographically. We discuss how this is done for DISTAG, and how we obtain a shorter list of aggregate plots for use in MSDM. Some of DISTAG's plots are removed from detailed consideration, and instead treated more simply, while the remainder are aggregated to form the shorter list. Finally we present the two agricultural scenarios used in MSDM, each represented by a different list of aggregate plots, and each representing a different amount of cultivated land under irrigation.

Next we examine how a single plot of cultivated ground is represented in DISTAG and the PLOT models. In Sec. 6.3 we discuss the hydrologic cycle and its relation to agriculture. It is on this cycle that we base our calculation of water consumption by crops. In Sec. 6.4 we discuss movements of salt in cultivated areas. Then, in Sec. 6.5, we outline the dependence of crop damage on the lack of water and the excess of salt.

Finally we discuss irrigation. The preferred method of irrigaton in the Netherlands is sprinkling (rather than flooding), and we have made the approximation that all irrigation is accomplished in this way. In Sec. 6.6 we discuss how the sprinkling systems are designed and operated. In Sec. 6.7, the last of the chapter, we calculate the value of sprinkling, and the loss likely to be suffered by a farmer if his supply of water for sprinkling is smaller than he would like.

6.2. THE GEOGRAPHY OF DUTCH AGRICULTURE

Of course, it is impossible to treat agriculture in the Netherlands as a single, uniform activity. Different crops are grown on different types of soils using different sources of water in different parts of the country. To take these differences into account, we have partitioned the entire area of the Netherlands into districts. Districts have been further partitioned into subdistrics, and subdistricts into plots. The plot is the smallest area we consider, and agriculture within a plot is taken to be homogeneous in every way. DISTAG considers the Netherlands to consist of more than 1200 plots.

MSDM cannot possibly deal with 1200 plots and still operate economically. Accordingly, we have simplified and aggregated the geographical aspects of Dutch agriculture for MSDM. All 1200 plots are represented in MSDM in at least a very simple form, but only the plots on which crops grown in the open air (instead of under glass) are irrigated using surface water (rather than irrigated using groundwater or not irrigated at all) are represented in a detailed fashion. Even these plots are aggregated, reducing the number of aggregated plots to receive detailed consideration in MSDM to less than 60.

In this section we discuss the geographical aspects of Dutch agriculture. We wish to emphasize how the districts, subdistricts and plots were chosen, and how they are related to the aggregate plots used in MSDM.

6.2.1. Districts

Different parts of the Netherlands draw their water supplies from, and drain their excess water into, different parts of the national and regional infrastructure. It is on this basis that the PAWN districts were chosen. That is, if two different points receive their water from the same node of the Distribution Model network, and discharge their excess water to the same node, they should be in the same district. We partitioned the Netherlands into 78 districts, as shown in Fig. 6.1.

Districts are related to the networks of the Distribution Model Vol. XI and MSDM as follows. In DISTAG, each district is provided with a surface water pool that is separate from any node in the Distribution Model network. The distribution model network is taken to represent the national and major regional infrastructure; hence flows between the surface water pool of a district and the national and regional infrastructure are the connection between agriculture and local water management, as represented by DISTAG, and regional and national water management, as represented by the Distribution Model.

Fig. 6.1--PAWN districts

In MSDM, no such surface water pool is provided for an aggregate plot. Instead, the aggregate plot takes its water directly from an MSDM network node, called the supply node, and drains its excess water to a second MSDM node, called the drainage node. In many instances the same node serves for both supply and drainage for an aggregate plot. Furthermore, a node may serve as the supply or drainage node for several aggregate plots simultaneously.

6.2.2. Subdistricts

We partitioned each district into one or more subdistricts. The criteria for selecting subdistricts was that they had to be homogeneous with respect to interactions between soil and water. What these interactions are, and what factors govern them, is best understood in terms of the hydrologic cycle.

6.2.2.1. The Hydrologic Cycle. Consider a plot of ground on which a farmer is growing a crop. The rain will fall on the soil, and some of it will evaporate. The remaining moisture may be retained in the soil, or it may drain into the surface water system. Conversely, if the farmer possesses the necessary equipment, he can take water out of the surface water system, or from deep in the ground, and sprinkle it on his crops. These movements of water are elements of what is called the hydrologic cycle. The elements of this cycle of greatest interest for agriculture are shown in Fig. 6.2.

6.2.2.2. Two Moisture Pools. In this figure, we distinguish two pools of moisture, root zone moisture and subsoil moisture. (A third moisture "pool," the surface water system, is shown, but merely serves as a source and sink for water moving into and out of the root zone and subsoil pools.) To define the root zone and subsoil moisture pools, it helps to think of the plot of ground as a vertical column of soil. The top 50 centimeters or so constitute the region from which the roots of the crops planted in the column will obtain most of their water. We call this region the root zone, and the moisture in it is the root zone moisture. Of course, the root zone need not be 50 centimeters deep; depending on the crop and soil type it may be deeper or shallower.

Below the root zone is the subsoil. Deep in the subsoil is a region saturated with water; the depth at which a saturated condition is first encountered is called the groundwater level. Between the groundwater level and the bottom of the root zone is an unsaturated zone. Moisture in both the saturated and unsaturated zones is included in the subsoil moisture pool.

6.2.2.3. Soil Types. The most important characteristic of the two moisture pools is their capacity to hold water. This is especially important in the case of the root zone. A crucial factor in determining the capacity of a column of soil to hold water is its soil type, which is therefore an important factor in our choice of subdistricts.

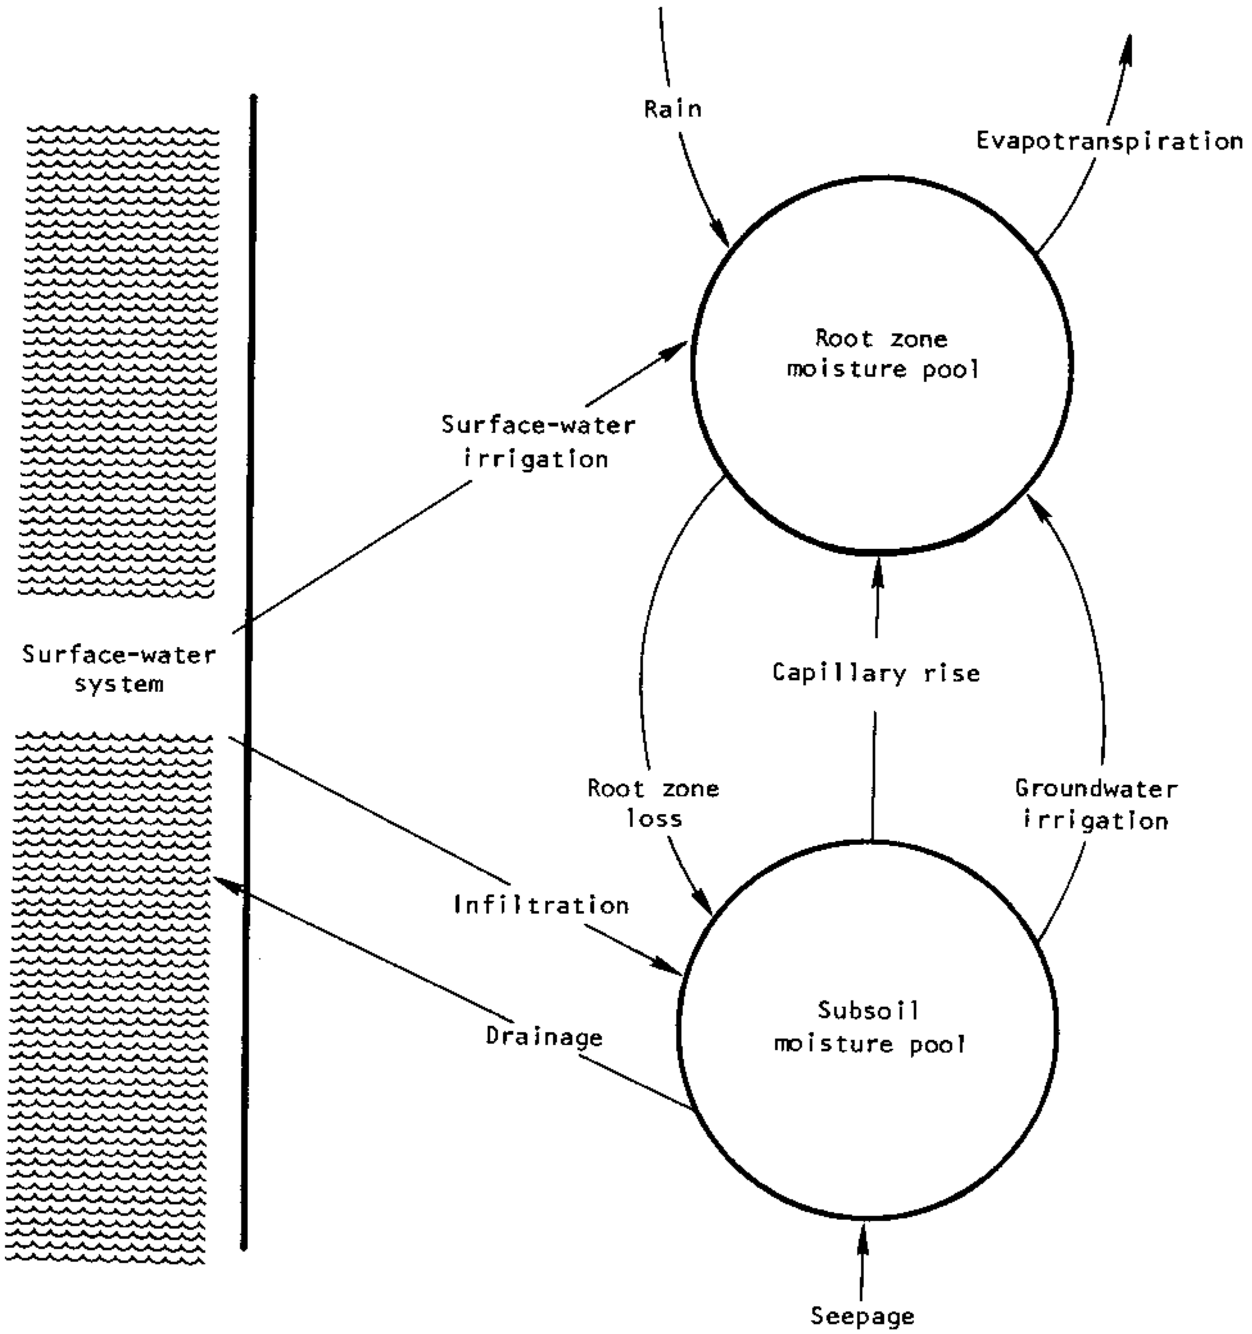

Fig. 6.2--Elements of the hydrologic cycle

We measure the capacity of the root zone as the maximum amount of moisture it can make available to plants. The driest the soil can become is its wilt point; if it contains only this much water, plants can extract none of it. The wettest the soil can be is its saturation capacity; any addtional moisture flows into the subsoil, beyond reach of the roots of plants. (Between these two points, but closer to the saturation capacity, is the field capacity, which is the maximum amount of water that the soil can hold indefinitely against the force of gravity. In DISTAG, the field capacity was chosen as the measure of soil capacity; in MSDM, the saturation capacity. The difference, which was inadvertent, hardly matters because in MSDM we are always concerned with situations in which the soil is quite dry; thus the soil capacity plays no role.) Differences in the water capacities of soils can be large; for example, plants can extract a maximum of only 3 mm of water from a saturated 10 mm column of sandy soil, while they can extract nearly 5 mm from a similar column of peat (which is mostly organic matter).

6.2.2.4. Pathways for Water Movement. The pathways for water movement are represented by the labeled arrows in Fig. 6.2. In this section we will briefly describe them and mention the major factors that govern the rates of water movement on each pathway.

The sources of water for the root zone are rain, surface water irrigation, groundwater irrigation, and capillary rise. Rain is specified exogenously and need not be considered in partitioning the districts into subdistricts.

Surface water and groundwater irrigation are under the control of the farmer. When the farmer irrigates, his goal is to maintain the root zone moisture within reasonable limits, so as to avoid both flooding the field (as might happen if he sprinkled and immediately afterward it rained) and allowing the field to become so dry that the crops suffer damage. The amount of water he applies will therefore depend on the root zone capacity, which depends in turn on the soil type, as noted above.

Capillary rise is the upward movement of water from the subsoil to the root zone. Much as oil rises in a wick, so moisture will flow upward through the smaller soil pores. The rate depends most strongly on the depth of the groundwater. Hence it is important to distinguish areas with typically shallow groundwater levels from areas whose groundwater levels are generally deep. The rate also depends on the amount of moisture found in the root zone.

The loss of water from the root zone is due largely to evapotranspiration. Transpiration is the evaporation of water from plants. In addition to transpiration, water is lost by evaporation directly from the bare soil. This loss is much smaller than the loss by transpiration from an equal area, except when the top few millimeters of soil are very wet. The sum of evaporation and transpiration losses is called (appropriately) evapotranspiration. The rate of evapotranspiration depends on the amount of moisture in the root zone in comparison to the root zone capacity, on the type of crop being cultivated, and on climatological conditions.

Root zone loss occurs when the root zone moisture exceeds the root zone capacity. Gravity then draws the water out of the root zone and into the subsoil. Should large amounts of water be applied to the root zone, for example by a heavy and prolonged rainfall, the water lost from the root zone may eventually reach the saturated zone and raise the groundwater level. Root zone loss occurs mostly during the winter months, when evaporation losses are low. While there can be root zone loss during the months of a wet growing season, root zone loss is essentially zero during a dry growing season, such as we have analyzed with MSDM.

The remaining water movements shown in Fig. 6.2 involve only the subsoil moisture pool. The subsoil gains water through infiltration and seepage, and loses water through drainage.

Infiltration and drainage are movements between the subsoil and nearby surface water bodies, under the influence of the difference between the surface water level and the groundwater level. Movement from nearby surface water into the soil is infiltration, while movement in the reverse direction, from the soil into the open water, is drainage. Drainage occurs in essentially all parts of the Netherlands, but significant infiltration only occurs in the lowlands. These areas possess extensive networks of closely spaced ditches, whose walls offer a large aggregate area at which soil and water are brought into contact.

Seepage is the underground movement of water into the subsoil from remote locations. It is driven by the difference between the water level at the remote location, and the local groundwater level. For example, the groundwater levels in many of the low-lying areas of the Netherlands are several meters below the mean sea level. This level difference results in a pressure gradient that drives sea water tens of kilometers inland into the subsoil.

In principle, seepage can occur in the reverse direction as well, from the subsoil to more remote locations on the national or regional infrastructure. In Vol. XII, this is called outside drainage, and is modeled by diverting some of the drainage term to surface water at locations remote from the plot of ground under consideration. We have discovered some areas of the Netherlands, particularly those in the south and northeast at higher elevations, where it was necessary to assume significant (60 mm/year) water losses by this mechanism in order for DISTAG to match the observed regional water balances. (We note that the prevailing opinion among Dutch hydrologists is that no more than 10 mm/year is lost in this manner. Where the remaining water goes is unresolved).

6.2.2.5. Partitioning Districts into Subdistricts. According to the above discussion, the interactions between soil and water are governed almost completely by the root zone and subsoil soil types, whether the normal groundwater level is deep or shallow, and whether significant infiltration occurs.

We distinguish areas with different soil types in the root zone and subsoil. DISTAG considered a total of 26 soil types. If all possible combinations of root zone and subsoil soil types had occurred, there would have been 26x26, or 676 categories. Instead, only 16 combinations were actually chosen.

We distinguish areas whose average groundwater level is deep (more than one meter below the soil surface) from those whose level is shallow (one meter or less below the soil surface). Areas with deep groundwater never experience significant capillary rise, while those with shallow groundwater regularly have significant capillary rise. The groundwater is generally shallow in locations close to open water, and hence at relatively lower elevations than locations where the water table is deep. Significant infiltration only occurs in the lowlands, where the groundwater is always shallow. Thus the last two properties were combined into a <u>landform</u> criterion, which called for distinguishing lowlands, low highlands, and high highlands.

Partitioning the 78 districts on the basis of soil type and landform results in 144 subdistricts, or about two per district. Some districts have only one subdistrict, while others have four, five, or even six. The maximum number of subdistricts in a district is six.

6.2.3. Plots

The greatest degree of geographical refinement in DISTAG is at the level of plots. We partitioned each subdistrict into one or more plots, arriving finally at 1267 of them. The criteria for choosing plots were homogeneity of crop and of irrigation. We distinguished fourteen crops, one of which combines nature and fallow land, and thirteen of which represent actual crops or groups of crops cultivated for market. Two of the thirteen actual crops are vegetables under glass and flowers under glass, which must be irrigated, but the other eleven market crops might be either irrigated or not, at the discretion of the farmer. Further, some plots might be irrigated with ground water, and others with surface water. Table 6.1 shows the various possible types of plots within a single subdistrict. A count shows that the maximum number of plots in a subdistrict is 38.

No subdistrict actually contains 38 plots. Some subdistricts have as many as fifteen plots, while others have as few as two. The average number of plots per subdistrict is slightly less than nine.

Table 6.1

POSSIBLE PLOT TYPES

		Source of Irrigation Water		
	Crop Name	None	Surface Water	Ground Water
0.	Nature and fallow ground	Y	.	.
1.	Grass	Y	Y	Y
2.	Consumption potatoes	Y	Y	Y
3.	Milling potatoes	Y	Y	Y
4.	Seed potatoes	Y	Y	Y
5.	Sugar beets	Y	Y	Y
6.	Cereals	Y	Y	Y
7.	Cut corn (for fodder)	Y	Y	Y
8.	Bulbs	Y	Y	Y
9.	Vegetables in open air	Y	Y	Y
10.	Pit and stone fruits	Y	Y	Y
11.	Trees and ornamental shrubs	Y	Y	Y
12.	Vegetables under glass	.	Y	Y
13.	Flowers under glass	.	Y	Y

6.2.4. Aggregating and Simplifying Plots for MSDM

For the purposes of MSDM, it was impossible to deal with 1267 plots. Therefore we eliminated some of them, and aggregated the remainder to obtain a more manageable number. In this section, we briefly outline our methods for doing this; in App. F of Vol. VA we discuss them in detail. First of all, we eliminated district number 78 from consideration, and with it one subdistrict containing eight plots. This district represents the Markerwaard, land that has yet to be reclaimed from the Markermeer, a section of the IJssel lakes. In none of the MSDM runs was this reclamation assumed to have been done.

In MSDM, no plot was considered in detail if it was not sprinkled from surface water. That is, all plots were eliminated that were either not sprinkled or were sprinkled from groundwater. This does not mean that plots not sprinkled from surface water were ignored completely. Rather, it means that we can calculate the flows of water and salt between the surface water system and these plots before running MSDM. (From Fig. 6.2, we see that for plots with no surface water irrigation, these flows are infiltration and drainage.) The precalculated water and salt flows can then be presented to MSDM as a stream of input data. By this means, we eliminated all but 300 plots.

(Actually, the managerial strategy can influence the salt concentration in the water supplies to some districts, notably those in the midwest. For these districts, therefore, it is not strictly true that salt flows can be precalculated. We calculate the salt flows using a managerial strategy that approximates present Dutch practice (the RWS strategy, as described in Vol. XI). Because the salt concentrations predicted under this strategy are little different from those predicted under the "superior" strategies found by MSDM, we feel that precalculating the salt flows does not significantly distort our results.)

Because of their very high value and low aggregate demand for water, we have assumed that vegetables and flowers under glass (crop types 12 and 13) are given all the water they demand. Also, being under glass these crops have no other source of water than irrigation. The fact that we give these crops all the water they demand means we need only consider damage they may suffer from excess salt. (Crops under glass are the most salt sensitive crops we consider.) To represent crops under glass in MSDM, therefore, we attach a cost function to each of the twenty nodes that supplies water to crops under glass, which depends only on the chloride concentration at that node. This cost function estimates the salt damage to crops under glass supplied from the node in question (see Chap. 4, and App. D of Vol. VA). By introducing twenty cost functions into MSDM, this simplified representation eliminates a further 90 plots from more detailed consideration, leaving only 210.

Two hundred ten plots is still an inconvenient number, so we have reduced the number further by aggregation. First, we note that less geographical detail is represented in MSDM than in the distribution model (Vol. XI), and that this suggests that some districts might be combined. Then we argue that little is lost for the purposes of MSDM by considering fewer crop types and soil types than are used in DISTAG. Aggregating the plots on these grounds results in 57 plots, which is a manageable number for MSDM.

6.2.5. Two Sprinkling Scenarios

Note that the number and size of plots can be changed by installing more irrigation equipment. This possibility is accounted for in our models, as explained in Vol. XII. In the PAWN study we generated several sprinkling scenarios, each reflecting a different projection for the area that would be brought under irrigation in the foreseeable future. Each scenario resulted in a different list of plots. The various scenarios are described, along with the procedure for generating the associated lists of plots (the <u>plot files</u>) in Vol. XIV. The minimum area irrigated (approximately 1600 sq km) occurred in the 1976 sprinkling scenario, and gave rise to 1267 plots. The greatest expansion of irrigated area involved almost a quadrupling (to approximately 6000 sq km) over the situation of 1976, and resulted in 1513 plots. The result is called the <u>maximum sprinkling scenario</u>. The two sprinkling scenarios are discussed in App. F of Vol. VA. More details may be found in Vol. XIV.

Because more plots are sprinkled from surface water in the 1990 scenario, there are more aggregate plots after all the aggregation steps described above have taken place. Aggregating the plot file for the maximum scenario resulted in 60 aggregate plots. As mentioned earlier, the 1976 scenario resulted in 57 aggregate plots.

6.3. CALCULATING WATER FLOWS IN AGGREGATE PLOTS

In the foregoing section, we have identified approximately 60 aggregate plots which MSDM must treat in detail. In this and the following sections, we will describe the detailed treatment for a single aggregate plot. Our starting point is the hydrologic cycle. Figure 6.3 depicts the representation of the hydrologic cycle that we use in MSDM.

Note several differences between Fig. 6.3 and the earlier Fig. 6.2. First, we have omitted groundwater irrigation. As discussed in Sec. 6.2.4, plots irrigated by groundwater are not treated in detail in MSDM, and therefore such plots never become parts of MSDM aggregate plots. Second, we have replaced the representation of the surface water system, which was a single block in Fig. 6.2, by two MSDM network nodes in Fig. 6.3. One node--the supply node--serves as a source of supply for both infiltrating water and irrigation water, while the other--the drainage node--serves as a sink for drainage water. Third, we have introduced a separation between the left half of the diagram, which denotes the surface water system, and the right half, which denotes the aggregate plot.

We introduced the separation because most processes take place entirely on one side or the other of this separation. Only surface water irrigation, infiltration, and drainage involve both sides. The computations in MSDM that relate to aggregate plots take advantage of this near-complete separation. Thus in any decade, computation takes place in three steps. First, computations are carried out in the context of the aggregate plot (the right half of the diagram) that result in estimates of infiltration and drainage. At the same time estimates are made of the value of different amounts of irrigation water to the farmer. None of these computations depend in any way on the state of the surface water system at any time during the decade.

Second, computations are made in the context of the surface water system (the left half of the diagram). These computations yield flows in all links of the MSDM network, and water levels at all nodes with storage. Chloride concentrations at every node are also calculated. Among the flows calculated are the amounts of irrigation water actually delivered to each aggregate plot during the decade, and among the chloride concentrations are the concentrations in the irrigation water. These results depend on the infiltration and drainage, and the value of irrigation water, calculated in step one, but are otherwise independent of the state of the aggregate plot at any time during the decade.

Third, computations made in the context of the aggregate plot determine the amounts of water and chloride remaining in the root

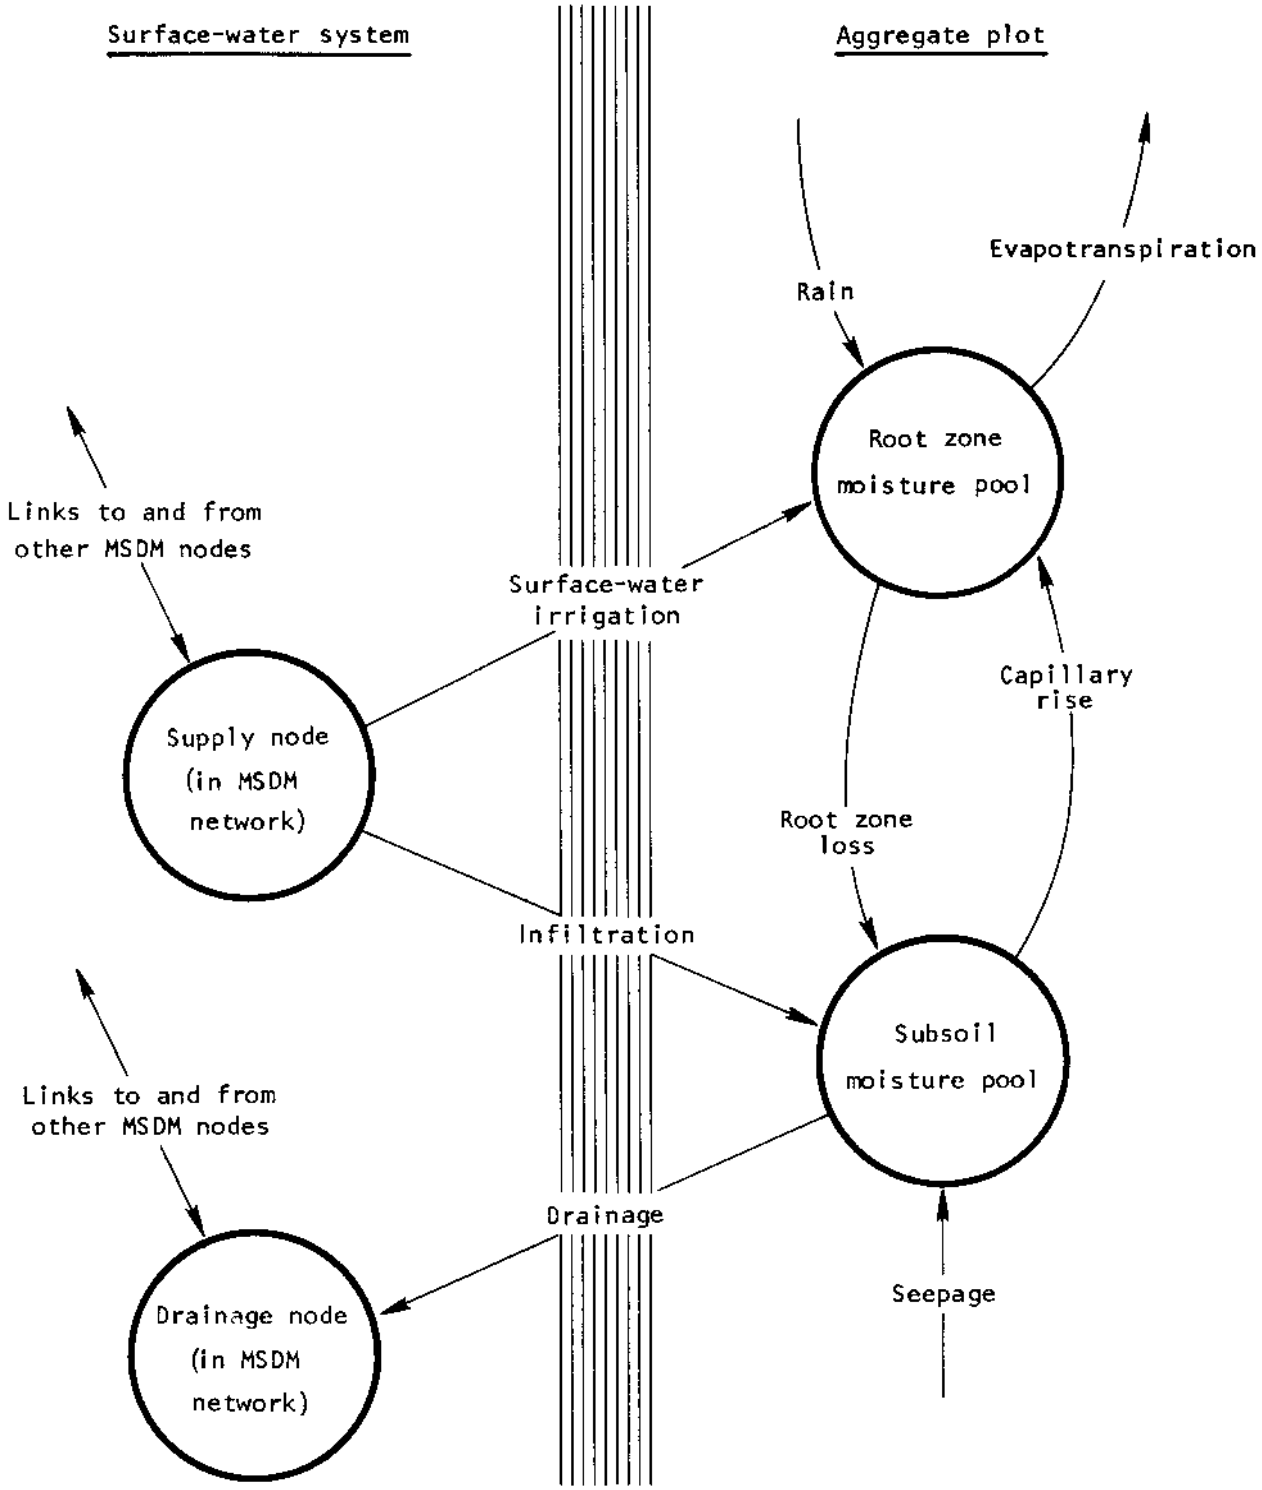

Fig. 6.3--Representation of the hydrologic cycle in MSDM

zone and subsoil at the end of the decade. These depend on the infiltration and drainage, calculated in step one, and the amount and chlorinity of irrigation water delivered, calculated in step two, but they are otherwise independent of the state of the surface water system at any time during the decade.

The next several sections discuss these computations. The present section deals with calculating the flows of water within the aggregate plot, plus infiltration and drainage. These are the flows shown in the right half of Fig. 6.3, except that surface water irrigation is deferred until Secs. 6.5 through 6.7. Section 6.4 discusses the salt flows associated with these water flows.

As before, let us think of an aggregate plot as a vertical column of soil. The top 50 centimeters or so constitute the root zone, the region from which the roots of the crops planted in the column will obtain most of their water. Below the root zone is the subsoil. Deep in the subsoil is a region saturated with water; the depth at which a saturated condition is first encountered is called the groundwater level. Between the groundwater level and the bottom of the root zone is an unsaturated zone.

Because the root zone is bounded both above and below, it is straightforward to define the amount of moisture in the root zone. But the subsoil has no natural lower limit. Obviously, the amount of water in the subsoil depends on how deep the subsoil is considered to be. We have assumed that when completely saturated with water, the subsoil can hold a column of water two or three meters in height. Because the soil takes up some room, a two- or three-meter column of water actually occupies a column of soil perhaps five meters high. This figure was chosen on the advice of one of our Dutch colleagues.[2]

In order to represent an aggregate plot in MSDM, we must be able to calculate the amount of water in the root zone and subsoil moisture pools at the start and end of each decade, and the rates at which water moves into, out of, and between the pools during each decade. This is done by means of two water balance equations, one corresponding to each of the two pools, and a number of auxiliary formulae for calculating various flow rates. The two water balance equations can be constructed easily by referring to Fig. 6.3. Each arrow in the figure corresponds to the movement of water; thus for each arrow extending to or from a moisture pool there must be a term which adds water to or subtracts water from that pool in the corresponding balance equation.

The balance equation for the root zone is:

$$(6.1) \qquad RZM_f = RZM_i + (RAIN - EA + CAP - RZL + SWI)*RZA*10$$

where RZM_i = initial root zone moisture content (i.e., moisture content at the start of the decade), in m^3;

RZM_f = final root zone moisture content, in m^3;

RAIN = rainfall during decade, in mm;

EA = actual evapotranspiration during decade, in mm;

CAP = capillary rise during decade, in mm;

RZL = root zone loss during decade, in mm;

SWI = surface water irrigation during decade, in mm;

RZA = area of root zone, in ha.

The multiplier "10" in Eq. (6.1) converts the product of mm and ha to m^3.

The balance equation for the subsoil pool is:

$$SSM_f = SSM_i + (RZL - CAP + SPG + INF - DRN)*RZA*10 \tag{6.2}$$

where SSM_i = initial subsoil moisture, in m^3;

SSM_f = final subsoil moisture, in m^3;

SPG = seepage into subsoil during decade, in mm;

DRN = drainage during decade, in mm;

INF = infiltration during decade, in mm.

All other terms in Eq. (6.2) have been defined earlier. Note that we multiply the water movement rates to and from the subsoil by the area of the root zone, RZA. This area is the same, of course, as the area of the subsoil, so we need not introduce another variable.

Each of the variables appearing in either balance equations is either provided to MSDM as an input, or must be calculated internally. Table 6.2 shows the source of each term. We discuss each of the calculated terms in App. F of Vol. VA. Irrigation is also treated in Sec. 6.6. We will not discuss the terms which are provided as inputs; for more information, see Vol. XII.

Table 6.2

SOURCES OF TERMS IN AGRICULTURAL WATER BALANCE EQUATIONS

Term	Source
RZM_i	Calculated as final root zone moisture from previous decade, or (for initial decade in a computer run) provided as input
RZM_f	Calculated from Eq. (6.1)
RAIN	Provided as input
EA	Calculated internally; see App. F, Vol. VA
SWI	Becomes a variable in MSDM associated with the node in the MSDM network that supplies this aggregate plot. We call such a variable an <u>irrigation variable</u>; see Sec. 6.6, and App. F, Vol. VA
CAP, RZL	Calculated internally; see App. F, Vol. VA
RZA	Provided as input
SSM_i	Calculated as final subsoil moisture from previous decade, or (for initial decade in a computer run) provided as input
SSM_f	Calculated from Eq. (6.2)
SPG	Provided as input
DRN, INF	Calculated internally; see App. F, Vol. VA

6.4. CALCULATING SALT FLOWS IN AGGREGATE PLOTS

In PAWN, we measure salt in terms of dissolved chloride. In MSDM, salt can be found in both the root zone and the subsoil. As with water, in order to represent an aggregate plot we must be able to calculate the amounts of salt in the root zone and subsoil at the start and end of each decade, and the rates at which salt moves into, out of, and between the pools during each decade.

6.4.1. Salt Balance Equations

The amounts of salt in the two pools are calculated by means of two <u>salt balance</u> equations. Each relates the salt in its respective moisture pool at the end of a decade to the salt present at the start of the decade and the rates of flow of salt into and out of the pool during the decade. The movements of salt to, from, and between the pools plot generally follow the movements of water, with some important exceptions. For the root zone, the balance equation is:

$$\text{(6.3)}\quad RZM_f * CRZ_f = RZM_i * CRZ_i + (RAIN * CRAIN * RZA * 10) + (SWI * CSW * RZA * 10) - (RZL * CRZ_{avg} * RZA * 10) + LOAD$$

where CRZ_i = the initial salt concentration in the root zone (at the start of the decade), in ppm chloride;

CRZ_f = the final salt concentration in the root zone (at the end of the decade), in ppm chloride;

CRZ_{avg} = the average salt concentration in the root zone during the decade, in ppm chloride;

CRAIN = the salt concentration in rain, in ppm chloride;

CSW = the salt concntration in the surface water used for irrigation, in ppm chloride;

LOAD = a salt load introduced directly to the root zone, in grams.

All other quantities appearing in Eq. (6.3) have been defined previously.

According to this equation, irrigation water transports salt into the root zone. In addition, rain contains a small amount of salt, which accompanies it into the root zone. Finally, a salt load is introduced directly into the root zone, without any accompanying flow of water. This represents the salt contained in the fertilizers spread so heavily by farmers on virtually all the cultivated land in the Netherlands.

Salt only leaves the root zone when water is lost to the subsoil. Thus rain falling on the root zone can flush salt into the subsoil, once the rain has filled the root zone to field capacity and root zone loss begins. The same effect can be accomplished by sprinkling the root zone with enough water to cause root zone loss. This is called leaching; it can reduce the salinity in the root zone only to the level found in the sprinkling water.

The other water flows in Fig. 6.3 that involve the root zone do not carry salt. Water evapotranspirating from the root zone leaves as a vapor, and thus carries no salt. Also, we have assumed that water brought to the root zone by capillary rise is salt-free. For discussion of this point, see Vol. XII.

The salt balance equation for the subsoil is the following:

$$(6.4)\quad \begin{aligned} SSM_f * CSS_f = {} & SSM_i * CSS_i \\ & + (RZL * CRZ_{avg} * RZA * 10) \\ & + (INF * CSW * RZA * 10) \\ & + (SPG * CSP * RZA * 10) \\ & - (DRN * CSS_{avg} * RZA * 10) \end{aligned}$$

where CSS_i = the initial salt concentration in the subsoil (at the start of the decade), in ppm chloride;

CSS_f = the final salt concentration in the subsoil (at the end of the decade), in ppm chloride;

CSS_{avg} = the average salt concentration in the subsoil during the decade, in ppm chloride;

CSP = the salt concentration in seepage water, in ppm chloride.

The remaining quantities have all been defined previously.

This balance equation says that the subsoil gains salt by three pathways, root zone loss, seepage, and infiltration. The water entering from the root zone contains salt at the average concentration found in the root zone during the decade. Seepage water contains a concentration of salt that must be specified as an input. As mentoned above, seepage water often contains high concentrations of salt, which it transports into the subsoil. We assume infiltrating water comes from the same MSDM node as irrigation water; hence it has the same salt concentration.

The subsoil can only lose salt in its drainage water. When drainage occurs, it transports salt at the average concentration found in the subsoil during the decade.

6.4.3. Sources of Terms

We have previously discussed the sources of most of the quantities that appear in Eqs. (6.3) and (6.4). All of the water flows and amounts except irrigation were discussed in Sec. 6.3. We will discuss irrigation in Sec. 6.5. Only the salt concentrations and the quantity LOAD remain to be discussed in this section. Table 6.3 indicates the source of each of these remaining quantities.

The average salt concentrations in the root zone and subsoil are calculated in the third step of our three-step procedure, after leakage, drainage and the value of irrigation water have been calculated (step 1), and after the amount and salt content of the irrigation water to be delivered are calculated (step 2). Thus, all terms in Eqs. (6.3) and (6.4) are known except for

$$CRZ_f\ , CRZ_{avg}\ , CSS_f\ , CSS_{avg}\ .$$

Appendix B of Vol. VA deals with the behavior of a pollutant stored in a body of water when both water and pollutant are being added at constant, known rates, and water is flowing out of the body at a constant known rate, at each instant transporting the pollutant with it at the instantaneous concentration in the water body. The equation developed there relates the average pollutant concentration in the water body during the decade to the initial and final concentration, and the various rates of addition and removal of water and pollutant.

In the third step of agricultural calculations in MSDM, we first apply the equation from App. B of Vol. VA to the root zone. This equation and Eq. (6.3) together constitute a pair of simultaneous equations involving two unknown quantities, namely

$$CRZ_f\ , CRZ_{avg}\ .$$

Table 6.3

SOURCES OF TERMS IN AGRICULTURAL SALT BALANCE EQUATIONS

Term	Source
CRZ_i	Calculated as final root zone salt concentration from previous decade, or (for initial decade in a computer run) provided as input
CRZ_f	Calculated from Eq. (6.3)
CRZ_{avg}	Calculated internally; see below, this section
CRAIN	Provided as input
CSW	Calculated by MSDM as the salt concentration at the supply node for this aggregate plot; see Chap. 4
LOAD	Provided as input
CSS_i	Calculated as final subsoil salt concentration from previous decade, or (for initial decade in a computer run) provided as input
CSS_f	Calculated from Eq. (6.4)
CSS_{avg}	Calculated internally; see below, this section
CSP	Provided as input

The two equations can be solved simultaneously to yield values for both. Once the average root zone salt concentration is known, we can do the same for the subsoil. The equation from App. B of Vol. VA is paired with Eq. (6.4), and the two are solved simultaneously to obtain

$$CSS_f , CSS_{avg} .$$

6.5. CROP DAMAGE DUE TO DROUGHT AND SALT

In our models, we have considered two kinds of damage that crops can suffer, damage due to drought, and damage due to excess salt. Our interest in crop damage arises from the fact that it can be reduced to some extent if agriculture is supplied with irrigation water. Indeed, this is the entire justification for irrigation.

6.5.1. Drought Damage

It has been widely observed that an environmental water deficit, also called water stress, reduces plant growth. In simplest terms, water stress occurs whenever the loss of water by transpiration exceeds the rate at which the water is extracted from the soil. When this occurs, the plant's water content is reduced, and its stomata (pores in the leaves) close, reducing the actual transpiration and bringing it in balance with water extraction from the soil. Thus, one measure of water stress is the amount by which the actual evapotranspiration rate falls short of the potential rate. In our models, we assume that the drought damage in a decade is a piecewise linear decreasing function of the ratio of the actual to the potential evapotranspiration rates. Figure 6.4 shows an example of such a function for grass.[3] This function applies to any decade of the growing season. We chose grass because more land is devoted to raising grass for fodder than is devoted to any other crop.

Note that the damage in Fig. 6.4 is expressed as the percentage reduction from a remaining potential yield. In any decade, the remaining potential yield will be the maximum possible yield of the crop reduced by any damage or harvest that has occurred in previous decades.

6.5.2. Salt Damage

Plants are subject to various kinds of damage from high salt concentrations. In the most extreme case, the plant will die or never germinate. But many crops are subject to leaf scorch, tip burn, blossom end rot, or other damage that may cause them to lose most or all of their commercial value. In our models, we have related damage due to salt to the chloride concentration in the root zone. An example salinity damage function for grass is shown in Fig. 6.5. The damage is

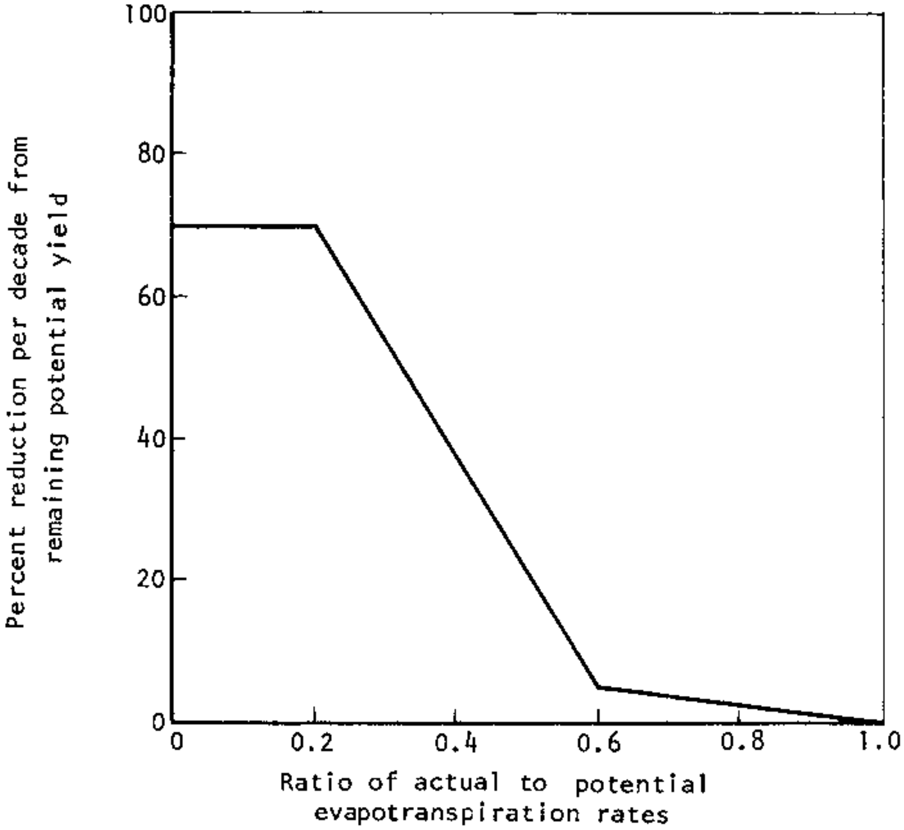

Fig. 6.4--Typical drought damage (for grass in all summer half-year decades)

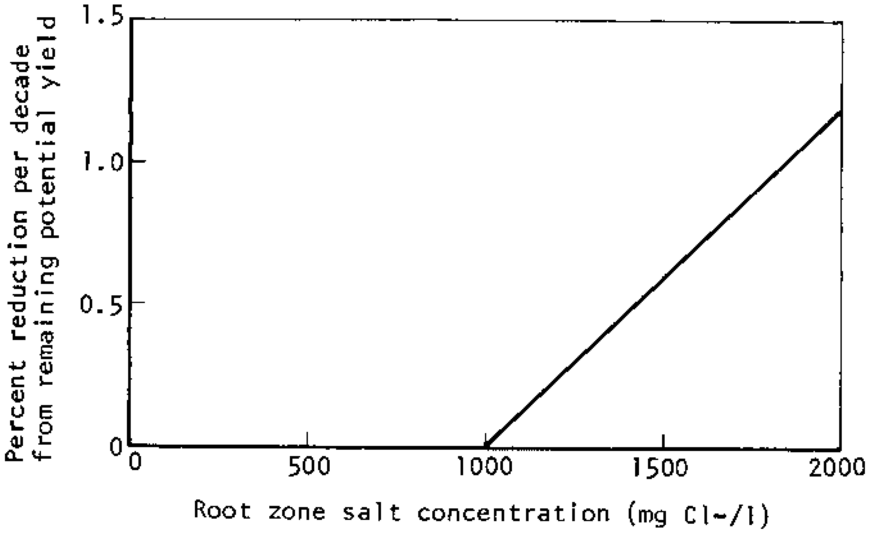

Fig. 6.5--Typical salt damage function (for grass in any decade)

expressed as a fractional reduction in yield from a remaining potential yield, and must be converted to an absolute amount as is the drought damages.

Note that in Fig. 6.5, the scale for percent reduction in yield runs from zero to 2.5 percent only, not from zero to 100 percent. At 2000 mg/l chloride concntration in the root zone (a high value, rarely attained), only 1.2 percent of the crop will be lost per decade to salt damage. If this persists for the entire growing season, the damage will accumulate to approximately 20 percent of the maximum potential yield, so this damage is not negligible. However, in any one decade the potential for salt damage is much less than the potential for drought damage.

Different crops have different sensitivities to salt. Grass is among the least sensitive, while flowers and vegetables grown under glass are the most sensitive. In addition, all crops grown in the open air receive some of their water from rain, which contains almost no salt and hence dilutes the salt in the root zone. By contrast, glasshouse crops receive all of their water from irrigation, which often contains relatively high salt concentrations. For the most part, therefore, salt damage to agriculture is confined to crops grown under glass.

It is almost never necessary to combine the two kinds of damage. Drought is chiefly a problem in areas at higher elevations, where access to rivers and other surface water is limited. In these areas, salt concentrations are ordinarily low. Salt is a problem in lower areas, where seepage of highly saline water occurs. There, water is always available, although not always of good quality.

6.5.3. The Effect of Sprinkling on Damage

It is clear how sprinkling can reduce damage due to drought. Drought damage occurs when the actual evapotranspiration falls below the potential evapotranspiration. But this occurs only when the root zone becomes dry. Thus, by keeping the root zone wet, sprinkling prevents or reduces drought damage.

It is less easy to see why sprinkling could affect salt damage. There are three possibilities. First, suppose that the root zone contains a considerable quantity of salt. If it is allowed to dry out the salt concentration will rise, possibly resulting in salt damage. By keeping the root zone moist sprinkling will prevent this from happening. Second, the farmer may decide to sprinkle so heavily that the root zone reaches saturation capacity, and root zone loss occurs.[4] This will flush some of the root zone salt into the subsoil. Using sprinkling in this way is called leaching.

Finally, sprinkling can cause salt damage. This can happen if the water used for sprinkling contains salt, which most water does.

Sprinkling water is added to the root zone. It evapotranspires, leaving its salt behind. More sprinkling water is added, containing its own salt, which is again left behind when the water evapotranspirates. This concentrating mechanism permits even sprinkling water with very modest salt concentrations to eventually cause salt damage, if leaching is not practiced and if rain does not intervene to dilute the accumulated root zone salt.

6.6. DESIGN AND OPERATION OF SPRINKLER SYSTEMS

Surface water irrigation is accomplished either by deliberately flooding the field, or by sprinkling. In the Netherlands, sprinkling is the method of choice. The purpose of sprinkling is to prevent damage to crops, or looking at it more positively, to help crops achieve more of their potential yield. The farmer will then sell the additional yield, thus realizing an income greater than his neighbor's whose field is not sprinkled. How we calculate the additional yield was dealt with in Sec. 6.5.

However, sprinkling is not free. The farmer must first invest in the equipment that makes sprinkling possible--pipes, pumps, sprinkler heads, etc. Then, whenever he uses his sprinkling system to apply water to his field, he must pay for the energy and labor needed to operate the system. If the system is to pay for itself, the increased yield it makes possible must provide enough additional income to cover both the investment and operating costs of the system.

But the farmer's choice is not simply whether to install a sprinkling system, but which system to install (if any), and how to operate it. The question of what system is appropriate for which crops on what size field consisting of what soil type is addressed in Vol. XIII; we will not discuss it here. But it is necessary to briefly describe the kinds of sprinkling systems available in order to understand the trade-offs a farmer must make in deciding how to operate it. The design and operation of sprinkling systems are the subjects of Secs. 6.6.1 and 6.6.2. Sprinkler system operation is treated at greater length in App. F of Vol. VA.

The sprinkler operating policy discussed in Sec. 6.6.2 and App. F of Vol. VA provides the basis for calculating the demands of agriculture for sprinkling water. However, the fact that agriculture demands a given amount of water does not imply that all of it must be supplied. To decide how much should be supplied, it is necessary to estimate the loss that a farmer would suffer if some of the water were witheld. This problem is addressed in Sec. 6.7.

6.6.1. Sprinkling System Design

Two sprinkling systems, called buis and haspel, are representative of those typically used in the Netherlands. A buis system is a semiportable sprinkling system consisting of a number of movable pipes with sprinkler heads spaced along them. These pipes, called laterals, are usually connected to a larger main distribution pipe that runs through the middle of a field. To irrigate a field, laterals are moved manually from one location to another, tapping the main pipe at a different point each time. The buis system is relatively inexpensive to buy, but requires a substantial amount of labor to use. Figure 6.6 shows two ways a buis system might be laid out in a field.

A haspel system consists of a long flexible hose with a sled-mounted sprinkler at one end, a large take-up reel onto which the hose is wound during sprinkling, a motor for winding the reel, and a pump to deliver water through the hose to the sprinkler at the required pressure. To irrigate a field, the reel is placed by the side of the field, and a tractor is used to stretch the hose with the sled attached to the other side of the field. The pump and motor are then started, and the reel automatically winds up the hose, pulling the sled and the operating sprinkler along with it. Once the sled reaches the reel, one swath across the field has been irrigated. To irrigate another swath, the reel must be dragged by a tractor to another position by the side of the field, and the process repeated. The haspel system has a higher initial cost than the buis system, but it requires less labor to operate. Figure 6.7 shows two typical configurations of a haspel system.

Neither system will irrigate the farmer's entire field at one time. Rather, it is necessary to sprinkle a part of the field, called a set, then move the laterals of a buis system or the reel and hose of a haspel system, and sprinkle the next set. Sprinkling in this manner leaves the last set wettest, with those sprinkled earlier progressively drier. In order to prevent or limit drought damage, there must be few enough sets, and the equipment must be moved frequently enough, to prevent a set from drying out too much between sprinklings.

On the other hand, it is expensive to reduce the number of sets too much, because it requires each set to be larger. This in turn requires that there be more equipment. At the extreme, it would be possible to install a full-coverage system, with enough equipment to sprinkle the entire field from one position, but it would be very costly.

Designing an efficient system thus requires that a balance be struck between buying too little equipment and consequently suffering excessive drought damage, and buying too much equipment and underutilizing it. A critical parameter in choosing a design that strikes this balance is the rate at which the field is assumed to dry out in the absense of sprinkling. This parameter is called the design evapotranspiration rate. If it is set too high, the system must be designed to complete an irrigation cycle around the field too quickly, leading the farmer to buy too much equipment. If it is set

Configuration 1: Small (< 10 ha) field with point water source; likely to occur in highlands

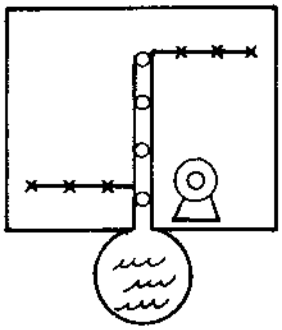

Configuration 2: Large (> 10 ha) field with extended water source (ditch); likely to occur in lowlands

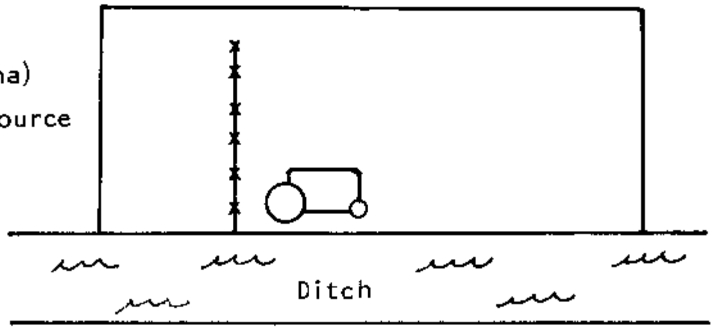

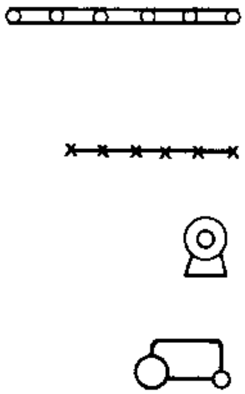

Main distribution pipe with fittings for attaching smaller lateral pipes

Lateral pipe with sprinkler heads

Electric motor with pump

Diesel tractor with pump

Fig. 6.6--Typical buis system configurations

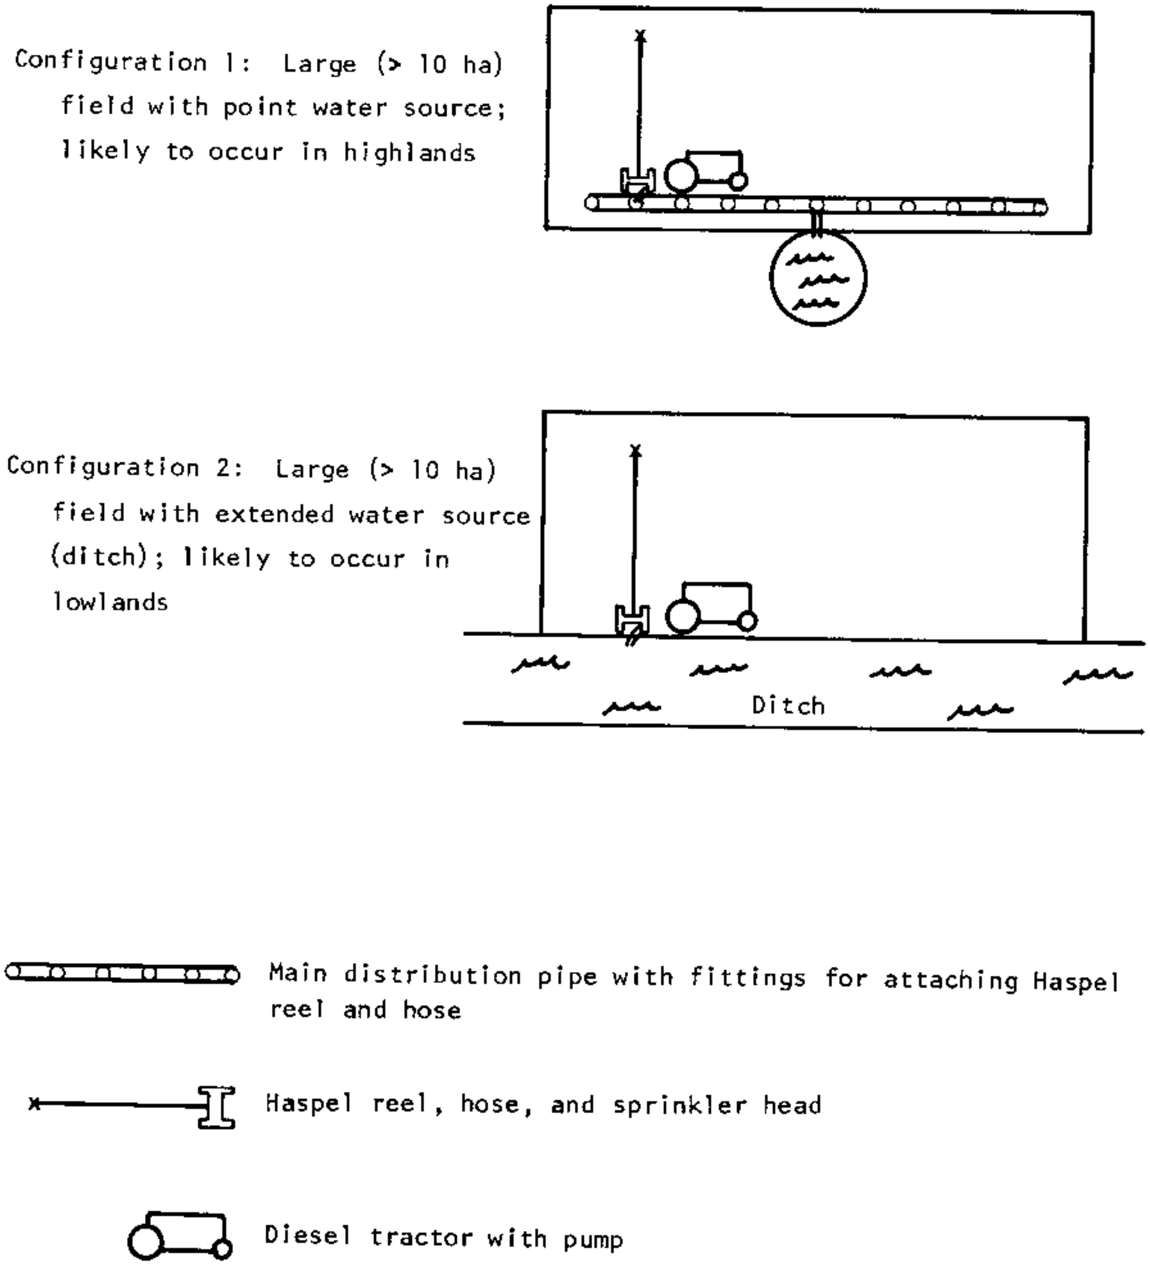

Configuration 1: Large (> 10 ha) field with point water source; likely to occur in highlands

Configuration 2: Large (> 10 ha) field with extended water source (ditch); likely to occur in lowlands

Fig. 6.7--Typical haspel system configurations

too low, the farmer will buy too little equipment, and will be unable to keep his field from drying out under the influence of the actual evapotranspiration rate (as opposed to the design rate) that he will experience during warm, rainless periods.

A reasonable balance is struck when the design evapotranspiration rate is about 2.5 mm/day. At this rate, sprinkling systems are designed to have from 8 to 16 sets, depending on how much moisture the soil can store for the plants and on the frequency with which the farmer is willing to move the equipment. The annualized investment cost of such a system is between 200 Dfl/ha and 300 Dfl/ha, depending on the field size and on how much must be invested to assure a supply of water (e.g., extra pipes).

6.6.2. The Sprinkler Operating Policy

Consider a farmer who must decide one day whether to sprinkle his field. As pointed out above, he cannot sprinkle the entire field in a single day. At most he will be able to sprinkle one or two sets, comprising perhaps 1/10 of his field. If he sprinkles on this day, the sets he sprinkles will be wetter. In addition, he will have moved his equipment into position to sprinkle the remaining sets earlier, so his entire field will be better protected from drought. However, the farmer has spent some labor and energy to put the water on his field. By contrast, if the farmer decides not to sprinkle that day, his field will be drier and hence more susceptible to drought, but he will have saved the cost of his labor and the energy needed to pump the water.

The results of the farmer's decision are uncertain, because at the time of the decision he does not know whether or how much it will rain. If the farmer sprinkles and there is substantial rainfall, the sprinkled part of his field could become so wet that some moisture is lost through root zone loss. Had he not sprinkled in this case, his crops would have been equally well protected from drought, and he would have saved the labor and energy of sprinkling. On the other hand, if the farmer chooses not to sprinkle, and it does not rain, his eventual harvest will suffer and his income will drop. Thus, the decision whether to sprinkle involves both an estimation of the chance of rain, and a trade-off between the cost of sprinkling and the damage to crops. It is interesting to note that the farmer should be willing to let his crop suffer some drought damage, if the damage is less than the labor and energy cost of sprinkling that he saves thereby.

Taking these factors into account, we developed a procedure for estimating the amount of moisture that a farmer would apply to his field in a decade, as a function of the total rain and evapotranspiration during the decade, as well as such parameters as the crop and soil types. We call this procedure the decade sprinkling policy. In spite of the many factors considered in the decade sprinkling policy, and in spite of the computational intricacies, the overall effect of the policy is easily described. It is as if there is a threshold value for the root zone soil moisture content. If the root zone at the start of the decade is dryer than the threshold, the farmer will try to sprinkle enough to maintain the root zone moisture constant, equal to that at the start of the decade. If, on the other hand, the root zone is wetter than the threshold value at the start of the decade, he will let it dry out to the threshold by the end of the decade. That is, the farmer aims at a target root zone moisture, which is the smaller of the starting root zone moisture and the threshold value.

A farmer who follows the decade sprinkling policy will meet the target during all but exceptionally wet decades. Exceptionally wet decades are decades with a large rainfall in comparison to potential evapotranspiration. During such decades, the farmer will sprinkle more than necessary to meet the target, and consequently will leave his root zone wetter than desired at the end of the decade.

6.7. THE VALUE OF WATER FOR SPRINKLING

Our method for estimating the value of sprinkling water is derived in App. F of Vol. VA. One can write an equation which expresses the farmer's total losses if his field starts a decade with a given root zone moisture content, and if a given rainfall and potential evapotranspiration occur during the decade. This equation expresses the total losses as the sum of four terms. The first two terms are the crop damage that can be attributed to this decade, and the cost of sprinkling during this decade. Both of these terms clearly depend on the amount of sprinkling water delivered during the decade.

The third and fourth terms are the expected future crop damage and the expected future cost of sprinkling. These expectations must be taken over all possible combinations of rain and potential evapotranspiration that may occur during the next decade. Both of these quantities depend on the root zone moisture content at the start of the next decade. But this is determined as the root zone moisture at the start of this decade, augmented by rain, capillary rise, and sprinkling, less losses due to root zone loss and the actual evapotranspiration. Thus, depriving agriculture of some or all of its desired sprinkling water may cause crop damage in future decades as well as during the present decade. Also, the farmer will tend to sprinkle more in future decades if his root zone is dryer, as it will be if present sprinkling is reduced. Thus, depriving agriculture of water in the present decade may result in higher future sprinkling costs. These future consequences must be counted in calculating the value of water to agriculture.

Mathematically, we can express the farmer's total loss as:

$$(6.5) \qquad \mathrm{TLOSS}_t = \mathrm{PDam}_t + \mathrm{PCost}_t + \mathrm{FDam}_t + \mathrm{FCost}_t$$

where "PDam" and "PCost" are the present crop damage and sprinkling cost, and "FDam" and "FCost" are the expected future crop damage and sprinkling cost. All four terms are functions of the rain, potential evapotranspiration, and sprinkling during the present decade (denoted by the subscript "t"), and of the root zone moisture content at the start of the present decade.

The procedure for estimating the expected future crop damage and sprinklng cost is of some interest. We calculate the terms recursively. In the last decade of the growing season they are zero, since there is no "future" left. In the last decade but one, the future cost and damage will consist of only whatever may occur during the last decade. This could be calculated with certainty if the last decade's rain, potential evapotranspiration, and sprinkling were known; its expectation can be calculated if it is known how probable different values of these quantities are.

Then we step back one more decade, to the last decade but two. Its future consists of the last decade but one plus the future of the last decade but one. These we can add together, and calculate an expected value of sprinkling cost or crop loss in the future of the last decade but two. We continue stepping backward, decade by decade to the start of the growing season. At the end of the procedure, we will have calculated the expected crop damage and sprinkling cost when the future consists of the entire growing season; this will be a measure of the average yearly loss to the farmer.

As mentioned above, in order to carry out these calculations, we must know the joint probability distribution of rain, potential evapotranspiration, and sprinkling in each decade. In App. F of Vol. VA we have discussed the distributions of rain and potential evapotranspiration. But the distribution of the amount of sprinkling water supplied depends on the managerial strategy adopted by the water manager.

If the water manager always delivers all of the water that agriculture demands, then the decade sprinkling policy can be used to calculate the amount of sprinkling water from the rain, potential evapotranspiration, and initial root zone moisture. This will provide an upper bound on the amount of water provided, and hence a lower bound on the farmer's losses. Since we expect there will only rarely be water shortages, we think this assumption provides a good approximation to the distribution of sprinkling water actually supplied under an optimal water managerial strategy.

At the other extreme, we could assume that no sprinkling water would be delivered in the future. This assumption provides unrealistic results except in the case that the farmer has no sprinkling equipment. In this case, the function TLOSS estimates, decade by decade, how much the farmer can expect to harvest at the end of the year, given the past and present situation.

There is one function "TLOSS" for each decade in the growing season. Of special interest is the function corresponding to the first decade. The root zone moisture at the start of this decade is always close to saturation, and sprinkling hardly ever occurs. Thus, this function gives the farmers total loss for the entire growing season for any combination of rain and evapotranspiration that may occur in that first decade. If we average the losses over all such combinations, weighting each according to its relative frequency of occurrence, we will have determined the average annual total losses for the entire growing season. If we carry out this calculation twice, once using the decade sprinkling policy throughout, and again using no sprinkling, we can estimate the value of sprinkling as the difference between the average annual total losses in the two cases.

We have carried just such an experiment using grass as our crop type and sand as our soil type. A larger area of grass is sprinkled than of any other crop, and a larger area of sand is sprinkled than of any other soil type, so our example can be thought of as typical. We have assumed that no capillary rise occurs, so our example must be thought of as being in the high highlands.[5]

The results of the exercise are as follows. If no sprinkling occurs, the farmer must expect to lose an average of 595 Dfl/ha annually from the potential yield of 3000 Dfl/ha. This compares well with results from DISTAG, which was calibrated to the meager amount of data available. (See Vol. XII for DISTAG validation results. Also, a Distribution Model simulation spanning 47 years indicates an average annual loss of grass from drought of 10 percent, or 300 Dfl/ha, for the northern half of the Netherlands. Much of this area is lowlands, with significant capillary rise that acts to reduce damage. See Vol. II.) If, on the other hand, the farmer sprinkles his grass using the decade sprinkling policy, and is given all the water he demands, he will apply an average of approximately 116 mm of moisture to his field annually, at a cost of 129 Dfl/ha. This will reduce his drought damage to less than 3 Dfl/ha. Thus the net operating benefit of sprinkling is 595 - 131 = 464 Dfl/ha annually.

We have called this the "net operating benefit" because it ignores the investment cost of the sprinkling system. The investment cost must be defrayed somehow, and if he is rational the farmer will insist that it be covered by the net operating benefit. From Vol. XIII we find that the annualized investment cost of a sprinkling system is 200-300 Dfl/ha per year, depending on whether the farmer chooses a buis or a haspel system, and on how large a field he expects to sprinkle. (The system we have assumed for our example is a Haspel, on an area between 10 and 20 ha.) In our example, therefore, the net operating benefit is ample to cover the investment cost of the system.

There is another interesting interpretation of the net operating benefit. It is the amount of money the farmer should expect to lose if he were deprived of all water for sprinkling. Thus, the average net operating benefit divided by the average amount of water applied is a measure of the value of water to the farmer. In our example, this ratio is 464/1160 = 0.4 Dfl/m^3. That is, on the average water is worth 40 cents per cubic meter to our example farmer.

But we wish to find more than the average value of sprinkling. We also wish to determine the marginal value of water to the farmer. By this we mean, how much will the farmer suffer if, during a single decade, we deprive him of some or all of the sprinkling water he requests? This question can be immediately answered by examining the function "TLOSS" for the desired decade, for the marginal value of sprinkling water is merely the negative of the derivative of TLOSS. If the amount of sprinkling water is reduced, the value of "TLOSS" will change, and the amount of change for a unit reduction in sprinkling water is the value of that water. In this connection, the farmer will expect deliveries of sprinkling water to return to normal in the next decade, which we have assumed means that the farmer will receive all that he demands. Thus, to calculate the marginal value of water, we use a function TLOSS derived using the decade sprinkling policy to estimate the distribution of future deliveries of sprinkling water.

The marginal value of water to the farmer depends on how much his sprinkling is reduced. The farmer will demand an amount of water that will reduce his expected total losses for the remainder of the growing season to a minimum. This means that the marginal value of water is zero at that point. (This is a well-known result from economic theory.) But if large reductions in sprinkling occur, the marginal value of water rises steeply. As an example, consider grass grown on sand in the high highlands. In the middle of the growing season, after half the year's yield has been harvested, a reduction of 20 mm from the desired sprinkling raises the marginal value of water to 0.1 Dfl/m^3 (i.e., the farmer would pay up to 10 cents per cubic meter to increase his allotment). In the same circumstances, a reduction of 40 mm raises the marginal value of water to 0.27 Dfl/m^3. Figure 6.8 shows the relation of the marginal value of water to the reduction in sprinkling.

Figure 6.8 refers to a particular kind of farm with a particular sprinkling system. With other crops, or other sprinkling systems, or other soil types, the marginal value will be different. Generally it will be higher, since our example has used a relatively low-value crop and a sprinkling system with low operating costs. It is quite possible to find circumstances under which the marginal value of water is as high as one guilder per cubic meter.

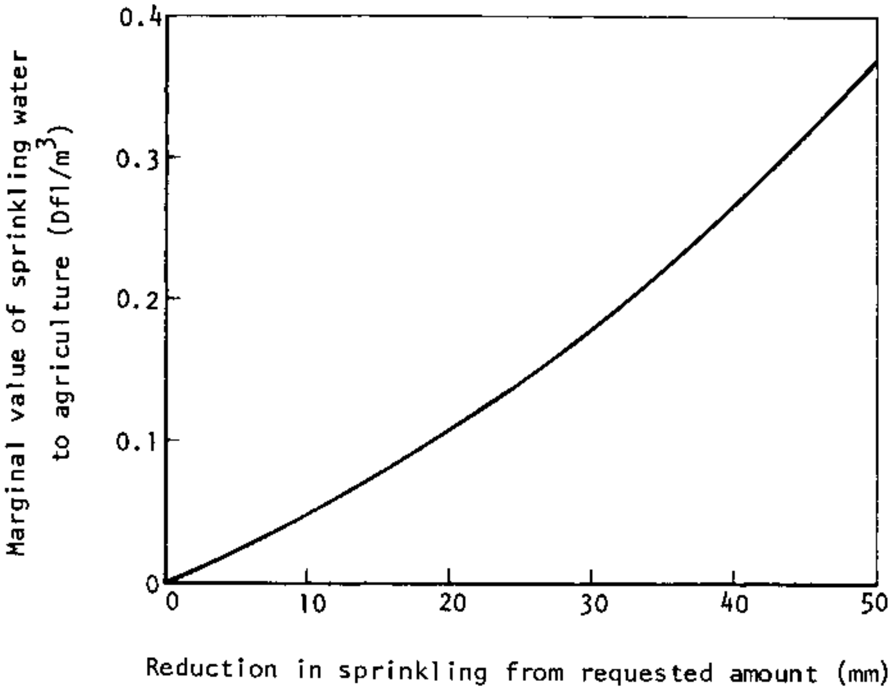

Fig. 6.8--Marginal value of sprinkling water for grass on sand in the high highlands

NOTES

1. Agricultural areas typically extract more water from the network than that needed for consumption. The excess is discharged into the network at points other than the point of extraction, or sometimes discharged into the North Sea. This has the effect of flushing polluted water out of the area, and replacing it with cleaner water. In time of need, flushing could be reduced, and in those instances where flushing water is discharged into the North Sea, this would increase the amount of water available (theoretically) to other users in the Netherlands. But flushing is not, strictly speaking, consumption of water by agriculture, and is not necessarily consumption of water at all.
2. Volume XII reports that this parameter was set to one meter of soil in DISTAG. Our figure was chosen before the DISTAG number had been made final, and due to an oversight, our figure was never adjusted. In any event, this parameter has hardly any effect on MSDM results, because MSDM was never run for more than a single time period in sequence. The initial subsoil salt concentration in a period has much more effect on the final concentration than does this parameter. The parameter is more important in DISTAG because the calculations are done for many time periods in

sequence, each taking as its initial subsoil salt concentration the final concentration from the previous period (except, of course, for the initial period, whose initial concentration must be specified exogenously). Thus, the effect of this parameter has time to become significant.

3. This is not the same function as is reported in Vol. XII; rather, it is an earlier version of the damage function used in DISTAG. When the damage functions were adjusted to their final forms, in the last few hectic weeks of the project, word somehow failed to reach us. In any case, the function used here and the function reported in Vol. XII differ significantly only for evapotranspiration ratios far below one. Such ratios can only occur after several consecutive dry decades in which no sprinkling is permitted, an occurrence so unlikely that the difference in damage functions should have hardly any effect on our results.
4. As we mentioned earlier, DISTAG (correctly) uses field capacity where we have used saturation capacity. This difference will affect the amount of sprinkling that we estimate is necessary to leach the root zone. However, the sprinkling policies considered in PAWN never called for leaching of soils growing crops in the open air. Farmers do leach soils growing crops under glass, but MSDM deals with these crops in an entirely different way than are crops grown in the open air. Indeed, the questions of root zone moisture content, amount of sprinkling, etc., never arise.
5. Actually, most grass is grown in the low highlands, where some capillary rise might occur, especially early in the growing season.

Chapter 7

SHIPPING LOSSES DUE TO LOW WATER FLOWS

The shipping industry depends on water. In order for ships to travel on waterways, the water must be deep enough to allow their keels a safe clearance above the bottom. This problem is worst on the uncanalized rivers Waal and IJssel, because their water levels depend on their flows. When nature fails to maintain a sufficient flow in the Rijn, the water levels in the Waal and IJssel drop. Water levels on these rivers can be varied by only a few decimeters by managerial tactics alone, but this is enough to significantly affect the cost imposed on shipping by low water.

On canals and canalized rivers (Maas and Neder-Rijn), the presence of locks and weirs allows the water levels to be controlled independently of the flows. Ordinarily, this enables water managers to prevent the depth dropping to the point where it affects shipping. Nevertheless, minimum flows are necessary to enable normal operation of the locks. When flows drop still lower, measures must be taken to conserve water at locks, lest the water levels drop. The conservation measures unfortunately also result in delays to shipping.

In this chapter, we discuss both of these problems. Sections 7.1 through 7.6 deal with various aspects of shipping losses due to shallow water. Section 7.7 treats delays at highland locks necessitated by low flows.

7.1. CRITICAL POINTS

Somewhere along each link of the Distribution Model (DM) or MSDM network will be the point with a minimum depth. It is this minimum depth that will determine how large a vessel can traverse the link, and how heavily laden it can be. Of course, the minimum depths on some links are considerably larger than on other, neighboring links. The depths on such links do not impose any limitation on the size of ships or cargoes that can use the network. On the basis of experience, the Dutch have identified about a dozen links whose minimum depths may limit shipping.

7.1.1. Locations

The dozen points that may limit the size of ships and cargoes are called critical points. Two of them occur on the Rijn outside of the Netherlands, one near the mouth of the Ruhr and the other in the section of the Rijn between Cologne and Karlsruhe. (This section is canalized, and the depth is kept fairly uniform throughout, so a unique critical point cannot be identified.) There are also critical points on the Maas outside of the Netherlands.

There are numerous critical points along the Maas, one corresponding to each of the weir ponds (stuwpanden), in which the depth is controlled by regulating the discharges through the various weirs. The limiting depth on the Maas is, of course, equal to the smallest of the depths at these critical points. In the MSDM network, the Maas is represented as having five such points.

There are potentially three critical points on the Waal. One such point lies a few kilometers upstream of Tiel, and when no water is withdrawn from the Waal at Tiel and sent north along the Amsterdam-Rijnkanaal (link 28), this point sometimes has the smallest depth of the three. However, when water is withdrawn at Tiel in substantial quantities, the depth at a point a few kilometers downstream of Tiel then becomes smaller, because the amount of water flowing in that section of the Waal is reduced. (Even when no water is withdrawn at Tiel, the second point is sometimes shallower than the first, depending on the flow in the Waal.) Finally, if the St. Andries Connection were built, and water were withdrawn from the Waal at that point, the depth at a third point still farther downstream would become limiting.

In the MSDM network, only two of the critical points on the Waal need be considered. The first two points, those just above and just below Tiel, have nearly the same depth when there are no withdrawals at Tiel. When there are significant withdrawals at Tiel, the downstream point always has the smaller depth. Thus, little error is introduced by ignoring the depth at the point upstream of Tiel, and considering only the two downstream points. This is the approach we have adopted in MSDM.

The Dutch have identified five critical points on the IJssel. Which of them has the minimum depth will depend on the amount of water that is withdrawn from the IJssel at various points along its course, such at the Twenthekanaal. For the PAWN study, however, not all of the IJssel critical points had to be considered. When no withdrawals are being made along the IJssel, all four critical points have virtually identical depths. When withdrawals are being made, points downstream of the withdrawals always have the smallest depths. But no significant withdrawals from the IJssel occur downstream of the Twenthekanaal. Thus, all points farther downstream may be ignored. In addition, the IJssel depth at the mouth of the Twenthekanaal is always at least as small as the depth at the critical point farther upstream. Thus, only the critical point at the mouth of the Twenthekanaal is considered in MSDM.

The locations of the critical points inside the Netherlands that are considered in MSDM are shown in Fig. 7.1.

7.1.2. Shipping Depths

At each critical point, we calculate a shipping depth, which is the draught of the largest ship that can pass that critical point. The

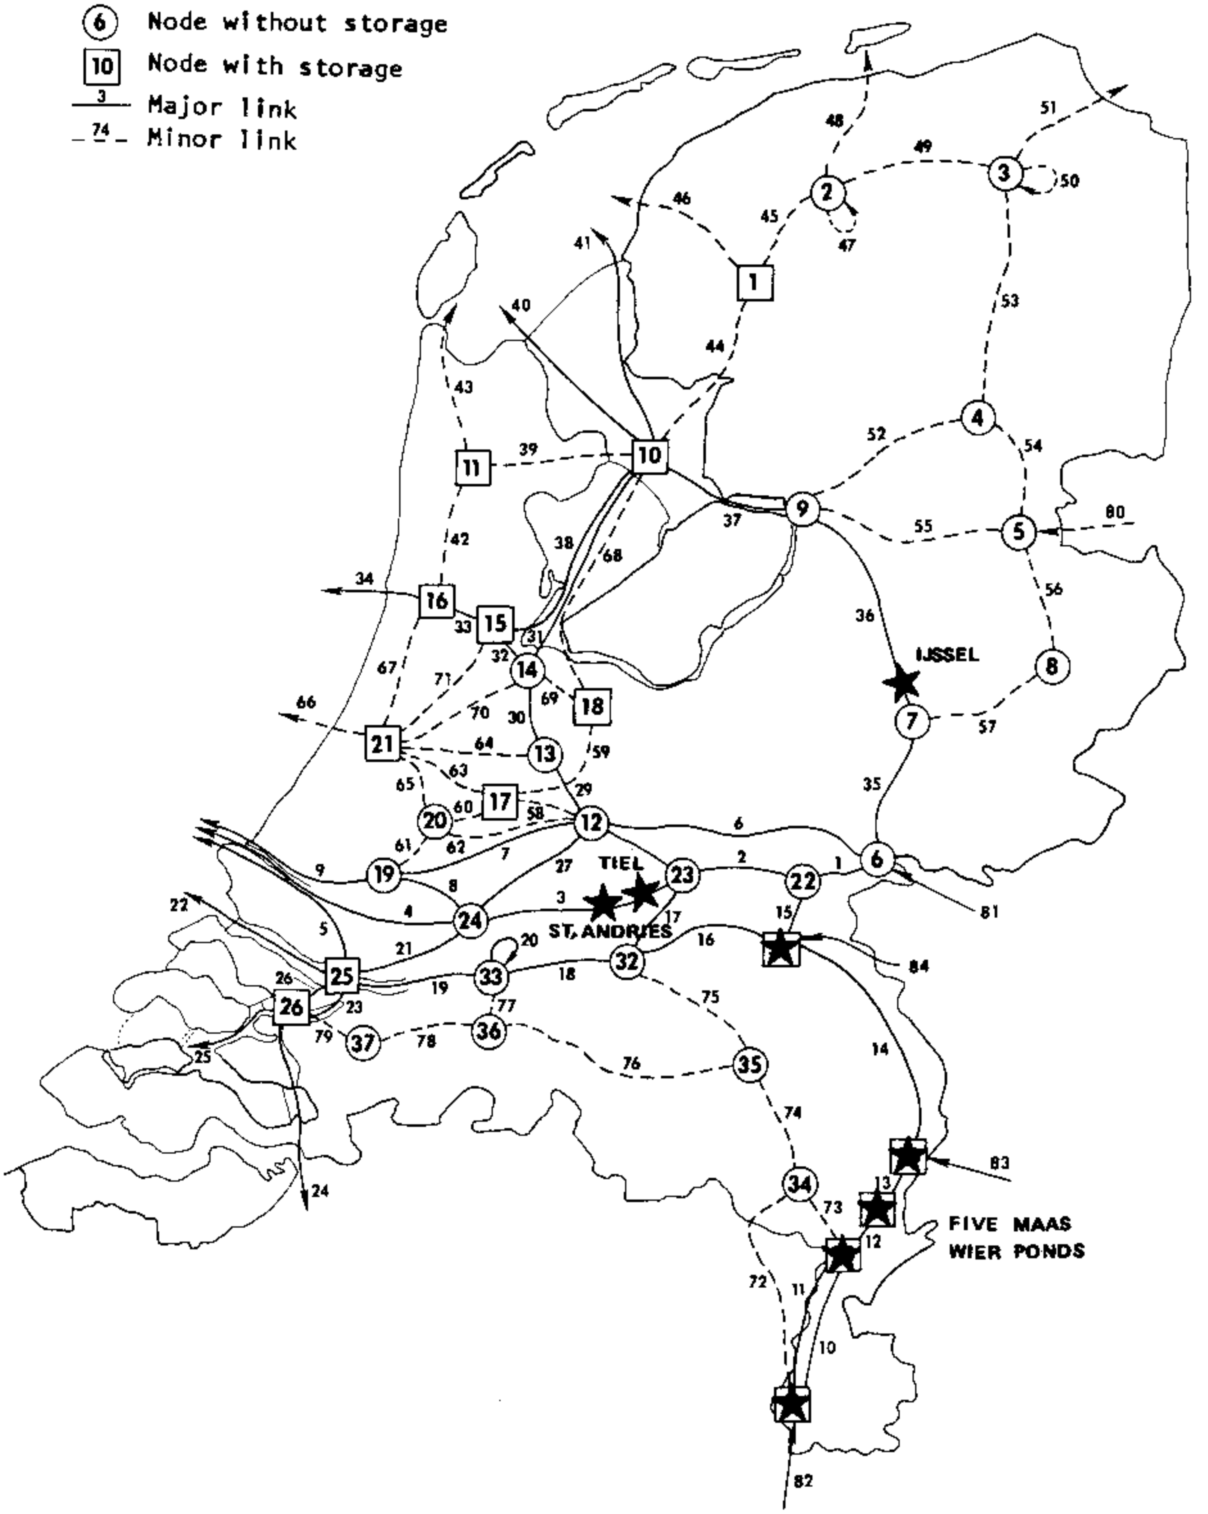

Fig. 7.1--Critical points for shipping in the MSDM network (excludes critical points on the Rijn outside the Netherlands)

shipping depth decreases by one centimeter for every centimeter that the water level falls, but it becomes zero long before all the water is gone from a node with storage (for the critical points on the Maas), or from a link. This is partly because the water depth (as opposed to the shipping depth) at the shallowest point is less than the average depth, and partly because the shipping depth includes an allowance for keel clearance. The allowance depends on the ship size and type, and ranges from 30 to 60 centimeters (Vol. IX).

At each critical point, the depth may be determined from the flows and water levels in the appropriate links and nodes of the network. For example, the shipping depth of the Waal below Tiel can be calculated as:

$$D(TIEL) = 5.986 - 0.0226*F(2) - 0.03*(F(2)-F(3)-F(17))$$

Here, F(2) is the flow in link 2 of the MSDM network, which represents the Waal upstream of Tiel. The expression (F(2)-F(3)-F(17)) calculates the withdrawal of water from the Waal at Tiel, including withdrawals used locally as well as those sent north along the Amsterdam-Rijnkanaal. The shipping depth D(TIEL) is expressed in decimeters (dm). The functions describing the shipping depths at all critical points are developed in Vol. IX. Their implementation in MSDM is described in App. G of Vol. VA.

7.2. SHIPPING ROUTES

Ships do not traverse single links. Instead, they carry cargo along a sequence of links called a <u>route</u>, which goes from the point at which the ship picks up a cargo (the <u>origin</u>) to the point where the ship delivers it (the <u>destination</u>). Each route passes a particular sequence of critical points, corresponding to the links of the network as discussed above, whose associated shipping depths determine the maximum size of the ship that can use the route and the maximum cargo it can carry.

There are actually thousands of distinct routes used by shipping in the Netherlands, by they can be grouped naturally into seven groups of routes. Routes in the same group all pass the same sequence of critical points, while routes in different groups pass different sequences. The development of these groups of routes is described in Vol. IX. The sequences of critical points associated with each of the seven routes, and the routes themselves, are shown in Figs. 7.2a-7.2g.

7.3. SHIPPING LOSSES DUE TO LOW WATER

Consider a ship on one of the seven routes. The minimum of the shipping depths at all critical points along the route will determine the maximum size ship and cargo which can safely use that route.

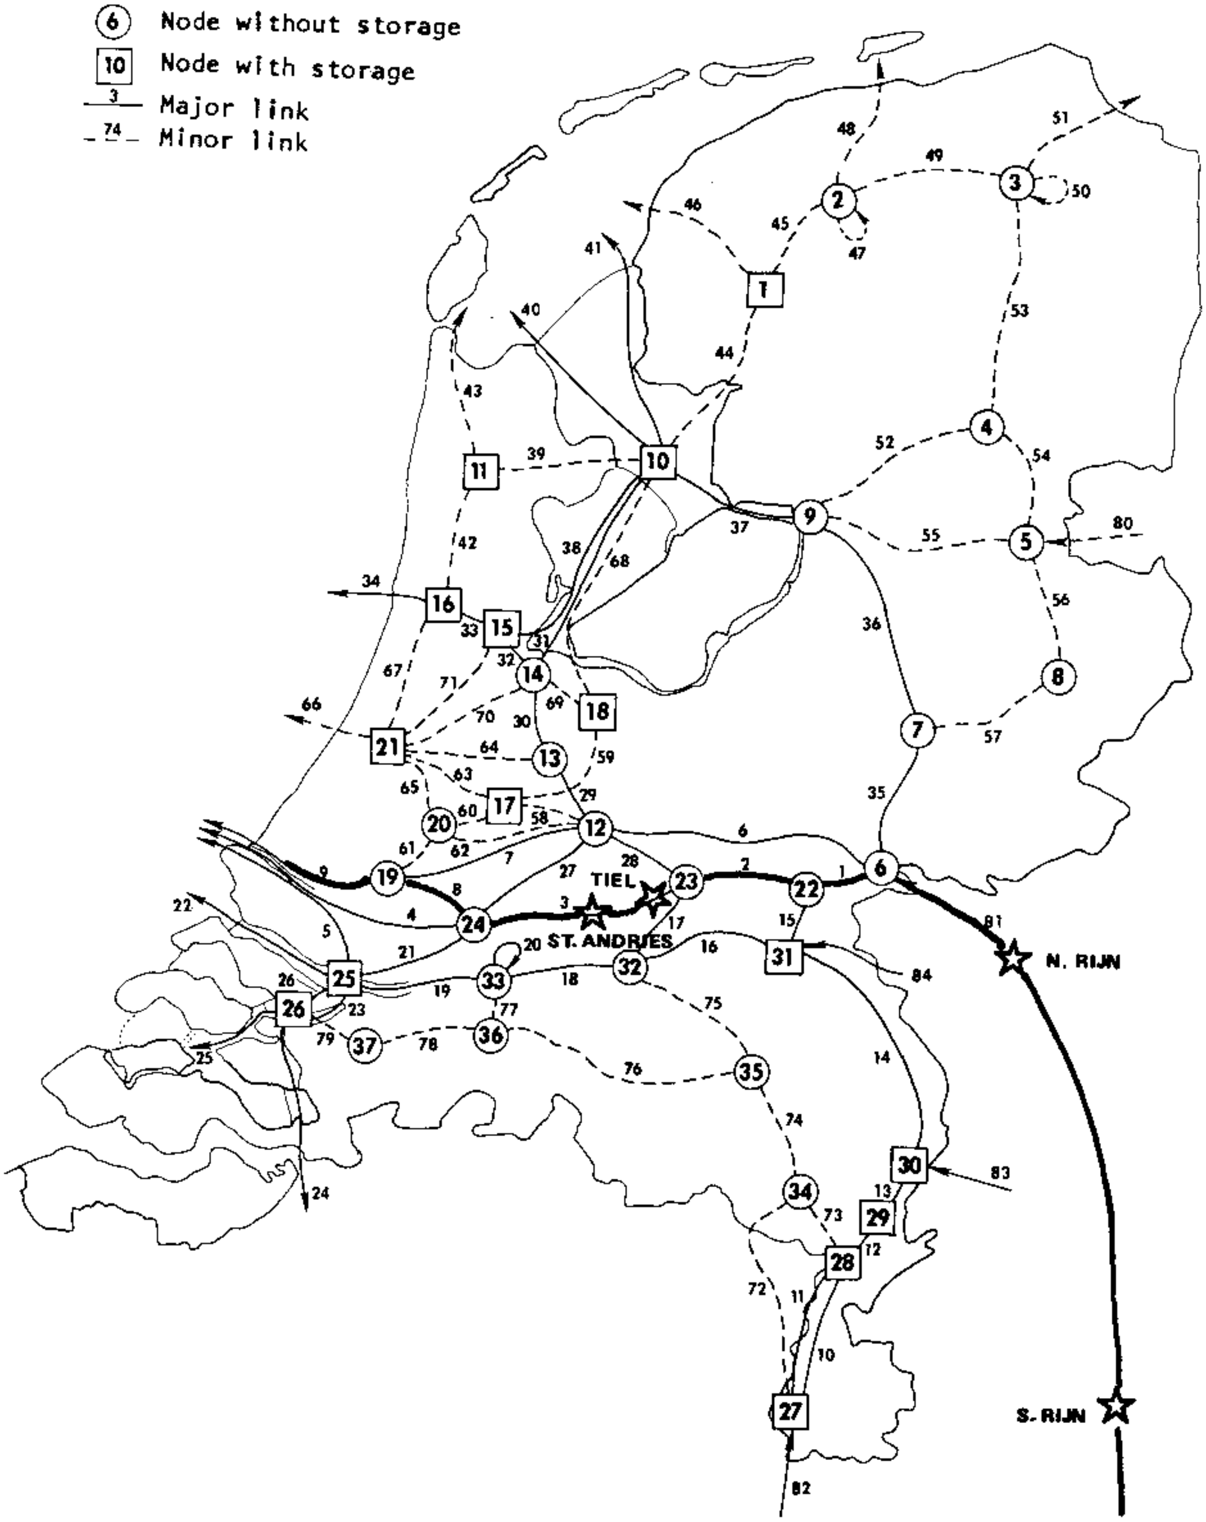

Fig. 7.2a--Seven shipping routes included in the PAWN analysis: southern Rijn-Waal (critical points passed by route shown as stars)

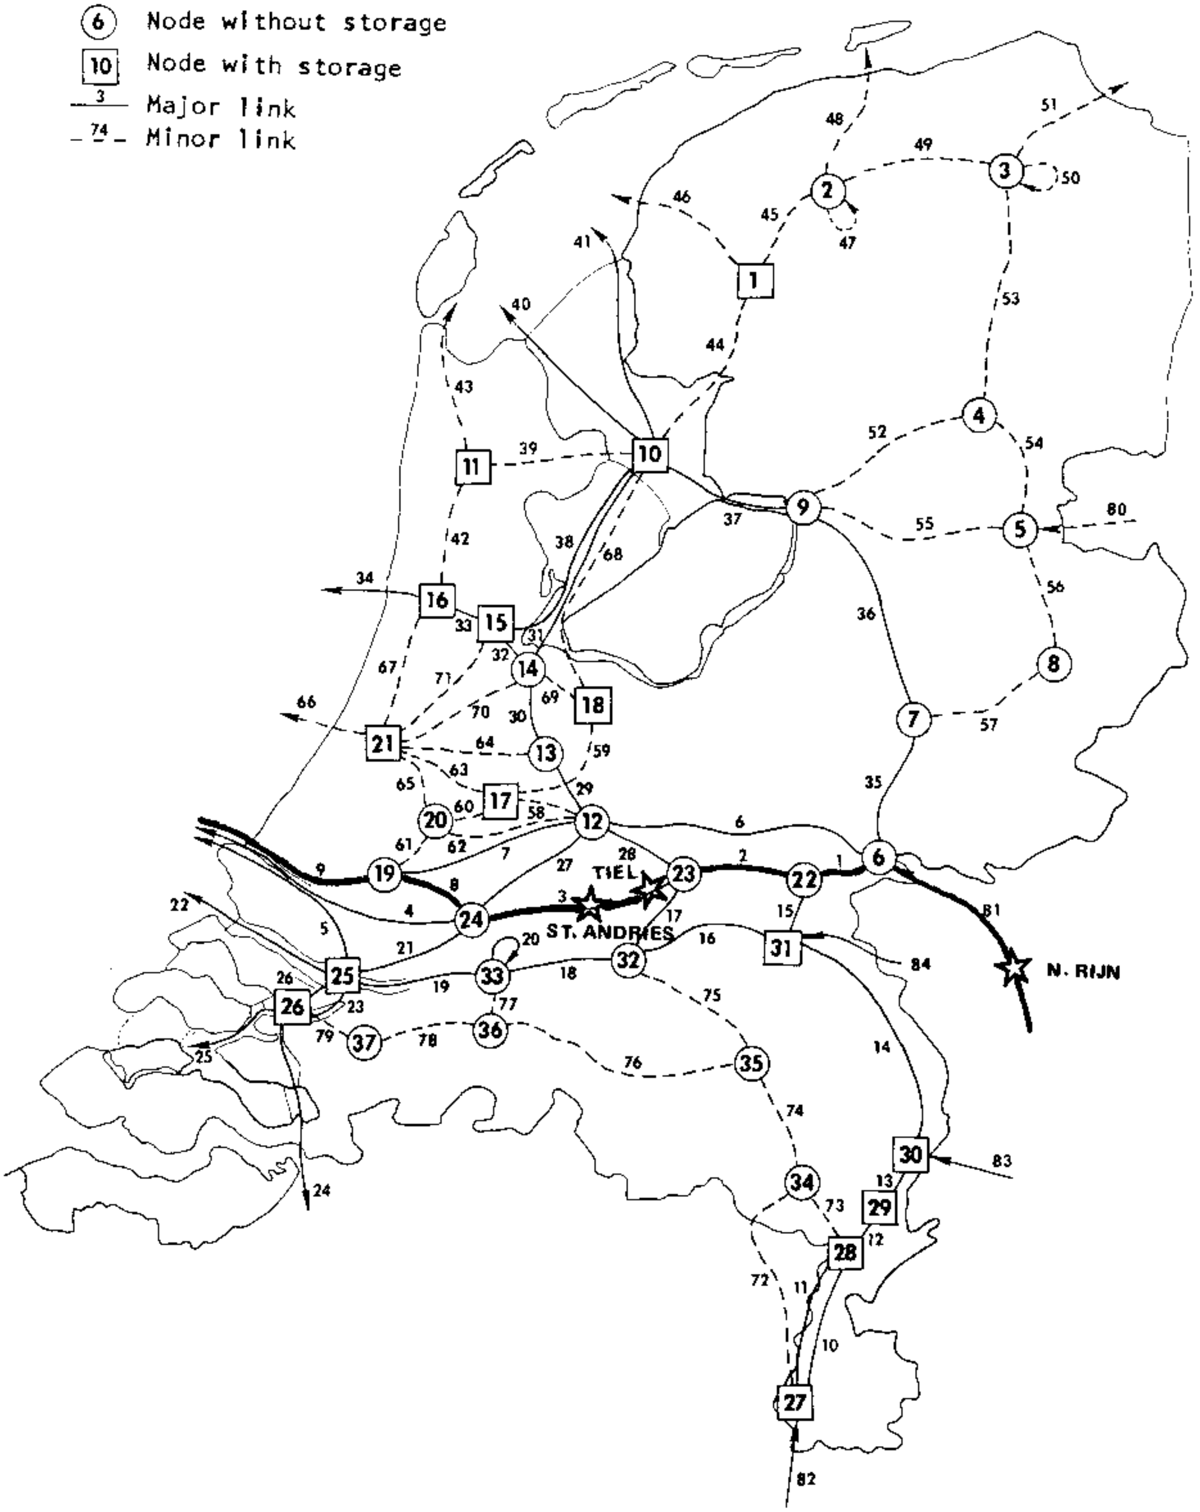

Fig. 7.2b--Seven shipping routes included in the PAWN analysis: northern Rijn-Waal (critical points passed by route shown as stars)

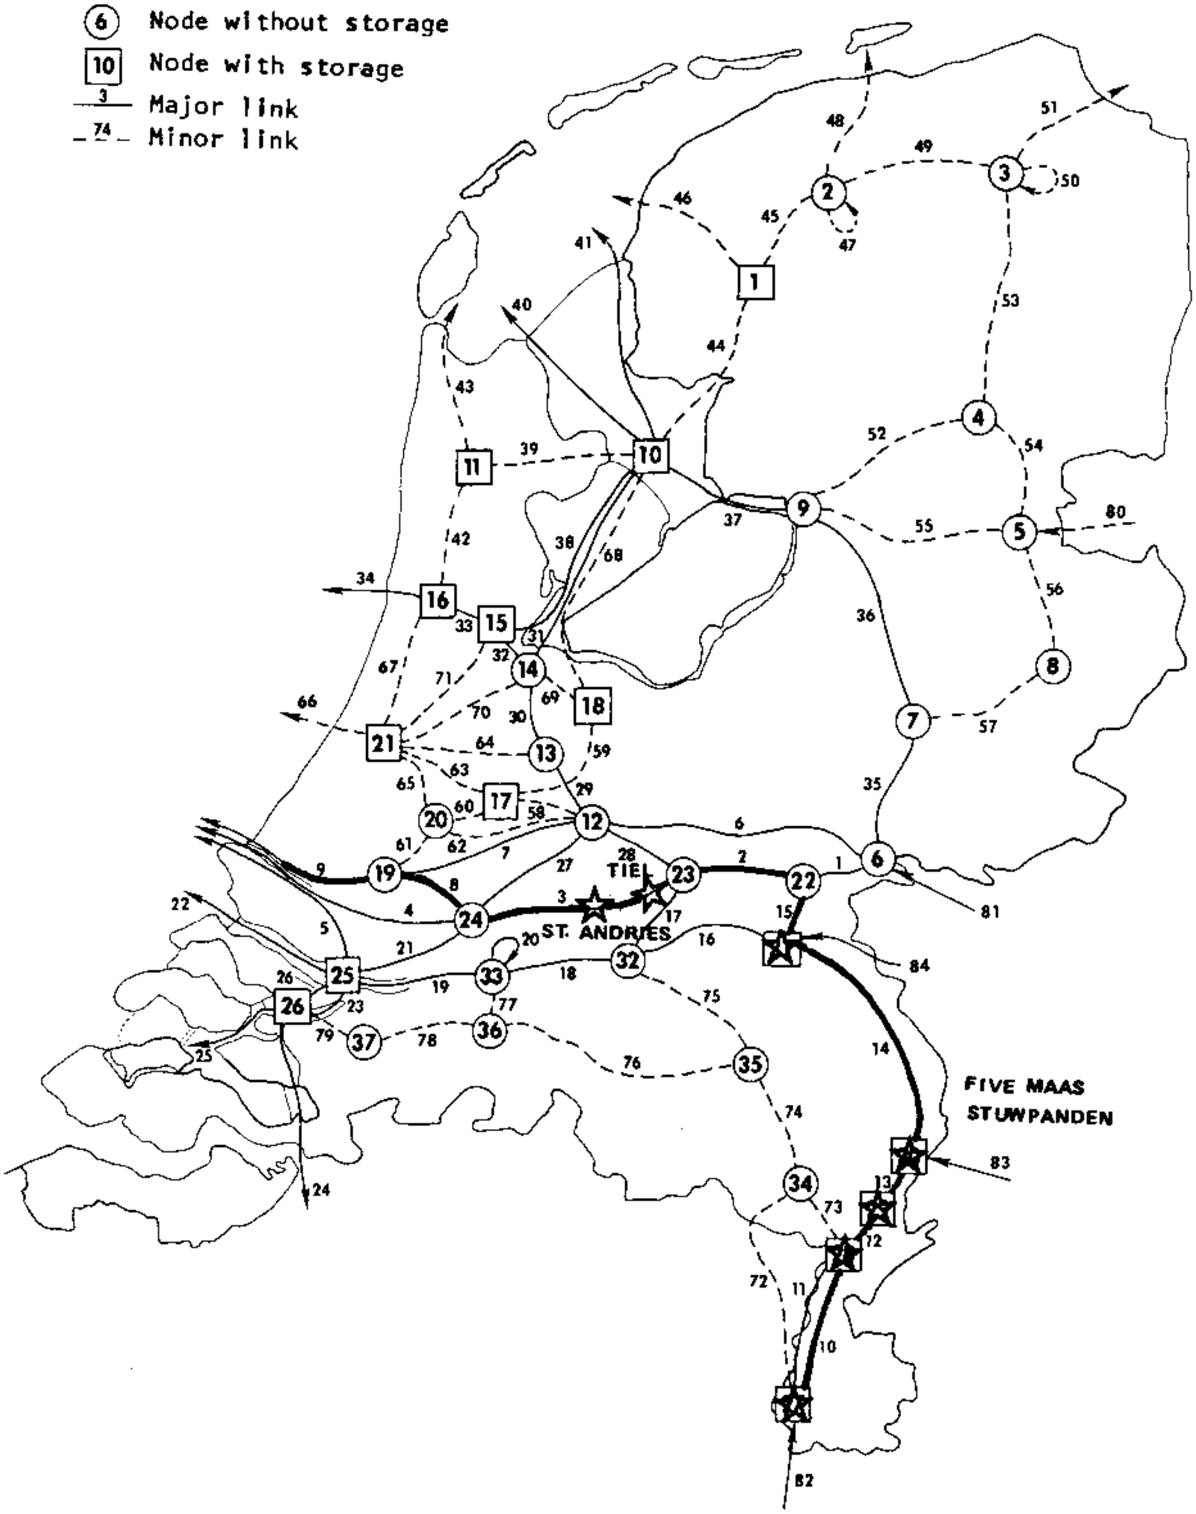

Fig. 7.2c--Seven shipping routes included in the PAWN analysis: Maas-Waal (critical points passed by route shown as stars)

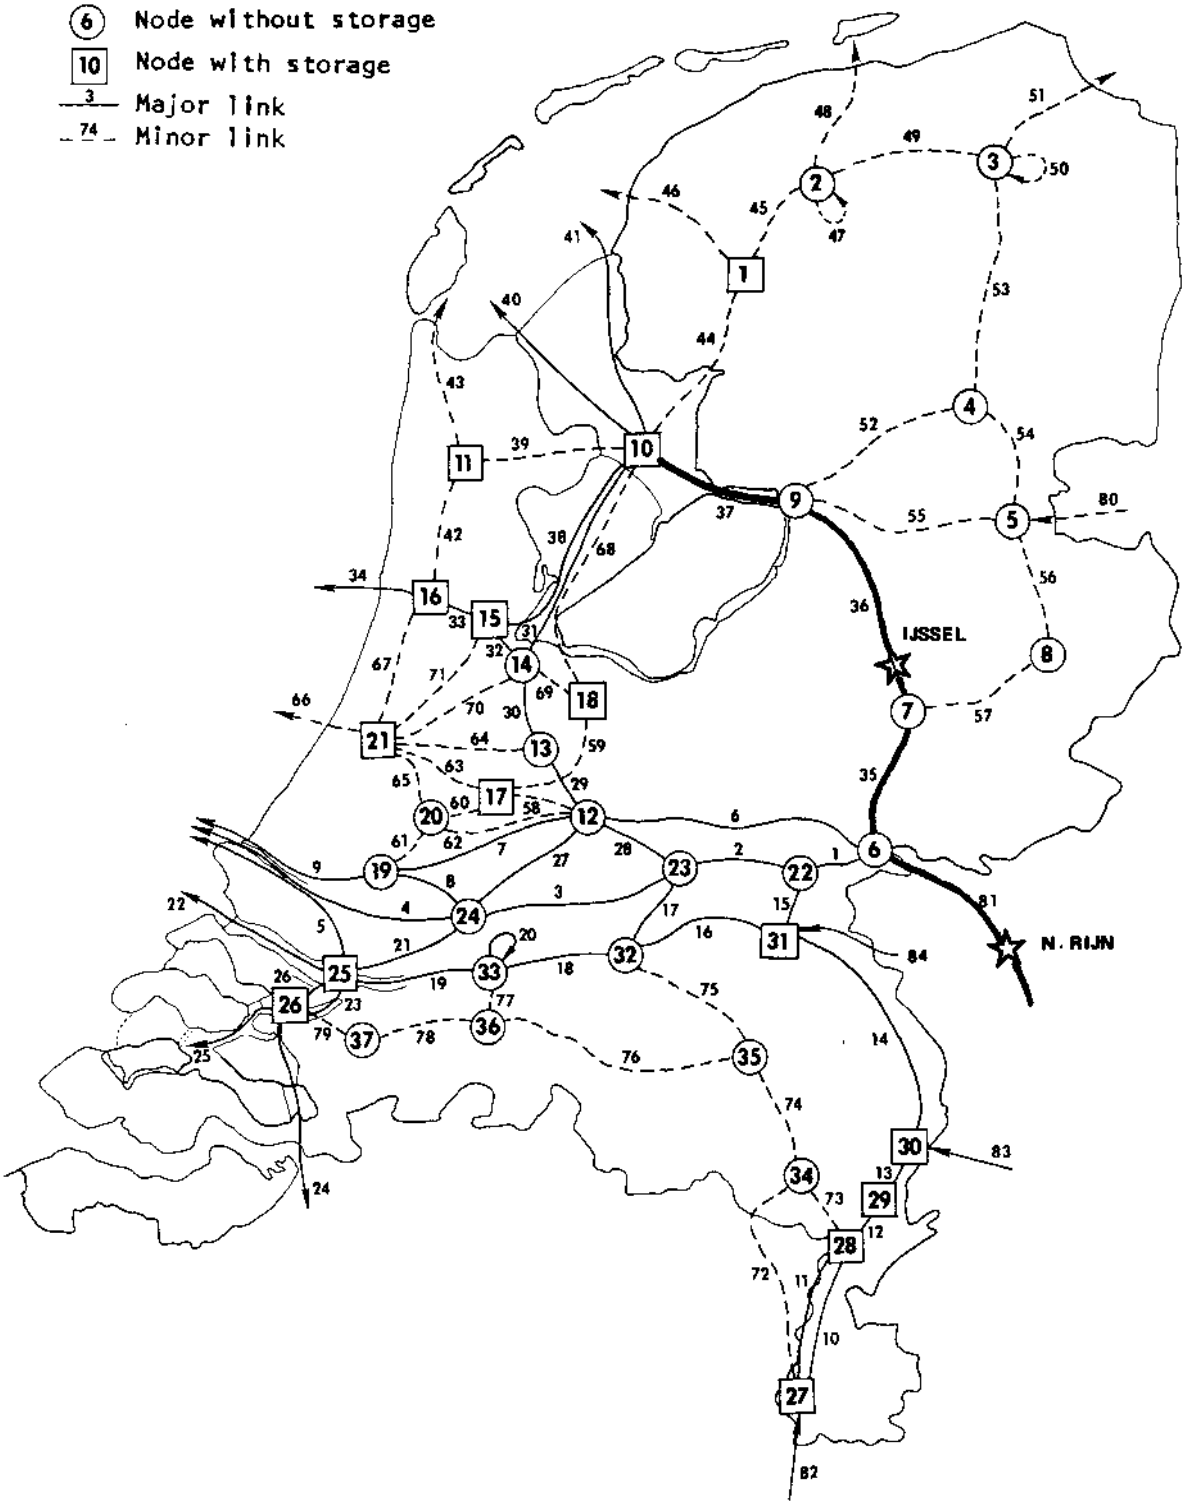

Fig. 7.2d--Seven shipping routes included in the PAWN analysis: northern Rijn-IJssel (critical points passed by route shown as stars)

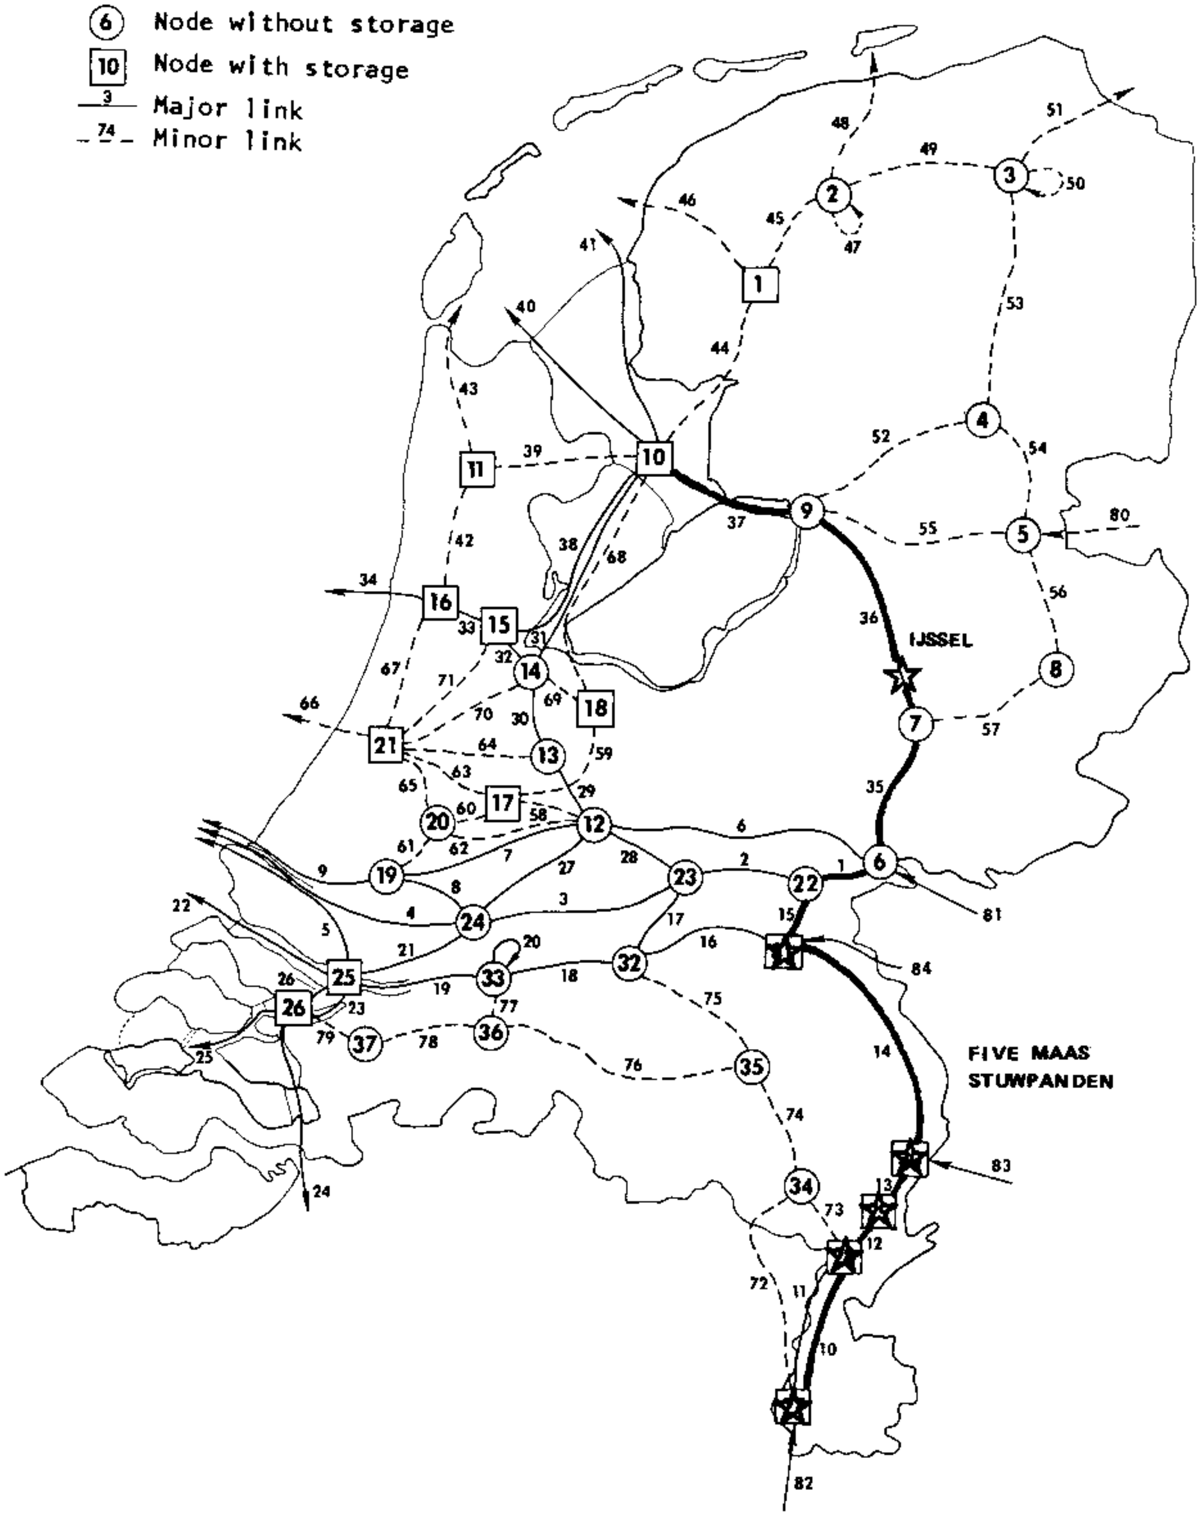

Fig. 7.2e--Seven shipping routes included in the PAWN analysis: Maas-IJssel (critical points passed by route shown as stars)

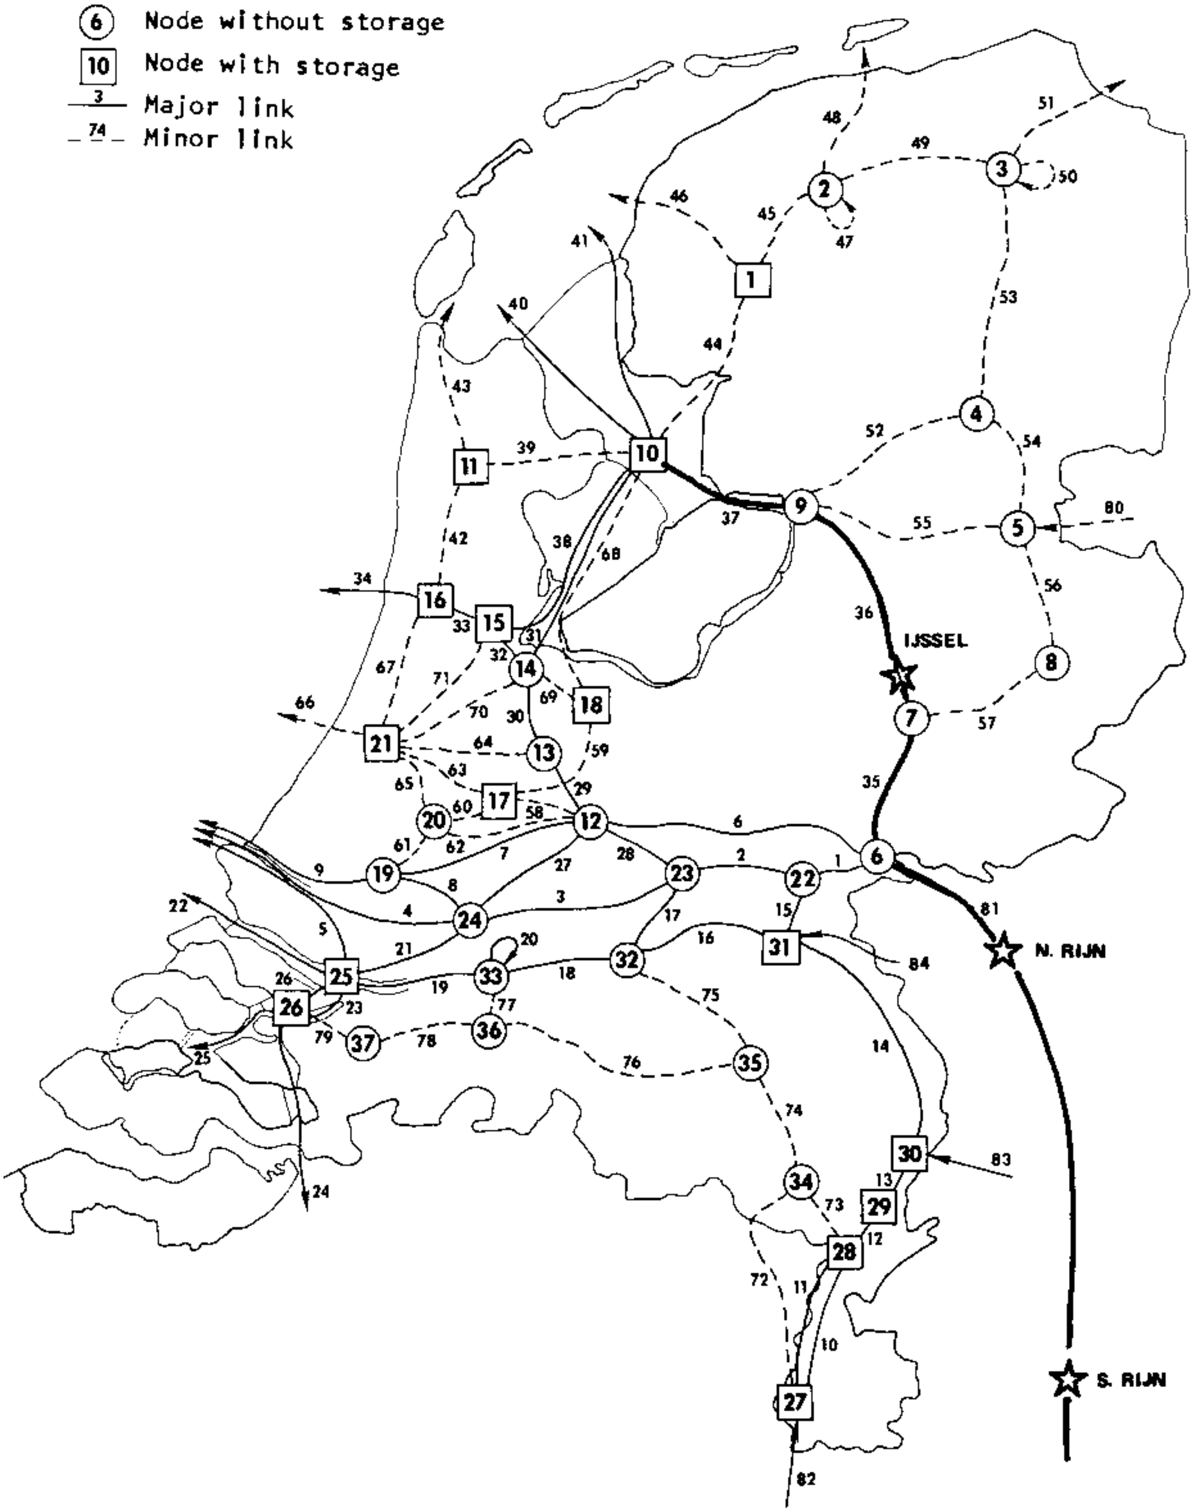

Fig. 7.2f--Seven shipping routes included in the PAWN analysis: southern Rijn-IJssel (critical points passed by route shown as stars)

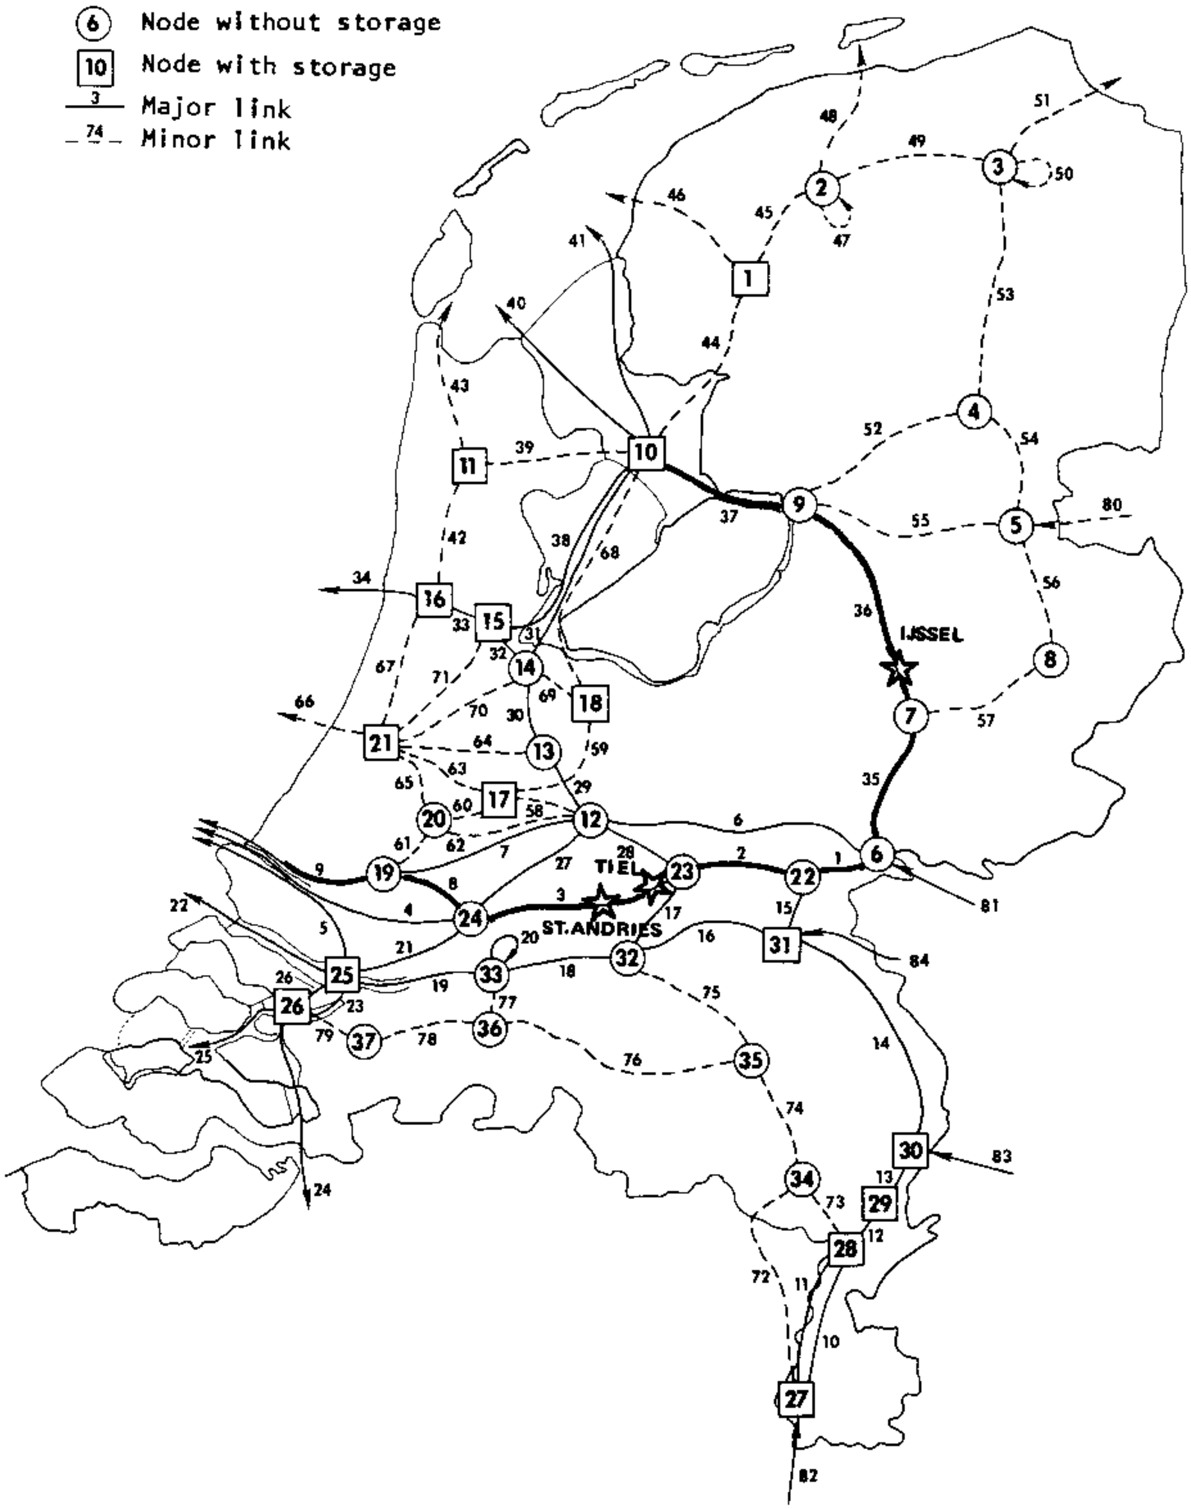

Fig. 7.2g--Seven shipping routes included in the PAWN analysis: Waal-IJssel (critical points passed by route shown as stars)

When that limiting depth is very large, the ship in question will be able to carry a full load on the route. As the limiting depth becomes smaller, that ship will be forced to offload some fraction of the cargo in order to reduce the draught of the ship. The consequence of this requirement is that the ship will have to take more trips in order to transport the same cargo, or equivalently, that more ships must be mobilized in order to deliver the same cargo in the same elapsed time. The use of extra ships results in higher costs for fuel and crews' wages, but no extra service is rendered (i.e., no extra cargo is delivered). Therefore, the occurrence of low water imposes a cost on the shipping industry.

The cost imposed on the shipping industry by low water is only the operating cost for the extra ships needed to transport all the cargo. The investment cost in these extra ships is not included. The reason for this is that the extra ships must be on hand at all times, even when the water is high, in order to be available for use when the water is low. Thus, the investment cost in the extra ships is incurred in all periods, and is not an increment to be paid only in periods of low water.

Notice that, with one exception, it is unimportant who bears the cost. Perhaps the fee for carrying cargo rises when the water falls. Perhaps the fee is constant, and the ship operator must bear the cost. Perhaps the government maintains a fund to reimburse the shipping industry for the extra costs of low water. Regardless of the financial arrangements, there is a real increase in resource costs--fuel and labor--as a result of low water.

The single exception is whether the cost is borne by the Dutch or by citizens of other countries. We have decided to adopt the viewpoint that only costs borne by the Dutch are relevant, and hence we ignore costs to non-Dutch individuals. As discussed in Vol. IX, it is quite difficult to separate Dutch from non-Dutch costs, but it appears that Dutch costs are approximately 62 percent of total shipping costs, in both the 1976 and the 1985 scenarios.

Of course, on any route there will be a mixture of ships of different sizes. Each size class of ships responds to a reduction of the limiting depth on the route in qualitatively the same way as described above. However, the depth at which off-loading must begin will be larger for large ships than for small ships. As the limiting depth on a route decreases from a very large value, at first there may be no additional cost to shipping. When the limiting depth reaches a certain point, determined by the largest size vessels that use the route, shipping will begin to incur a cost. The smaller the limiting depth becomes, the more any size vessel must off-load. Also, at smaller limiting depths, the smaller vessels will have to off-load as well as the larger. For both reasons, the cost increases as the limiting depth decreases, with the rate of cost increase per unit depth reduction accelerating at smaller limiting depths.

In PAWN, we considered two scenarios, one for 1976 and one for 1985. For each scenario, the routes and critical points were the same, but the cargoes to be carried and the fleet of ships available to carry them were different. As a result, the loss functions were different. The fleets and cargoes were specified for PAWN by the Dutch; for further information, see Vol. IX.

In Fig. 7.3 we show the 1976 low water shipping loss function for the Northern Rijn-Waal route, one of the routes with the largest potential losses due to low water. Although the absolute value of losses will differ between routes, the shape of this function is typical of all low water shipping loss functions. Although the function is expressed in Vol. IX as a smooth curve, for reasons of computational convenience we represent it in MSDM as a piecewise liner function. Both the smooth curve and the piecewise linear approximation are shown in Fig. 7.3. Tables can be found in App. G of Vol. VA for both scenarios showing the points that define the piecewise linear approximations used in MSDM for the loss functions for all routes.

7.4. STORAGE COSTS

As the limiting depths on the various routes decrease, more and more ships are needed to fulfill the demand for cargo capacity, according to the explanation we gave earlier. Although the fleet may have excess capacity under high water conditions, under conditions of low water the capacity may be too small to carry all the cargo. When this happens, we have assumed that the least valuable cargo is allowed to pile up at its origin, to be shipped later when the water has risen. But the cargo must be stored somewhere, often in a warehouse or otherwise protected from the weather, and this constitutes another cost of low water to the shipping industry.[1]

In Vol. IX it is shown that the storage costs can be well described as the sum of two functions, one a function of the minimum depth found at critical points on the Waal, and the other a function of the minimum depth on the IJssel. Tables showing the piecewise linear approximations of the storage cost functions used in MSDM can be found in App. G of Vol. VA.

7.5. SEDIMENTATION IN THE WAAL

There is another phenomenon that affects shipping losses due to low water. If water is withdrawn from the Waal at Tiel or at St. Andries, sand will settle to the river bottom just downstream of the withdrawal point. The reasons for this are that the water velocity below the withdrawal point is reduced along with the flow, so the water cannot suspend as much sand as before. Also, the water that is withdrawn is usually taken from near the surface of the Waal, where sand is in lower concentration than nearer to the bottom. Thus, the concentration of suspended sand is higher downstream of the

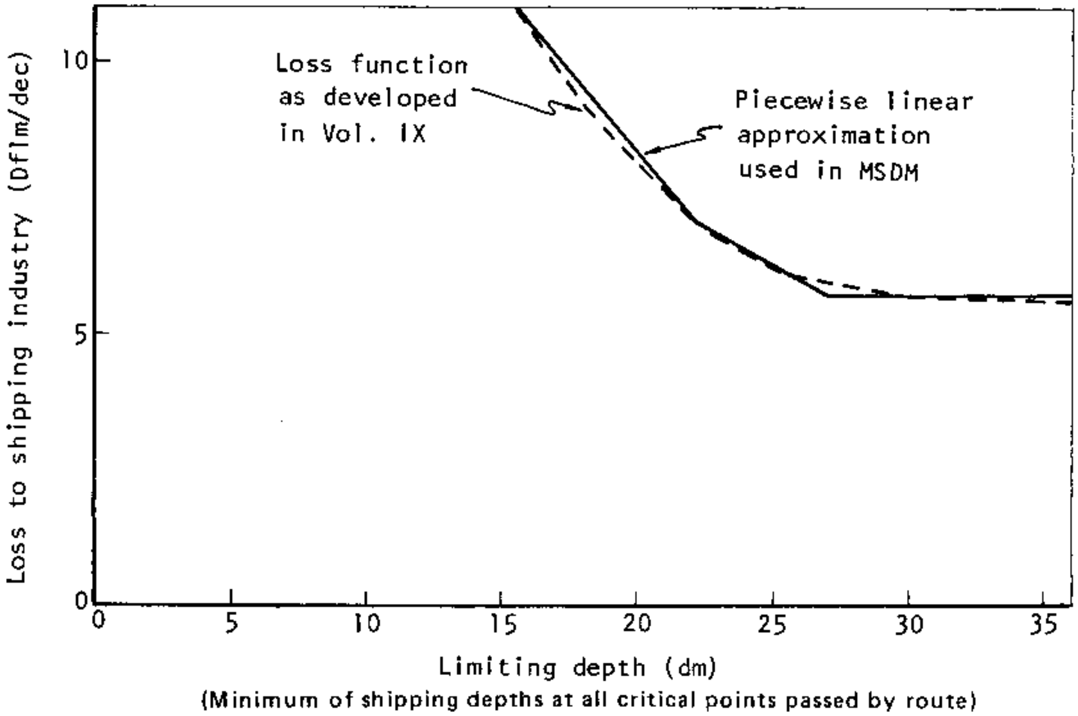

Fig. 7.3--Low water shipping loss function for the northern Rijn-Waal route, 1976 fleet, and demand for cargo transport

withdrawal point, and even had the velocity remained the same, the water would be incapable of suspending so much sand. Thus, a sandbar forms downstream of the withdrawal point, which reduces the depth of the water at that point. This has no immediate effect on shipping, because the point at which the depth has been reduced is not the critical point.

If nothing is done, the sandbar will be washed downstream until, several months later, it reaches the critical point, reducing the depth there. Moreover, several years will pass before the sandbar has completely passed by the critical point, so shipping will bear the cost of a reduced depth on the Waal for a considerable period of time. Alternatively, the sandbar can be removed by dredging before it reaches the critical point. This avoids the cost of a reduced shipping depth, and substitutes the cost of dredging.

7.5.1. Representing Sedimentation Costs in MSDM

The cost to shipping of withdrawal-caused sedimentation is accounted for in MSDM by assigning a cost to the withdrawal. It is understood that the cost will be incurred in later periods than that in which the withdrawal occurs, but clearly the cost must be attributed to the withdrawal, and not to any later event. According to the equations describing the sedimentation phenomenon (Vol. XVIII), the volume of sand deposited during a time period will depend on both the amount of sand already in place at the start of the period, and the magnitude of the withdrawal.[2] The amount of sand already in place can be specified as an initial condition, or, if MSDM is run for a sequence of consecutive decades (an option we have not exercised), it can be remembered as the amount of sand in place at the end of the previous period.

It happens that the cost attributed to withdrawals in the current period are largest if no sand is in place at the start of the period. The cost steadily decreases, the more sand is in place initially. In fact, if enough sand is in place, a withdrawal of a given magnitude in this period will have no effect on cost. Thus, it has proved useful to assume in many cases that a period under study begins with no sand in place. This assumption results in the sedimentation-caused costs being as large as possible.

In MSDM, we may choose either to remove the sandbar by dredging or to allow it to be washed downstream to interfere with shipping in the future. If one chooses to dredge the sand away, the cost of dredging is proportional to the amount of sand removed by dredging, so that the cost attributable to the withdrawal during this period is proportional to the amount of sand deposited during this period. If, on the other hand, one chooses to let the sandbar be swept away naturally by the river current, the cost is a function of the depth reduction at the critical point that one will eventually experience. Of course, we cannot attribute losses due to sand deposited in the past to the withdrawals of the present. Rather, we attribute to present withdrawals the difference between the shipping losses due to the depth reduction that will occur given the accumulation of sand through this period, minus the losses due to the depth reduction due only to the sand in place at the start of the period.

7.5.2. The Cost of Dredging

In Vol. XVIII, the cost of dredging is stated to be 16.2 Dfl per cubic meter of material dredged. If there is initially no existing sandbar, the sand will deposit over an area approximately 260 meters wide (the width of the river at the critical point below Tiel), and approximately 615 meters long. It will build up over this area at a rate which depends on the rate of withdrawal at Tiel, 0.06636 dm/dec for each cubic meter per second withdrawn. This can be converted into a cost of dredging attributable to each cubic meter of water withdrawn of 0.001971 Dfl/m^3. As explained above, because this figure assumes that the period begins with no sand in place, it is an upper bound on the dredging cost attributable to withdrawals at Tiel. We often use this figure because it is so small that adjusting it to reflect different initial conditions has an insignificant effect on our results.

7.5.3. The Cost of Depth Reductions Due to the Sandbar

The cost of the depth reduction due to the sandbar will occur only after the sandbar is washed downstream from the point of formation, near Tiel, to the critical point. This requires many months. The cost of the depth reduction will depend on the depth the Waal would have had without the sandbar--i.e., the depth from which the reduction has taken place. This nominal depth depends on the flow in the Waal, plus any withdrawals that are taking place concurrently. Over the life of the sandbar, the Waal will experience many different flows and hence nominal depths. When the nominal depth is large, the cost of the depth reduction is small; when the nominal depth is small, the cost of the reduction is larger. The total cost to shipping due to the sandbar is the sum of all these costs, taken period by period for the life of the sandbar.

One cannot actually know what the future flows in the Waal will be, so one cannot predict with certainty what the cost must be attributed to a sandbar-forming withdrawal that one wishes to make now. But, because the flows in the Waal follow a known frequency distribution, the expected future cost can be calculated.

If we assume as before that there is no sandbar in place at the start of the period, then we know that a sandbar will build up at a rate of 0.06636 dm/dec for each cubic meter per second of water withdrawn at Tiel. The future expected cost of the sandbar to shipping is not proportional to the height of the sandbar, but it is approximately so for sandbars of less than 1 dm height. Larger sandbars cannot be formed in a single decade by withdrawals at feasible rates. For small sandbars, the cost is 3.703 Dflm/dm of sandbar height for the

1976 scenario, and 6.313 Dflm/dm for the 1985 scenario. We can convert the above figures to costs attributable to each cubic meter of water withdrawn at Tiel of 0.0281 Dfl/m^3 in 1976 and 0.0478 Dfl/m^3 in 1985. Note that allowing the sandbar to interfere with shipping in the future is far more expensive than removing it by dredging, almost 15 times more expensive in 1976, and 25 times more expensive in 1985.

7.5.4. Effect of Infrastructure Changes

In Vol. XVI, two modifications of the groins in the Waal are described that would reduce the effect of sedimentation on shipping. One modification would only increase the depth below Tiel. The other would not only increase the depth, but would reduce the amount of sedimentation as well.

In principle, these tactics could be represented in MSDM by changing the relation of Waal flow to Waal depth, and the relation of withdrawals at Tiel to sedimentation rate. It is not known, however, by how much either modification of the groins would increase the depth of the Waal, or by how much the second modification would reduce sedimentation. In practice, therefore, we have not the information needed to implement these tactics in MSDM.

7.6. VALUE OF WATER FOR REDUCING LOW WATER SHIPPING LOSSES

7.6.1. Critical Points That Limit Shipping

In this section, we will estimate the value of water used to reduce shipping low water losses. The value of water to shipping in any situation depends on how the use of the water changes the depths at the critical points, and on which critical points on each route have the shallowest depths. For example, increasing the depth of the weir ponds of the Maas can at best benefit shipping on the Maas-Waal and Maas-IJssel routes, since only ships on these routes pass the Maas critical points. If the IJssel is shallower than the Maas, then increasing the Maas depth will not improve matters for ships on the Maas-IJssel route. Similarly, if the Waal is shallower than the Maas, increasing the Maas depth will have no effect on ships on the Maas-Waal route. To determine the value of water for preventing low water shipping losses, therefore, we must first determine which critical points on each route are likely be shallowest under what circumstances.

7.6.1.1. The Maas. Most of the time, the shipping depth on the Maas is greater than that on either the Waal or the IJssel, so that changes in the Maas level have no effect on low water shipping losses. If the level of the Maas is decreased enough, however (it can reach a minimum of 15 dm), it can be made somewhat shallower than the Waal. It is more difficult to find circumstances under which the Maas will be shallower than the IJssel, which is in turn shallower than the Waal. Because it is rare that the Maas is shallow enough to interfere with shipping, we will omit any analysis of such situations in what follows,

7.6.1.2. The Rijn Outside the Netherlands. The S. Rijn critical point is always shallower than the Waal, so that ships on the S. Rijn -Waal route are never benefited by changes in shipping depths in the Netherlands. But the N.Rijn point is deeper than the Waal critical point about half the time, and is never very much shallower. We will assume that the Waal depth is always the shallower depth on the N.Rijn-Waal route, which will result in an overestimate of the benefits to shipping of increasing the Waal depth.

7.6.1.3. The Waal. There are two critical points on the Waal, one below Tiel and one below St. Andries. As long as no withdrawals occur at St. Andries, however, the depth below Tiel will be the shallower of the two. Under the present infrastructure, no withdrawals are possible at St. Andries.

According to our assumptions, increases in the Waal depth will benefit shipping on the N.Rijn-Waal and Maas-Waal routes, and may in addition reduce Waal-related storage costs. We therefore take the sum of these three low water loss functions to be the total Waal-related low water shipping loss function.

7.6.1.4. The IJssel. The IJssel is virtually always shallower than every other critical point. Hence ships on any route using the IJssel will benefit from increases in its depth. Thus, increases in the depth at the IJssel critical point will benefit ships on the N. Rijn-IJssel, Maas-IJssel, S. Rijn-IJssel, and Waal-IJssel routes, and will also reduce the IJssel-related storage costs. We take the sum of these five functions to be the total IJssel-related low water shipping loss function.

7.6.2. Managerial Tactics

We can use the total low water shipping loss functions as defined above for the Waal and IJssel to estimate the benefit to shipping of managerial tactics, if we can determine the effect of the tactics on shipping depths in the Waal and the IJssel. (From the discussion above, we do not consider benefits due to changes in the Maas depth.) The tactics we will consider have the direct effect of changing flows on MSDM network links. The effect of changes in flows on depths can be calculated using the depth equations from App. G of Vol. VA. The three managerial tactics we examine are:

> Reducing withdrawals at Tiel, to increase the water level in the Waal. This also reduces sedimentation costs, but we ignore these in this section. The reader may decide whether

or not to deal with sedimentation by dredging, and add the appropriate incremental cost (from Sec. 7.5) to the numbers presented below.

- Reducing withdrawals at Eefde, to increase the water level in the IJssel.
- Closing the weir at Driel on the Neder-Rijn. This increases the flows, and hence water levels, on both the Waal and the IJssel.

The value of water employed in any of these ways will depend on the situation. We will treat these tactics as changes in flows from a nominal situation, in which the withdrawals at Tiel are 10 m^3/s, the withdrawals at Eefde are 10 m^3/s, and the weir at Driel is set so that the flow in the Neder-Rijn is 50 m^3/s. We will determine the value of water for this nominal situation for a range of flows in the Rijn at Lobith from 600 m^3/s (smaller than has ever been observed) to 1500 m^3/s (large enough that shipping will not benefit from further increases).

The effects of the three tactics on the shipping depths in the Waal and the IJssel are as follows. A reduction of the withdrawal at Tiel by 1 m^3/s for one decade amounts to a total reduction of 876,000 m^3, if we take each decade to be 1/36 of a year. That reduction, according to the equations relating depths to flows (see Vol. IX, or App. G of Vol. VA), should increase the depth of the Waal by 0.03 dm. These figures can be combined to determine the increase in depth per cubic meter reduction in the withdrawal. This can be multiplied by the slope of the total Waal-related loss function described above, which is the loss to shipping per decimeter reduction in depth.

Of course, the slope of the total Waal-related loss function differs according to the depth at which it is calculated. At large depths, a change in depth hardly affects shipping losses at all, whereas at smaller depths, the same depth change can be significant. Thus, the slope used in the product described above will depend on the depth in the Waal prior to reducing withdrawals at Tiel. This depth in turn may be calculated from the flow in the Rijn at Lobith. Thus, the outcome of our calculations will be the value of reducing withdrawals at Tiel, as a function of the Rijn flow at Lobith.

The effect of reducing withdrawals from the IJssel at Eefde is calculated in the same way. In this case, however, a reduction of 1 m^3/s (876,000 m^3/dec) in the withdrawal has the effect of increasing the depth by 0.1 dm (see App. G of Vol. VA). In addition, we use the total IJssel-related shipping loss function, and not the Waal-related function.

Reducing the flow in the Neder-Rijn by closing the weir at Driel affects both the Waal and IJssel depths. A reduction of 1 m^3/s (876,000 m^3/dec) in the Neder-Rijn flow increases the Waal flow by 0.58 m^3/s, and the IJssel flow by 0.42 m^3/s (see App. A of Vol. VA). For each m^3/s increase in the Waal flow, the Waal depth increases by 0.0226 dm. On the IJssel, one m^3/s increase in flow results in a 0.0922 dm increase in depth. These figures can be combined as before to yield the value of reducing the Neder-Rijn flow separately on Waal shipping and IJssel shipping, and the results can be added to yield the overall value.

7.6.3. Value of Water for Reducing Low Water Shipping Losses

In Fig. 7.4, we show the results of the above calculations for the tactic of reducing withdrawals at Tiel. The first thing to observe is that under the 1985 scenario, water taken from Tiel withdrawals is much more valuable to shipping than under the 1976 scenario. The second observation is that the value declines sharply as the Rijn flow rises. For Rijn flows below 700 m^3/s at Lobith, an event that occurs during only 0.3 percent of all decades (approximately one decade in ten years), the value of reducing withdrawals at Tiel is nearly 4 cents/m^3 in 1976, and 6 cents/m^3 in 1985. But by the time the Rijn flow rises to a value that could be expected to occur reasonably often, say three decades during the year (a flow of approximately 1060 m^3/s), the values of water have dropped to 0.3 cents/m^3 in 1976 and 1.0 cents/m^3 in 1985. The reader should keep in mind, of course, that these values do not include an allowance for sedimentation, which will also be reduced by a reduction in the withdrawal at Tiel. This additional value is especially significant if the sandbar is not removed by dredging.

The results of our calculations for withdrawals from the IJssel at Eefde are shown in Fig. 7.5. Again the value of such reductions is much higher under the 1985 scenario than under the 1976 scenario. It is perhaps surprising that the value is considerably higher than the corresponding value of reductions in withdrawals at Tiel, especially when one considers how much greater is the value of shipping using the Waal than that using the IJssel. The explanation lies in the fact that the IJssel depth is over three times more sensitive to changes in withdrawals than is the Waal depth. Finally, also unlike Fig. 7.4, the value of reducing withdrawals remains high even at relatively large, and hence likely, Rijn flows. At a flow of 864 m^3/s, for example (lower flows have occurred in 47 decades during the 47 years from 1930 through 1976), the value of reducing withdrawals at Eefde is about 6 cents/m^3 in 1976, and over 15 cents/m^3 in 1985.

Finally, in Fig. 7.6 is shown the value of diverting water from the Neder-Rijn into the Waal and the IJssel. As might be expected, this tactic has results intermediate between those of the previous two tactics. As before, the value of the water diverted by this tactic is larger in 1985 than in 1976. The value does decline as the flow in the Rijn increases, but much of the value is retained even at rather large and likely flows. It should be noted that exercising this tactic has no effect on sedimentation, and so no allowance should be made for this phenomenon.

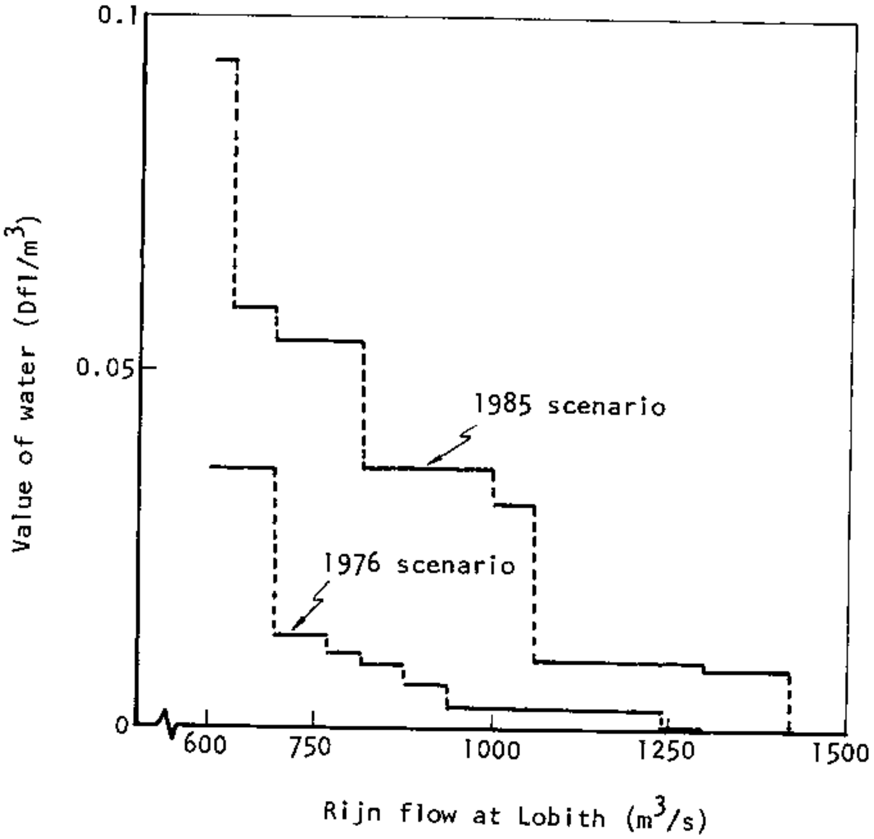

Fig. 7.4--Value to Waal shipping of reducing withdrawals at Tiel

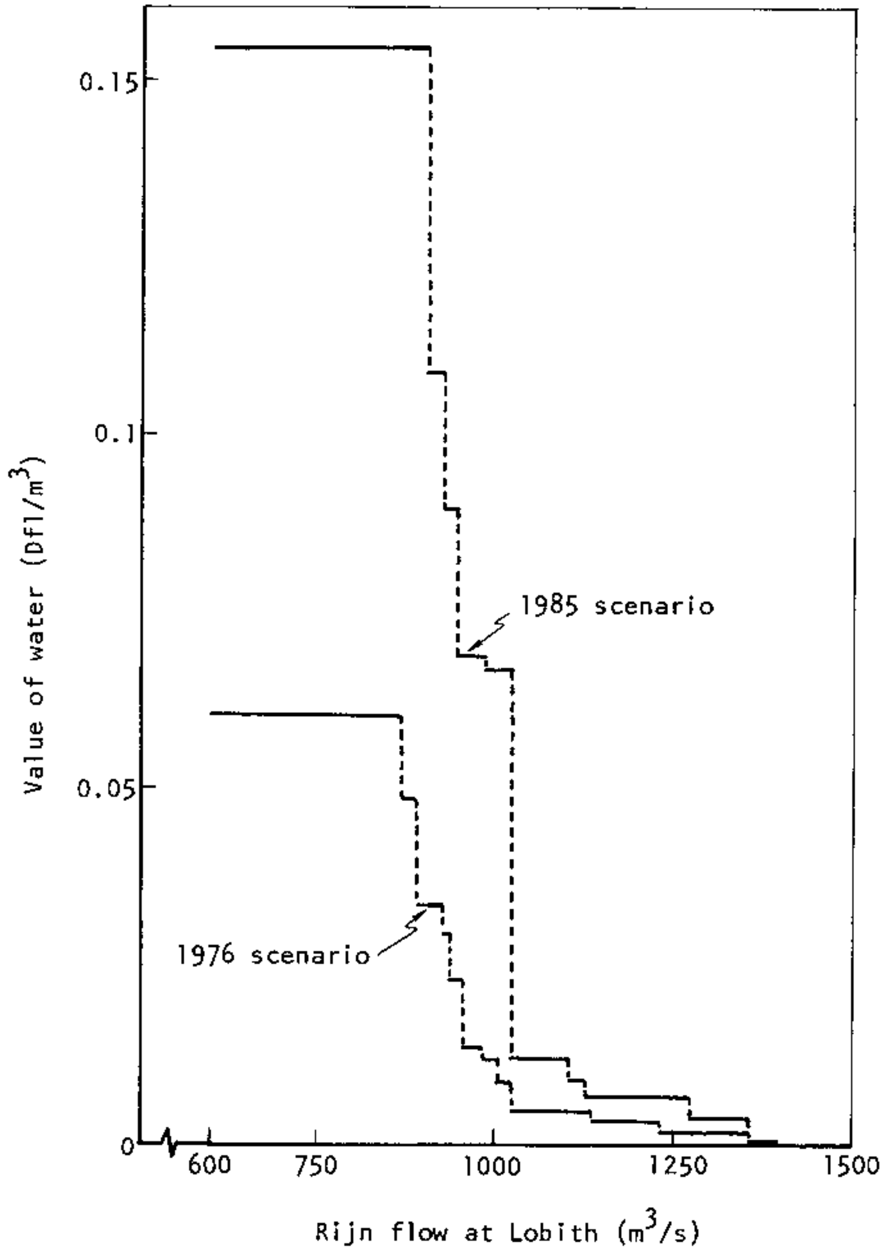

Fig. 7.5--Value to IJssel shipping of reducing withdrawals at Eefde

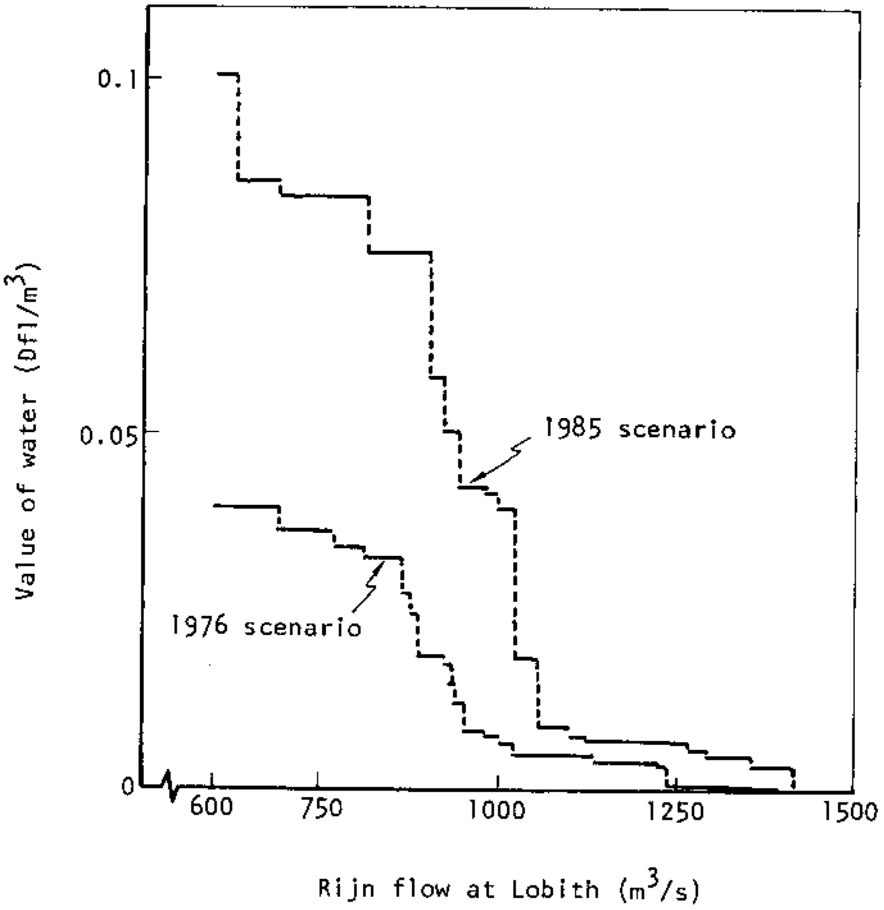

Fig. 7.6--Value to Waal and IJssel shipping of diverting water from the Neder-Rijn by closing the weir at Driel

7.7. SHIPPING DELAYS AT FRESH-FRESH LOCKS

In the Netherlands there are two types of lock complexes, the salt-fresh complex that separates salt water (usually the North Sea) from fresh water, and the fresh-fresh complex that separates fresh water at one elevation from fresh water at another elevation. In each case, the complex consists of a barrier (e.g., a dam or weir) that separates the two bodies of water, pierced by a lock that permits ships to pass from one side to the other. From a water management perspective, the purpose of a salt-fresh lock complex is to prevent the salt water on the one side from contaminating the fresh water on the other side, while permitting the passage of ships. The purpose of a fresh-fresh lock complex is to enable the desired depth to be maintained in the water bodies on both sides of the barrier without requiring an inordinately large flow of water from the higher side to the lower, and without denying passage to shipping. In this section, we are concerned only with fresh-fresh lock complexes; salt-fresh lock complexes were dealt with in Chap. 4. The fresh-fresh lock complexes in MSDM at which tactics affecting shipping are modeled are shown in Fig. 7.7. They are all located in the Southeast Highlands, on seven different links of the MSDM network. Although in reality there may be several locks on a single link, in MSDM we represent the effect of all of the locks together by a single function relating delay costs on the entire link to the water loss rate.

7.7.1. Normal Shipping Delays at Locks

A ship approaching a lock is always delayed, in comparison to the time it would take to pass that location if there were no lock. First, the ship must wait until the lock operator signals him to enter the lock chamber. If the lock is in the middle of a cycle when the ship reaches the lock, it may be twenty minutes or more before a chamber is ready to accept a new load of ships. Then the ship must creep slowly and carefully into the chamber, and tie up to the side so that the filling or emptying of the chamber will not cause the ship to swing across the chamber. Cycling the lock itself takes time. Once the opposite door of the lock is opened, the ship must cautiously leave the vicinity, moving more slowly than usual because of the natural traffic congestion near the lock.

However, shipping is willing to pay this price because the lock complexes often enable shipping to use a route that it otherwise could not. For example, the weir associated with a fresh-fresh lock permits the water depths on both sides of the lock to be maintained without requiring large flows. This enables ships to use the waterways even when flows are quite low.

7.7.2. Tactics for Conserving Water at Locks

At fresh-fresh locks, some loss of water through the lock is inevitable. A minimum background loss occurs due to simple leakage,

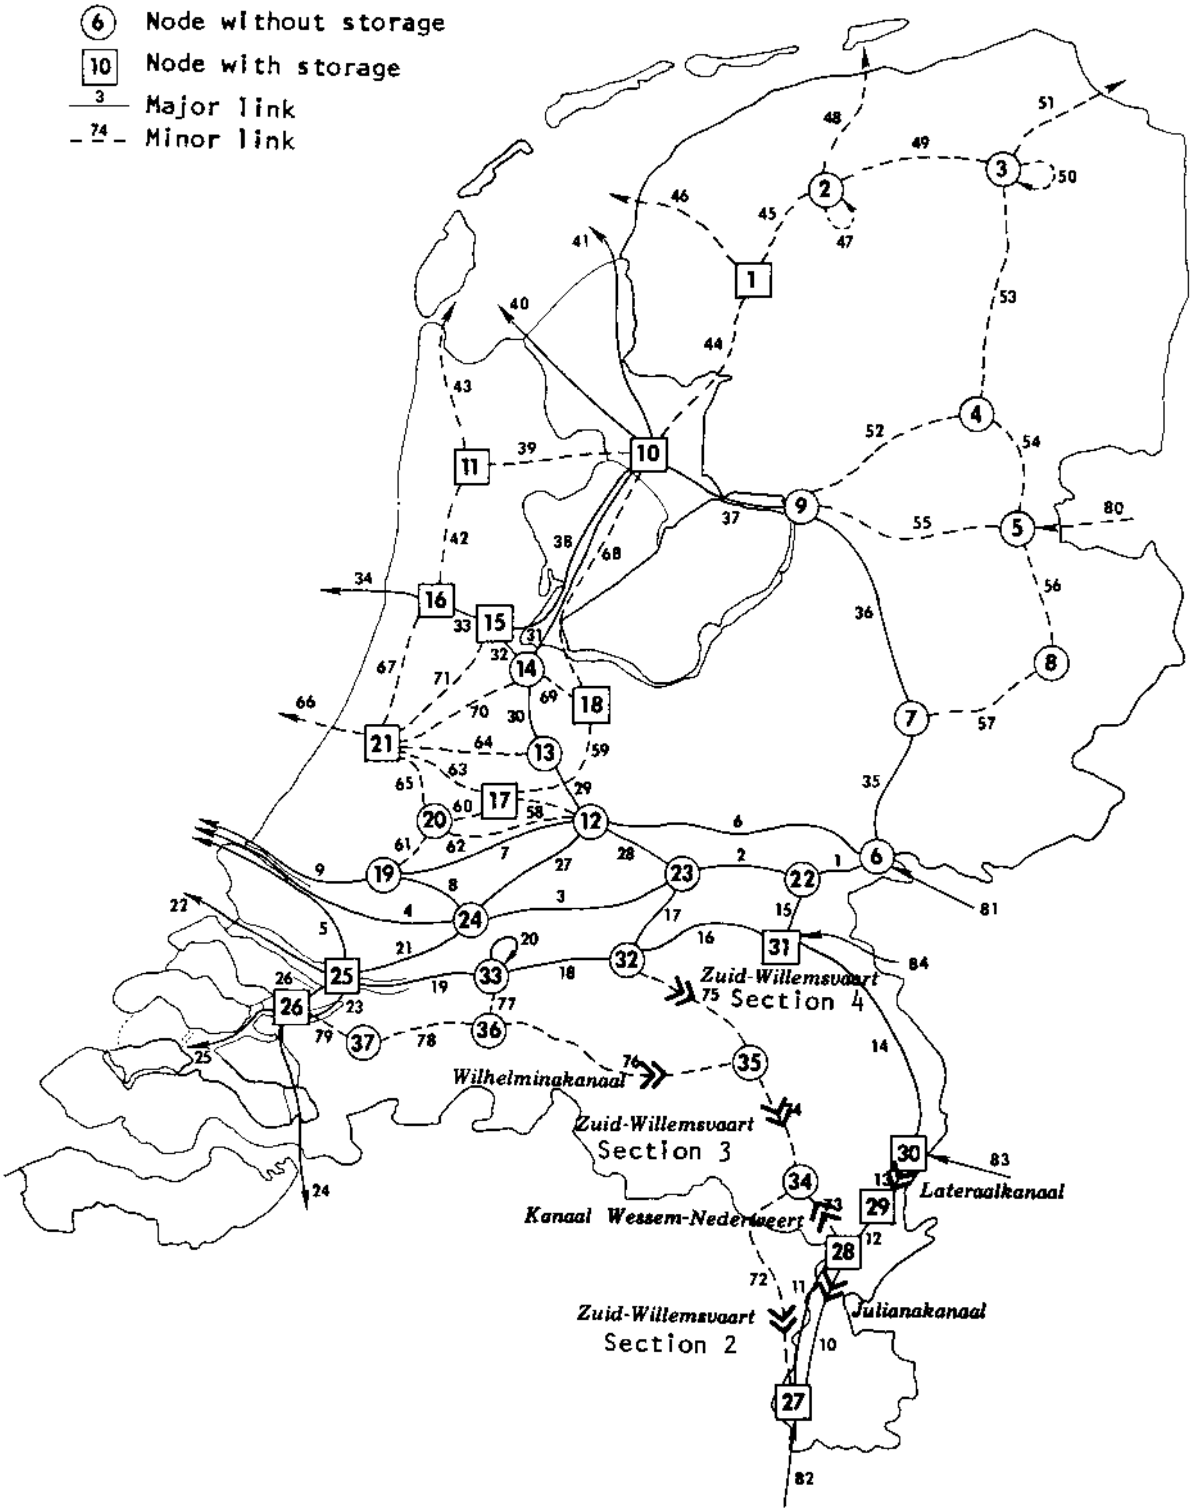

Fig. 7.7--Fresh-fresh locks considered in MSDM

even when the lock is not in use. In addition, during each cycle of the lock, a certain volume of water must be released from the higher side to the lower. In times of very low flow, it may be necessary to reduce this loss of water. A number of methods are available for doing so.

7.7.2.1. Plug Leaks. First, the leakage can be reduced. The weirs on the Maas, for example, normally leak at a rate of 15 cubic meters per second. When this much water is not available, the cracks where leakage occurs can be plugged with cotton wool, reducing the leakage to about five cubic meters per second [7.1]. The further measure of dropping fine particles (e.g., ash) in the water clogs the pores in the cotton wool and virtually eliminates the remaining leakage.

7.7.2.2. Pump Back Water. The second tactic is to pump water from the downstream side to the upstream side. At the expense of a little energy, any loss of water due to leakage or to cycling of the lock can be replaced by pumping. This solution has been used on the Maas, notably at Maasbracht during 1976. A portable pump was used to reduce by five cubic meters per second the net amount of water needed to pass ships through the Julianakanaal.

7.7.2.3. Synchronous Operation. Third, two parallel lock chambers can be connected, and operated in synchrony. When one is cycling ships upstream, the other will cycle ships downstream. The one cycling ships downstream can release its water into the other lock chamber, instead of into the water body on the downstream side, at least until half the water has been released. At this point the two lock chambers will have equal water levels, and the remainder of the cycle must take place in the normal fashion, with the downstream-cycling lock wasting the remainder of its water to the downstream side, and the upstream-cycling lock receiving the remainder of its water from the upstream side. Still, half the water needed to cycle the lock chambers has been saved.

Several locks, for example the lock on the Wessem-Nederweertkanaal at Panheel, are equipped to be operated in this fashion. However, ship operators dislike the tactic, because it causes them extra delay. Much of the delay occurs because the head difference that drives the filling of the upstream-cycling chamber is reduced by this procedure, so the filling occurs more slowly. Some of the delay occurs because the two chambers must be operated synchronously, so that if one chamber is ready to cycle, it must still await the other.

7.7.2.4. Intermediate Storage Pond. A variation on the tactic of synchronous operation is to construct a storage pond at a level midway between the water levels at the two sides of the lock. This pond can then receive water from the lock chamber when it cycles downstream, and can supply water to the chamber when it cycles upstream. As before, this saves half the water that normal operation would use, but it does not require synchronous operation. In fact, it does not even require that there be two lock chambers. One chamber alone will suffice.

7.7.2.5. Fewer Lock Cycles. The final method is to cycle the lock less often. Ships arriving at the lock are forced to wait until the chamber is full, or at least filled to some specified fraction of capacity, before they can be cycled through. At locks without much traffic, this may entail a wait of many hours. (At very busy locks, the chamber is probably full or nearly full for every cycle under normal operation, so this tactic would have no effect.) Again, ship operators prefer not to be delayed, and hence dislike this tactic.

7.7.3. Costs of Water Conservation at Locks

The actual costs of conserving water at locks are derived in Vol. IX. On four of the MSDM links, we use these costs directly. The four links are 72 ZUIDWLM2, 74 ZUIDWLM3, 75 ZUIDWLM4, and 76 WILHEKAN. At the locks on other links, certain complications arise. At links 10 JULCANL1 and 73 WESNVERT, there is the possibility of pumping water back from the downstream to the upstream side of the lock. The interaction of this tactic with lock management tactics is not considered in Vol. IX. Finally, at link 13 MAAS3, there are two parallel paths for shipping, one of which (the Lateraalkanaal) has only one lock for ships to pass, while the other (the Maas proper) has two lock in series. The interactions of the two paths are considered in Vol. IX, but we include a brief summary here.

7.7.3.1. The Simple Locks. The simple locks are the locks on the four links for which no complications arise. Figure 7.8 shows a typical example of the relation between the water loss at a lock and the accompanying cost to shipping. The link chosen for this example is 76 WILHEKAN. The figure shows the piecewise linear representation of this cost function used in MSDM. This function was determined using the lock simulation model described in Vol. IX, the same model used for salt-fresh locks (see Chap. 4).

Several things must be noted about the delay cost functions at these locks. First, the same curve applies for both the 1976 and 1985 scenarios. Second, a distinction can be made between Dutch and non-Dutch traffic at the locks. In MSDM we use the costs for the Dutch portion of the traffic alone, which is 62 percent of the total except on link 76 WILHEKAN, where it is 100 percent. Tables describing all of these cost functions can be found in App. G of Vol. VA.

7.7.3.2. Locks with Pumping. On two links, 10 JULCANL1 and 73 WESNVERT, it is possible to reduce water losses by pumping water from the downstream to the upstream side of the locks, as well as by lock management tactics. In MSDM, we wished to combine pumping with lock management in such a way as to achieve any selected water loss rate at minimum cost. Figure 7.9 illustrates how this is done for the lock at Panheel, on the Wessem-Nederweertkanaal (link 74 WESNVERT).

In the figure, the dashed curve on the right shows the cost of achieving each water loss rate using only lock management tactics.

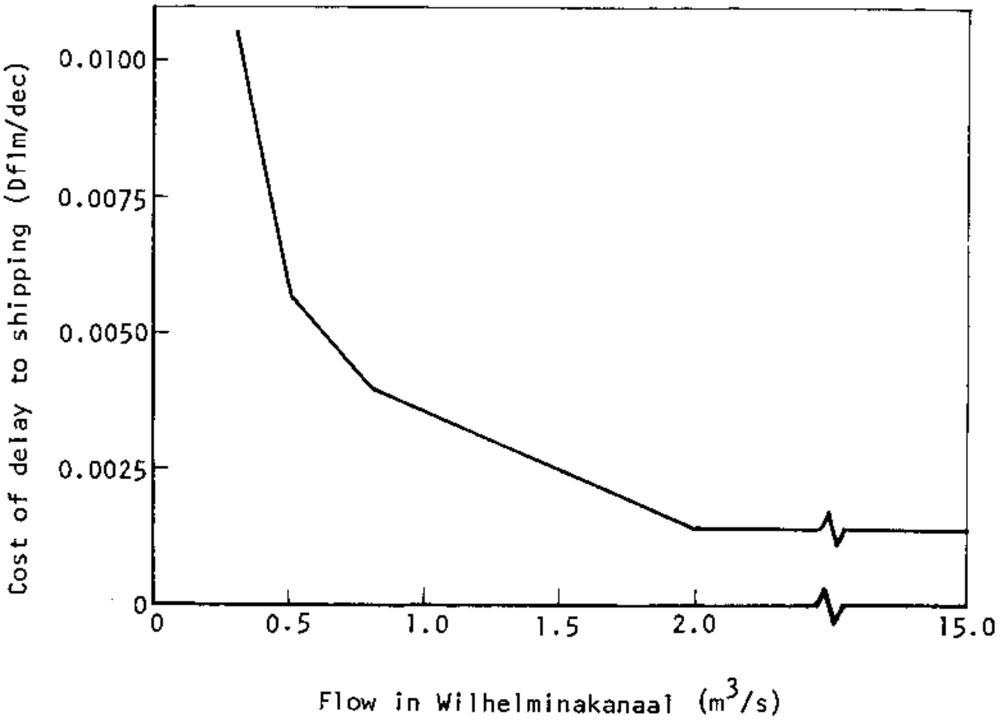

Fig. 7.8--Delay cost to shipping of conserving water at locks on the Wilhelminakanaal (link 76 Wilhekan)

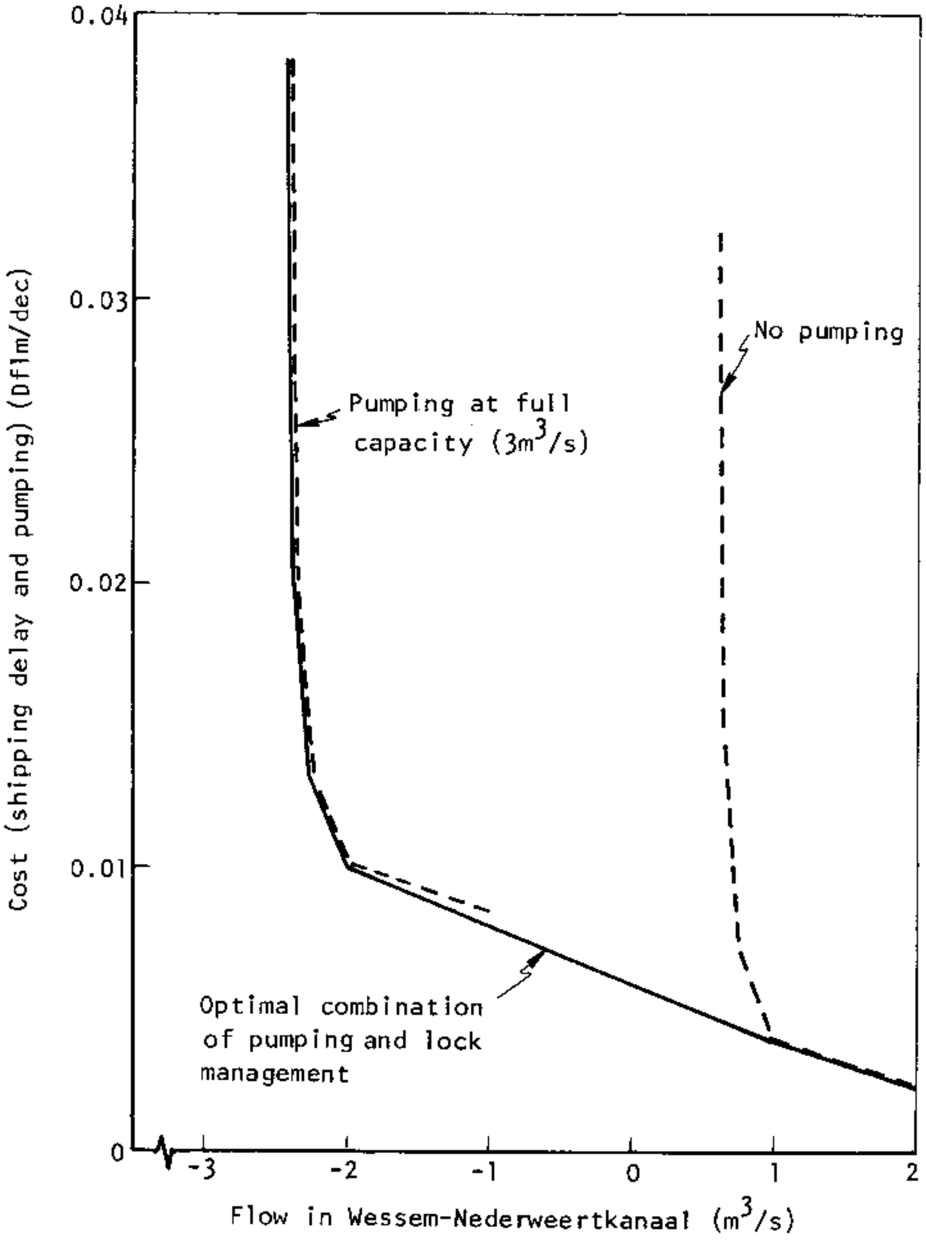

Fig. 7.9--Optimal combination of pumping and lock management at Panheel lock on Wessem-Nederweertkanaal (link 73 Wesnvert)

This curve is obtained in exactly the same way as the curve in Fig. 7.8. The dashed curve on the left shows the cost of achieving each water loss rate if the pumps at Panheel are operated at full capacity, and lock management tactics are used in addition. This curve can be obtained from the right-hand dashed curve by translating it 3 m^3/s to the left (the pumping capacity), and 0.00615 Dflm/dec upward (the energy cost of operating the pumps at capacity for one decade). Of course, it is not necessary to operate the pumps at full capacity. Any intermediate pumping rate is possible. By varying the pumping rate we can achieve any point intermediate between the two curves. The least expensive way of achieving each water loss rate is given by the lower envelope of all the intermediate curves, which is shown by the solid curve in the figure.

The situation on the Julianakanaal (link 10 JULCANL1) is similar to that on the Wessem-Nederweertkanaal, except that there are two locks in series, at Born and Maasbracht, each of which possesses some capacity to pump water upstream. The pumps at Born have a capacity of 13 m^3/s, and an energy cost when operated at capacity of 0.0368 Dflm/dec. At Maasbracht, the present pumping capacity is 5 m^3/s, and the energy cost at capacity is 0.0151 Dfl/dec. The loss of water from each lock must be the same, or the canal section between them would either overflow or run dry. But one may choose the least expensive combination of pumping and shipping delays at each lock in order to achieve the desired rate of water loss. To derive the cost as a function of the water loss rate on this link, we derived curves for the two locks separately in the same fashion as described for the Wessem-Nederweertkanaal, and then added them to obtain the total cost function for the link.

A tactic is described in Vol. XVI that would increase the pumping capacity at Maasbracht. This would change the cost function for link 10 JULCANL1. The method for obtaining the function, however, would remain the same.

Tables showing the cost functions used in MSDM for these links can be found in App. G of Vol. VA. For link 10 JULCANL1, there are separate functions for the 1976 and 1985 scenarios. For link 73 WESNVERT, a single curve is used for both scenarios. In both cases, the delay costs are the costs to Dutch shipping only, which is taken to be 62 percent of the total shipping (Vol. IX).

7.7.3.3. The Lateraalkanaal. At the Lateraalkanaal (link 13 MAAS3) there are two parallel paths for shipping. Ships using the Lateraalkanaal must pass only one lock, while ships using the Maas must pass two locks in series. Therefore, shipping prefers to use the Lateraalkanaal unless its destination is along this section of the Maas. However, with two locks in series the water lost from the upstream lock can be used to cycle the downstream lock, and so water losses are less if shipping uses the Maas route instead of the Lateraalkanaal route. The problem here is to apportion both the available water and the shipping traffic between the two routes to achieve the desired water loss rate at the minimum delay cost to shipping.

The procedure for doing is described fully in Vol. IX, but briefly it was the following. The lock simulation model was applied to each route individually, giving the delay cost to shipping on that route as a function of both the traffic and the water loss on the route. For any traffic level, the result would be a curve such as that shown in Fig. 7.8; since we varied the traffic level, we obtained a family of such curves for each route. In symbols, we can write these functions as:

$$CL = DL(TL,WL) \qquad CM = DM(TM,WM)$$

where CL, CM = delay costs on the Lateraalkanaal and the Maas routes, respectively;

TL, TM = traffic levels on the Lateraalkanaal and the Maas routes, respectively;

WL, WM = water losses on the Lateraalkanaal and the Maas routes, respectively;

DL, DM = functions relating water losses and traffic levels to delay costs, on the Lateraalkanaal and the Maas routes, respectively.

Then the problem of apportioning the water and traffic between the two routes may be expressed as a standard mathematical programming problem, in which TL, TM, WL, and WM are to be chosen so that TL + TM equals the known total traffic level, WL + WM equals the known total water flow, and CL + CM is a minimum. That is,

$$\text{Minimize: } (CL + CM) = DL(TL,WL) + DM(TM,WM)$$

Subject to:

$$\begin{aligned} TL + TM &= (\text{Total traffic}) \\ WL + WM &= (\text{Total water flow}) \\ TL &\geq (\text{Minimum Lateraalkanaal traffic}) \\ TM &\geq (\text{Minimum Maas traffic}) \\ WL &\geq (\text{Minimum Lateraalkanaal flow}) \\ WM &\geq (\text{Minimum Maas flow}) \end{aligned}$$

The minimum allowable traffic on the Lateraalkanaal is zero, but the minimum allowable traffic on the Maas is greater than zero, since some ships have their destinations along that section of the Maas. The minimum flows on the two routes are the smallest flows achievable by available lock management tactics and traffic reductions.

This mathematical programming problem is solved in Vol. IX. The cost function it yields looks much the same as the curve in Fig. 7.8. We present a table of values describing this function in App. G of Vol. VA. The same function is used for both the 1976 and 1985 scenarios, and it includes only the delay cost to Dutch shipping, which is taken to be 62 percent of the total.

7.7.4. The Value of Water for Reducing Conservation Costs

In one sense, it is straightforward to calculate the value of water used to reduce the cost of conserving water at locks. One simply calculates the derivatives of the cost functions described above, and scales them to the units of Dfl/m^3. The result is the savings that would accrue if an additional cubic meter of water were made available to operate each of the locks.

However, increasing the water available at a lock is not by itself a valid managerial tactic. The water must be brought from somewhere, and once it has passed through the lock, it must go somewhere. For simplicity we could assume that there were extractions just upstream of each lock, and that the source of our extra water were reductions in those extractions. But the problem of disposing of the water is not so simple.

At the locks on links 13 MAAS3, 75 ZUIDWLM4, and 76 WILHEKAN, the water can be disposed of simply by letting it flow downstream. At the other four links, however, water that flows downstream will pass through additional locks. For example, if extra water is provided to the Julianakanaal (link 10 JULCANL1), it will flow downstream to link 13 MAAS3, reducing conservation costs there. Water that passes through the locks on link 72 ZUIDWLM2 may flow downstream along links 73 WESNVERT and 13 MAAS3, or along links 74 ZUIDWLM3 and 75 ZUIDWLM4, or links 74 ZUIDWLM3 and 76 WILHEKAN. Moreover, it is not clear what extractions and discharges may occur between one lock and others downstream of it, so we cannot know in general how much water is available to operate the downstream locks, as a function of the amount of water available to operate the upstream lock.

In calculating the value of water for reducing conservation costs, we will therefore look at each lock in isolation, recognizing that the values we will calculate are not the values of water used according to any valid managerial tactic. The results are shown for all seven links in Fig. 7.10. To calculate the effect a managerial tactic will have on water conservation costs at locks, one must first define the initial amounts of water available at each of the locks, and specify how each of those flows will be changed by the tactic considered. The effect at each lock can be obtained from the figures, and the total effect can be calculated by summing the effects at the individual locks.

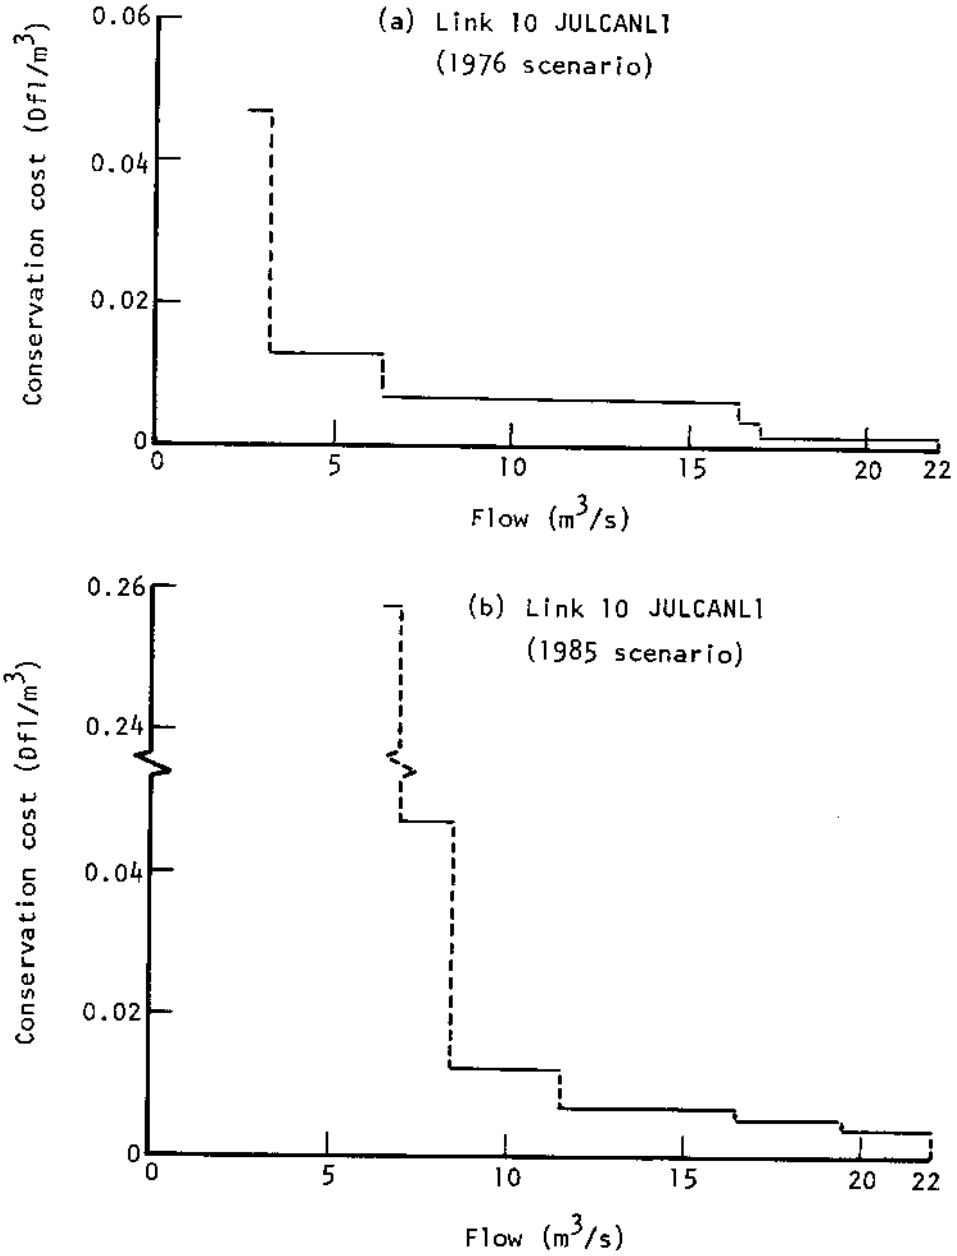

Fig. 7.10a--Value of water for reducing conservation costs at fresh-fresh locks: link 10 JULCANL1 (1976 and 1985 scenarios)

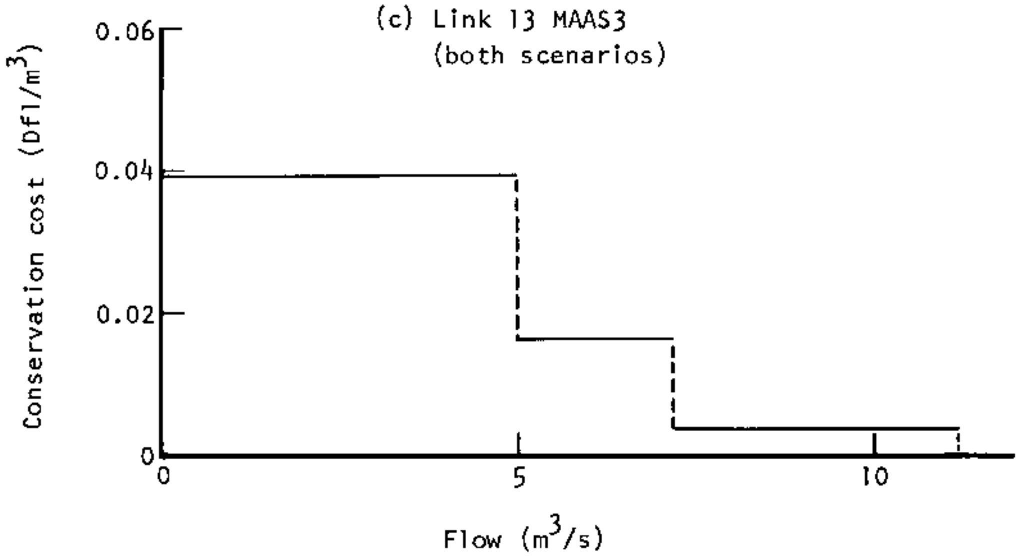

(c) Link 13 MAAS3
(both scenarios)

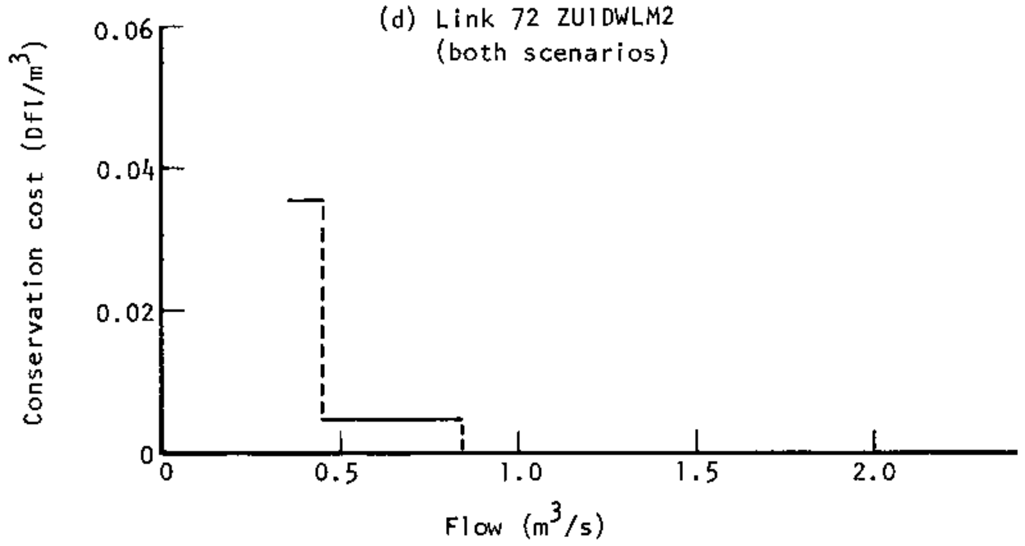

(d) Link 72 ZUIDWLM2
(both scenarios)

Fig. 7.10b--Value of water for reducing conservation costs at fresh-fresh locks: link 13 MAAS3 and link 72 ZUIDWLM2 (1976 and 1985 scenarios)

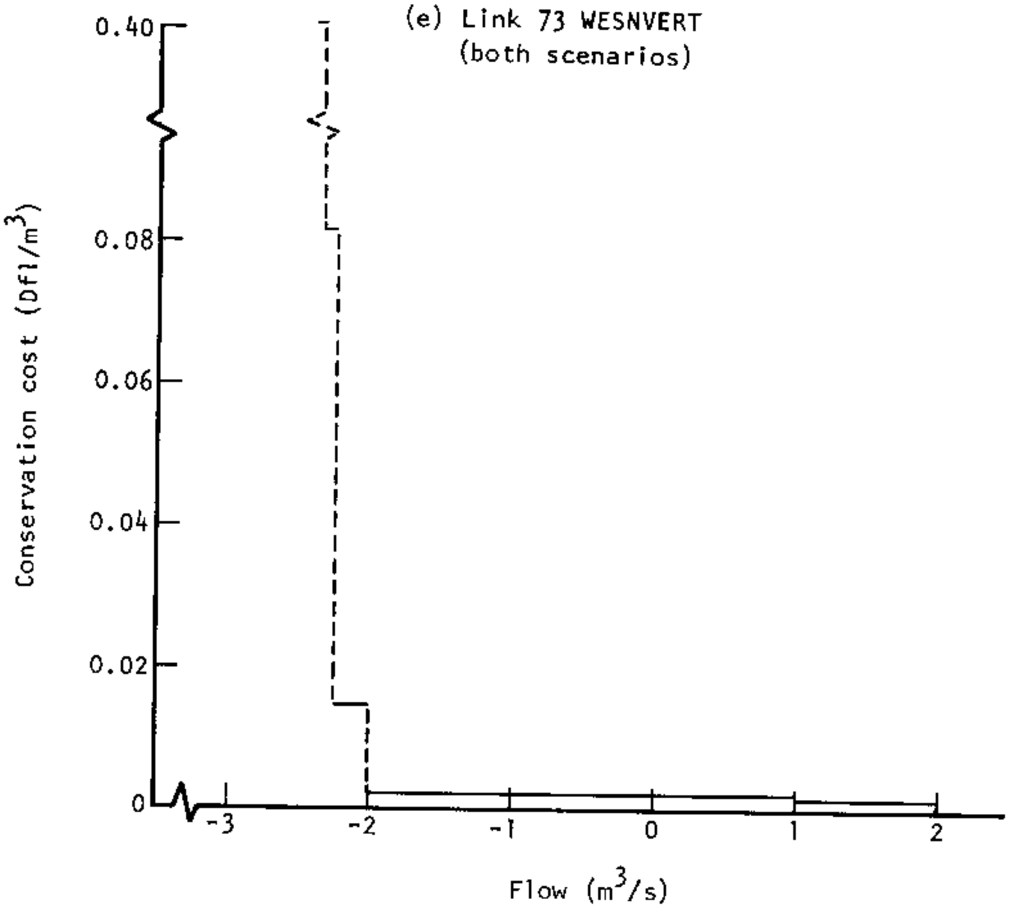

Fig. 7.10c--Value of water for reducing conservation costs at fresh-fresh locks: link 73 WESNVERT (1976 and 1985 scenarios)

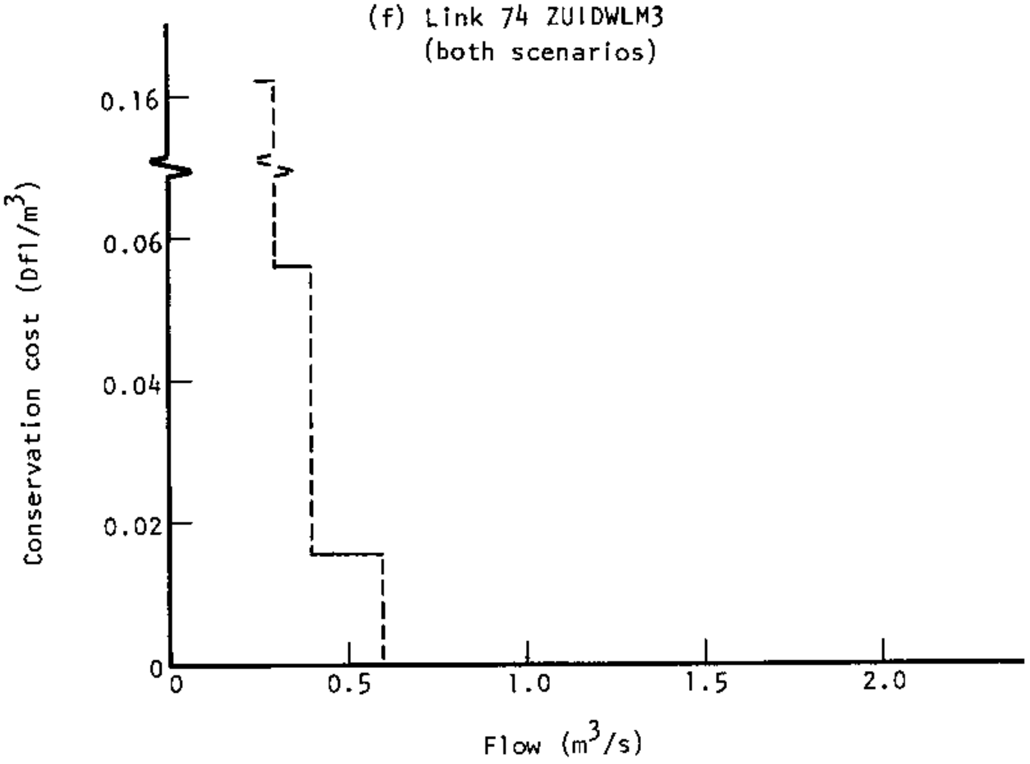

Fig. 7.10d--Value of water for reducing conservation costs at fresh-fresh locks: link 74 ZUIDWLM3 (1976 and 1985 scenarios)

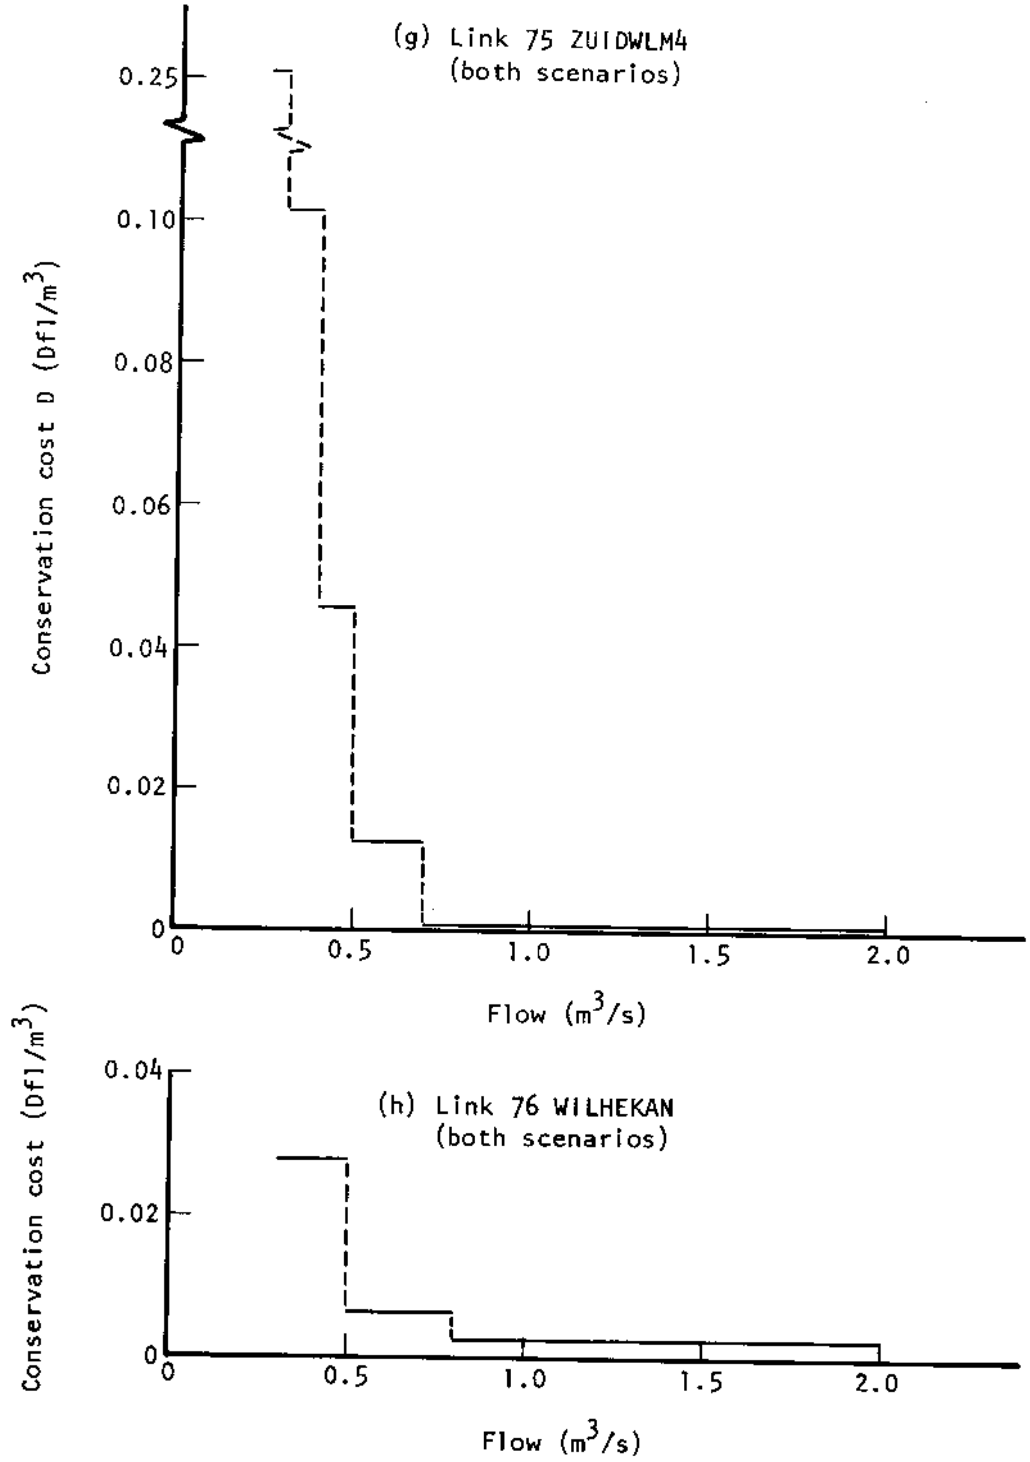

Fig. 7.10e--Value of water for reducing conservation costs at fresh-fresh locks: link 75 ZUIDWLM4 and link 76 WILHEKAN (1976 and 1985 scenarios)

As one can see from the figure, the marginal savings realized from increasing the flow at some of these locks can be considerable. On link 73 WESNVERT, for example, increasing the flow of water can save more than 40 Dutch cents for each additional cubic meter of water. But this is only true when one is attempting to conserve the maximum possible amount of water at the lock. As the flow is increased, the value of increasing the flow still further drops rapidly. Thus, the total amounts of both money and water at stake are small.

NOTES

1. Ideally, these functions should include the cost of storing cargo for as long as it will be delayed, i.e., for as long as river flows remain sufficiently low. However, in Vol. IX, where these functions are derived, only one decade's worth of storage costs were included. In our analysis, we simply used the storage cost functions from Vol. IX, without adding costs for extra decades of storage.

   The omission probably has at most a minor effect on our detailed results, and no effect on our general conclusions. This is because the water used to benefit shipping has a value to competing uses that is either already less than the value to shipping, or is enough greater that inflating the storage functions by a reasonable factor (e.g., less than ten, which assumes that on the avarage there will be no more than ten consecutive decades with low river flows) will not cause shipping to become a higher-valued use. Unless there occurs a switch in the relative value of water to two competing users, there will not be any change in the allocation of water due to an inflation of the storage functions.
2. Strictly speaking, the volume of the sediment deposit will also depend on the flow in the Waal, but as shown in Vol. XVIII, this dependence is very weak, and we ignore it in MSDM.

REFERENCE

7.1. MW-179 (unpublished PAWN memorandum), "Cost and Discharge Reduction Procedures Related to the Cottonwool Method at the Maas Weirs," June 1976.

Chapter 8

METHOD OF SOLUTION

8.1. INTRODUCTION AND SUMMARY

In this chapter, we briefly discuss the method used to solve the water managerial problem formulated in the earlier part of this volume. Unavoidably, much of the discussion will be technical. For the reader not interested in such details, only the initial, introductory section need be read.

In the previous chapters, the water managerial problem has been presented as six separate subproblems. They are:

- The Water Distribution Problem, Chap. 2.
- The BOD Problem, Chap. 3.
- The Phosphate Problem, Chap. 3.
- The Heavy Metal Problem, Chap. 3.
- The Salt Problem, Chap. 4.
- The Thermal Pollution Problem, Chap. 5.

Each subproblem has its own constraints, variables, and objective function. The connection between them is that the constraints of the five water quality subproblems (BOD, Phosphate, Heavy Metal, Salt, and Thermal Pollution) depend on the flows and water levels in the water management infrastructure, and these flows and water levels are variables of the water distribution problem.

It can be shown that the water managerial problem is not necessarily convex. What it means for the problem to be nonconvex is unimportant here; it is described in Sec. 8.4 below. What it implies is that there may be several <u>local</u> optimal solutions to the problem. Each local optimum will consist, in part, of flows and water levels in the links and nodes of the network. Starting from a local optimum, modest changes in the flows and water levels will result in an increase in the total costs incurred in the water distribution problem plus the five water quality subproblems. (That is, the local optimum is the best solution in its own locale.) But without examining all the local optima, it is impossible to determine which is the best overall--i.e., which is the <u>global</u> optimum.

We have devised an algorithm that seeks a local optimal solution to the water managerial problem. The algorithm cycles through the six subproblems, first solving the water distribution problem, and then each of the water quality subproblems in turn. Whenever a water quality subproblem is to be solved, the flows and water levels from the most recent water distribution solution are used to define the water quality constraints. Whenever the water distribution problem is to be solved, the most recent solution to each water quality subproblem is used to estimate the effect of changes in flows and water levels on the corresponding water quality-related costs. This effect is introduced as a new cost term in the water distribution objective function. Thus a solution to the water distribution problem is used to define the constraints of each water quality subproblem; the solutions of the water quality subproblems are used to modify the water distribution objective function; this results in new flows and water levels, which can be used to redefine the constraints of the water quality subproblems; and so on.

This algorithm is not guaranteed to converge. In many instances, it will oscillate from one water distribution to another, and then back to the first, neither solution being a local optimum. (If the algorithm happens to strike a local optimum, it will simply repeat that solution ad infinitum.) To counter the oscillations, we manually alter the upper and/or lower bounds on the flows in certain key network links. This restricts the range over which the water distribution can vary in successive solutions to the water distribution problem.

This manual procedure has two advantages. First, if the flows in enough links (usually less than four, if properly chosen) are restricted, the algorithm will converge to a local optimal solution to the original problem. Second, by making several runs which differ only in the ranges to which the flows on key links are constrained, it is possible to generate several local optima. For any set of circumstances, as defined by river flows into the country, rainfall, and the amount of water in storage at the start of a time period, this procedure allows us to generate at least one, and sometimes several attractive alternative managerial strategies.

8.2. MATHEMATICAL STATEMENT OF THE PROBLEM

Let us consider a simpler version of the water managerial problem, having only one water quality problem in addition to the water distribution problem. We can write this problem as follows:

$$(8.1)\quad \begin{cases} \text{Minimize:} & g(x) + h(y) \\ \text{Subject to:} & Ax = b \\ & F(x)*y = d \\ & l0 \le x \le u0 \\ & l1 \le y \le u1 \end{cases}$$

where x = a vector of flows and water levels in the network, plus any other variables in the water distribution problem;

y = a vector of pollutant concentrations at each network node, plus any other variables in the water quality subproblem;

g(x) = the objective function of the water distribution problem;

h(y) = the objective function of the water quality subproblem;

A = the coefficient matrix for the constraints in the water distribution problem;

F(x) = the coefficient matrix for the constraints in the water quality subproblem. As derived in App. B of Vol. VA, this matrix depends on the water levels and flows in the network;

b = the constant terms in the water distribution constraints;

d = the constant terms in the water quality constraints;

l0 = a vector of lower bounds on the variables x;

u0 = a vector of upper bounds on the variables x;

l1 = a vector of lower bounds on the variables y;

u1 = a vector of upper bounds on the variables y.

It is useful to introduce the following terminology. By a feasible solution we mean a pair of vectors (x, y) that satisfy all the constraints of Problem (8.1). The solution is also optimal if no other feasible solution yield a smaller value of the objective function.

The entire water managerial problem is nothing more than an extended version of Problem (8.1). It would include vectors of variables y1, y2, ..., y5 for all five water quality subproblems, in place of the single vector y. There would be five terms h1(y1), h2(y2), ..., h5(y5) in the objective function in place of the single term h(y). Five sets of constraints F1(x)*y1 = d1, F2(x)*y2 = d2, ..., F5(x)*y5 = d5, would appear in place of the one set of constraints F(x)*y = d. And there would be lower and upper bounds for all the vectors y1, y2, ..., y5, instead of the single vector y.

Note that there is only one interaction in Problem (8.1) between terms that relate to the water distribution problem, and terms that relate to the water quality subproblem. That interaction is the dependence of the coefficient matrix of the water quality subproblem, F(x), on the variables of the water distribution problem, x. If this dependence did not exist, it would be unnecessary to consider both problems simultaneously. Each could be solved individually, and the two answers would together constitute the optimal solution to the overall water managerial problem.

As it is, however, the optimal solution to each problem depends on the other. The solution to the water quality subproblem surely depends on the water distribution problem, through the dependence of its contraints on the flows and water levels, x. Thus, a change in the flows and water levels may well change the optimal solution to the water quality subproblem, and hence influence the value of the term h(y) in the objective function. While at first sight the water distribution problem is independent of the water quality subproblem, this is misleading. The effect of flows and water levels on the water quality objective function h(y) must be considered in choosing the optimal flows and water levels in the network.

8.3. OPTIMALITY CONDITIONS

Problem (8.1) is expresed as a mathematical programming problem. For such a problem, satisfying rather mild conditions, the classic paper of Kuhn and Tucker [8.1] gives conditions which the optimal solution must satisfy. There is one optimality condition for each of the variables in the two vectors x and y. For a typical x variable, the condition is:

$$(8.2) \qquad \frac{dg(x)}{dx_j} + w^T * A_j + q^T * \frac{dF(x)}{dx_j} * y \quad \begin{cases} \geq 0, \text{ if } & l0_j = x_j \\ = 0, \text{ if } & l0_j < x_j < u0_j \\ \leq 0, \text{ if } & x_j = u0_j \end{cases}$$

In condition (8.2), the subscript j refers to one of the components in the vector x. For j=1, we mean the first component; for j=2, the second; and so on. Thus:

$\frac{dg(x)}{dx_j}$ = the partial derivative of the function g(x) with respect to component j of the vector x.

A_j = column number j of the coefficient matrix of the water distribution problem.

$\frac{dF(x)}{dx_j}$ = the partial derivative of the coefficient matrix F(x) of the water quality subproblem with respect to component j of the vector x. This derivative is taken element by element, and hence the result is a matrix with the same number of rows and columns as F(x).

In condition (8.2), two new quantities have been introduced. They are:

w = the "dual variables," or Lagrange multipliers associated with the constraints of the water distribution problem. This is a vector having one component for each such constraint.

q = the "dual variables" associated with the water quality subproblem. This is a vector having one component for each such constraint.

The superscript "T" on the quantities w and q in condition (8.2) means "Transpose." All vectors are nominally taken to be column vectors. To make the matrix multiplications in condition (8.2) mathematically correct, these vectors must first be made into row vectors, which is accomplished by transposing them.

There are conditions analogous to (8.2) for the variables y in the water quality subproblem. These conditions are:

$$\frac{dh(y)}{dy_k} + q^T * F_k(x) \quad \begin{cases} \geq 0, \text{ if } & ll_k = y_k \\ = 0, \text{ if } & ll_k < y_k < ul_k \\ \leq 0, \text{ if } & y_k = ul_k \end{cases} \tag{8.3}$$

In condition (8.3), the subscript k refers to one of the components in the vector y. For k=1, we mean the first component; for k=2, the second; and so on. Thus:

$\frac{dh(y)}{dy_k}$ = the partial derivative of the function h(y) with respect to component k of the vector y.

$F_k(x)$ = column number k of the coefficient matrix of the water quality subproblem. Of course, this depends on the flows and water levels x.

The Kuhn-Tucker Theorem states that, if Problem (8.1) satisfies certain rather mild conditions, then a feasible solution (X, Y) to Problem (8.1) is also optimal only if there exists a dual pair of vectors (W, Q) such that when (X, Y) and (W, Q) are substituted into condition (8.2) and (8.3), all the conditions are satisfied. Here we have used capital letters to denote the particular values of the variables in the optimal solution, and small letters to denote the variables when particular values are not intended.

8.4. LACK OF CONVEXITY

Note that, while the Kuhn-Tucker Theorem states that the optimal solution to Problem (8.1) will satisfy conditions (8.2) and (8.3), it does not guarantee that the optimal solution (or solutions--there may be more than one) will be the only solution (or solutions) to satisfy these conditions. This would only be guaranteed if Problem (8.1) were convex. By this we mean that given any two feasible solutions (x1, y1) and (x2, y2), neither of which is necessarily optimal, and for any number "a" between zero and one, the intermediate point (a*x1+(1-a)*x2, a*y1+(1-a)*y2) is also feasible. In addition, it should be true that:

$$g(a*x1 + (1-a)*x2) \le a*g(x1) + (1-a)*g(x2) \qquad (8.4a)$$

$$h(a*y1 + (1-a)*y2) \le a*h(y1) + (1-a)*h(y2) \qquad (8.4b)$$

In Problem (8.1), conditions (8.4a,b) are true, but the condition that points intermediate between feasible solutions must themselves be feasible is false. To illustrate this, consider the simple network depicted in Fig. 8.1. A flow of 1 m^3/s enters node 1 (at the left), and may be divided in any proportions between the two links. One kg/s of pollutant is discharged into each of the nodes 2 and 3. The only water available to dilute the pollutant at either node 2 or 3 is the water sent from node 1. Water exiting nodes 2 and 3 is discharged out of the network (e.g., into the North Sea).

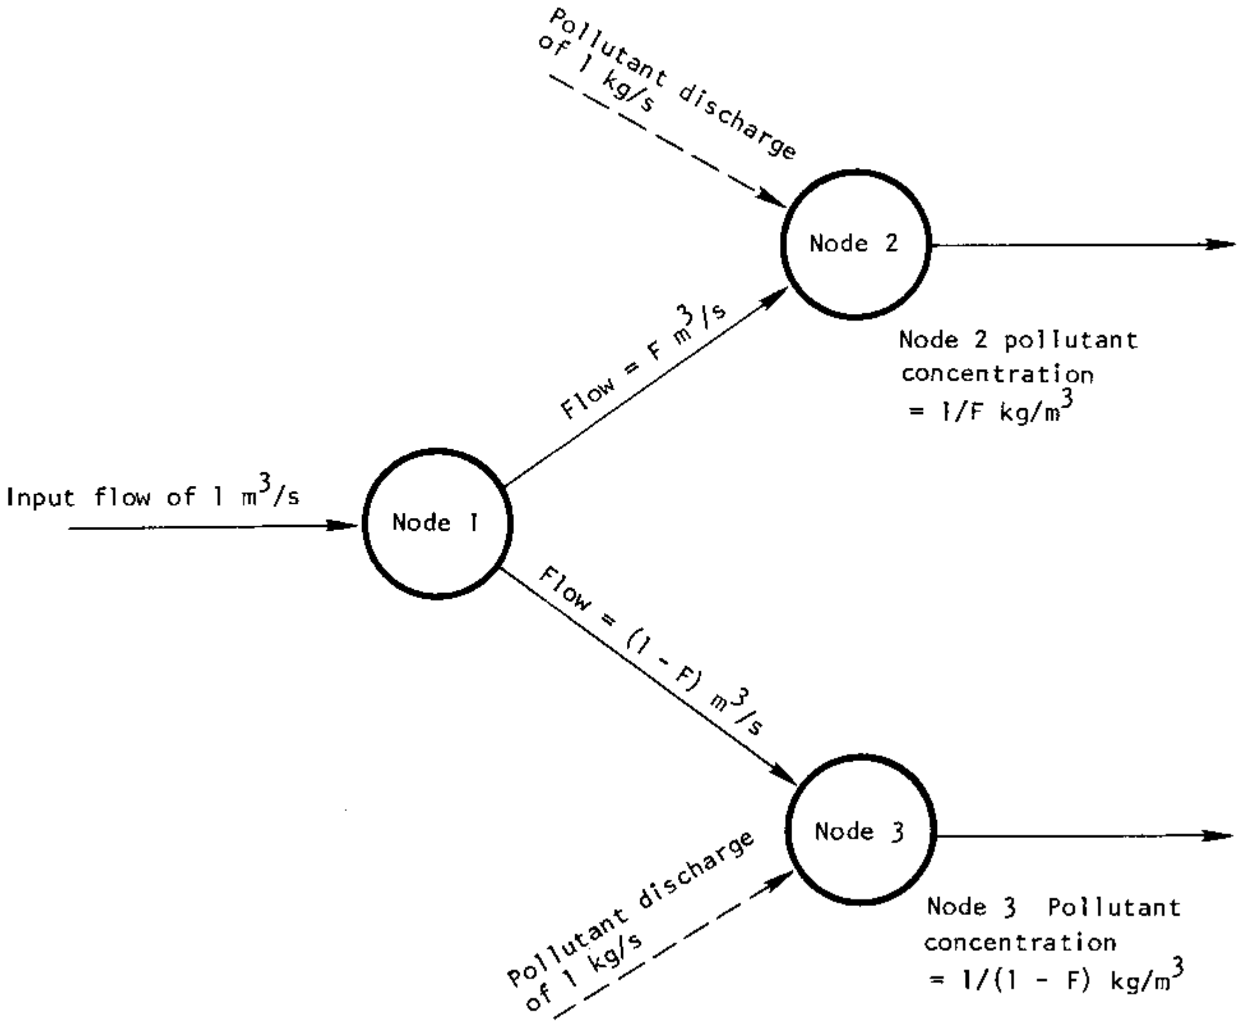

Fig. 8.1--Network illustrating nonconvexity

The vectors x and y have only two components of interest each. For the vector x, the interesting components are the flows from node 1 to 2, and from node 1 to 3. If we let "F" be the flow from node 1 to 2, then "1-F" must be the flow from node 1 to 3. Thus, all feasible vectors x may be written:

$$x = (F,\ 1-F)$$

For the vector y, the interesting components are the pollutant concentrations at nodes 2 and 3. Assuming the flow from node 1 to 2 is "F," the concentration at node 2 must be 1/F kg/m^3; since the corresponding flow from node 1 to 3 is "1-F," the corresponding concentration at node 3 must be 1/(1-F) kg/m^3. Thus, all feasible vectors y may be written:

$$y = (1/F,\ 1/(1-F))$$

Now consider two feasible solutions, one with F=0.1, and the other with F=0.9. Then:

$$x1 = (0.1,\ 0.9), \quad y1 = (10.0\ ,\ 1.11)$$

$$x2 = (0.9,\ 0.1), \quad y2 = (\ 1.11,\ 10.0)$$

If we take a point intermediate between these two feasible solutions, let us say the midpoint x = (0.5, 0.5) and y = (5.56, 5.56), we find that while x is a feasible solution to the water distribution problem, y is not a feasible solution to the water quality subproblem, and certainly not the feasible solution corresponding to the vector x = (0.5, 0.5).

In spite of the lack of convexity of Problem (8.1), the method we use to solve our problem seeks a feasible solution that satisfies the optimality conditions (8.2) and (8.3). We know that such a solution will be locally optimal, so that small variations from such a solution cannot yield a lower value of the objective function. Furthermore, we do not rely solely on the method. We solve the problem several times, changing the constraints each time, particularly the upper and lower bounds on the flows and water levels. By this means we generate a variety of good solutions to Problem (8.1), rather than a single solution. We have some hope, therefore, that we may discover a truly optimal solution. And even if the truly optimal solution should escape our scrutiny, we will have generated at least one reasonable good solution.

8.5. THE SOLUTION METHOD

Our solution method is based on the fact that the effect of any constraint in Problem (8.1) can equally well be represented by a term in the objective function. Thus, suppose we happen to know the values of (X, Y) and (W, Q) of the optimal solution. (As before, we use capital letters to denote the particular values, and small letters to denote the variables when no particular value is intended.) Then, we would find X and W to be the optimal solution of the following problem:

$$(8.5) \quad \begin{cases} \text{Minmize:} & g(x) + \sum_{j=1}^{n} \left[x_j * Q^T * \dfrac{dF(X)}{dx_j} * Y \right] \\ \text{Subject to:} & Ax = b \\ & l0 \leq x \leq u0 \end{cases}$$

In fact, the optimality conditions for Problem (8.5) are exactly the same as conditions (8.2). Similarly, we would find Y and Q to be the optimal solution of:

$$(8.6) \quad \begin{cases} \text{Minmize:} & h(y) \\ \text{Subject to:} & F(X)*y = d \\ & l1 \leq y \leq u1 \end{cases}$$

It can be shown that the optimality conditions for Problem (8.6) are the same as conditions (8.3).

Of course, to construct problems (8.5) and (8.6) requires that we know the optimal solution (X, Y) and (W, Q), and if we knew these quantities there would be no need to go any further. However, we can construct problems that are approximations to (8.5) and (8.6) instead, and hope to improve the approximations by an iterative procedure.

If we adopt this approach, the value of introducing the two problems (8.5) and (8.6) becomes evident. Problem (8.5) is the water distribution problem with an extra term in its objective function which accounts for the effect of the water quality subproblem on the optimal water distribution, while Problem (8.6) is the water quality subproblem, in which the water distribution has been specified. Each of the problems may now be solved independently of the other. The only question is, what values of X, Y and Q should be used in successive approximations to Problem (8.5), and what values of X should be used in successive approximations of Problem (8.6)?

Our answer is the following. Initially, we assume that Q=0, so the water quality subproblem has no effect on the water distribution problem. We then solve Problem (8.5), which yields a first estimate of X. Call this first estimate X1. Now, substitute X1 for X in Problem (8.6), and obtain a solution Y1 and Q1, which are first estimates of Y and Q. Use these estimates in Problem (8.5), and solve for a second estimate X2 of X. In general, one uses the latest estimate of X in Problem (8.6), and solves to obtain new estimates of Y and Q. These are used in Problem (8.5) which is then solved to yield the next estimate of X.

In the simplified example problem (8.1), we considered only one water quality subproblem, whereas in MSDM we consider five. But it is straightforward to extend the iterative method outlined above to a problem with multiple water quality subproblems. Instead of having a single problem (8.6), MSDM has five of them. Each gives rise to its own new terms in the objective function of the water distribution problem (8.5). In the iterative scheme, all five water quality subproblems are solved, each individually, between successive solutions of the water distribution problem.

Ideally, the process will converge upon a solution to the original problem (8.1) that satisfies the optimality conditions (8.2) and (8.3). In practice, convergence is not guaranteed. We will discuss what we may do to obtain convergence in Sec. 8.7 below. Before we do so, however, we will outline the method used to solve the individual Problems (8.5) and (8.6).

8.6. SOLVING THE SUBPROBLEMS

Both Problems (8.5) and (8.6) can be expressed as linear programming problems. Both have linear constraints. In addition, as has been discussed in previous chapters, all the terms in the objective functions of both the water distribution problem and all the water quality subproblems are piecewise linear functions of a single variable. That is, we have nowhere used a function of two or more variables, nor a function we are unwilling to express in terms of several linear segments. Finally, all the terms in the objective functions are convex--i.e., they satisfy conditions (8.4a,b). Even the new terms in the objective function of Problem (8.5), that account for the effect of the water quality subproblems on the optimal water distribution, satisfy these conditions. In fact, the new terms are linear.

Given a problem that satisfies these conditions, the method described in Ref. 8.2, pp. 175-180, can be used to transform it into a linear programming problem. The standard method for solving such a problem is called the simplex method [8.3], of which there are many variants. The variant we have used was tailored especially for this problem. It takes advantage of the fact that a dual-feasible solution is immediately available, since every variable is bounded from both above and below. It treats the upper and lower bounds on the variables implicitly rather than explicitly. The computational scheme also takes advantage of the sparcity of the coefficient matrix, and the fact that the close relation of our problem to classical network problems implies that the linear programming basis matrix is likely to be nearly triangular at each step of the simplex method.

This variant of the simplex method has been used in several other parts of the PAWN study. It was the method used in the Response Design Model (RESDM) for designing long-run pricing and regulation strategies for drinking water companies and industries (Vol. IV), and for assessing impacts on drinking water companies and their customers (Vol. VII). In addition, it was the method used in the Electric Power Reallocation and Cost (EPRAC) model, which was used in conjunction with Distribution Model outputs for determining the changes in electric power generation costs due to water management decisions (Vol. XV).

8.7. CONVERGENCE

We mentioned above that the iterative scheme outlined in Sec. 8.5 need not converge to a solution of Problem (8.1) that satisfies the optimality conditions (8.2) and (8.3). We can use example network from Fig. 8.1 to illustrate this lack of convergence as follows.

Let the situation be as depicted in Fig. 8.1, and suppose the objective function is to minimize the sum of the pollutant concentrations at nodes 2 and 3. One sees immediately that the solution is to divide the flow evenly between the two branches, which results in concentrations at both nodes of 2.0 kg/m^3, and a value of the objective function of 4.0. But it is a property of the simplex method for solving linear programs that it always seeks extreme solutions, and so initially it will send as much water as possible along one of the branches, say from node 1 to node 2, and as little as possible along the other.

If we allow zero flow on either branch, the concentration at either node 2 or node 3 could become infinite. Thus, we will suppose that the flow on each branch must be at least 0.01 m^3/s. Our initial solution, therefore, is $x = (0.01, 0.99)$, and the corresponding concentrations are $y = (100.0, 1.0101)$. The value of the objective is 101.0101, clearly considerably larger than the optimal value of 4.0.

It can be shown that the new term we add to the objective function of the water distribution problem, (8.5), is equal to the derivative of the optimal value of the water quality objective function, taken with respect to the flow and water level variables. In our example, we can calculate the water quality objective function analytically; if $x = (F, 1-F)$, then the corresponding $y = (1/F, 1/(1-F))$, and the value of the objective is $h(y) = 1/F + 1/(1-F)$. The derivative must be:

$$\frac{dh}{dF} = -\frac{1}{F^2} + \frac{1}{(1-F)^2}$$

Evaluating this at F = 0.01, we obtain a derivative of -9998.98, which becomes the coefficient of F in the objective function of the revised problem (8.5). Clearly, since we are minimizing the objective function, such a large, negative coefficient will result in in F being made as large as possible in the next solution. Thus, our next solution will be x = (0.99, 0.01). But this is essentially the same situation we have just analyzed, with the roles of nodes 2 and 3 interchanged. At the next iteration, therefore, we will once again arrive at the first solution; next we will bounce back to the second; and so on. Convergence will never occur.

Two methods suggest themselves for breaking this cycle. First, we could make the water quality term in the objective function of the water distribution problem nonlinear, so it would better approximate the behavior of the optimal value of h(y) as the water distribution x is changed. Second, we could restrict the magnitude of the change in the water distribution from one iteration to the next, so that it would remain in the range where our approximation of the optimal h(y) is reasonably good.

We decided against the first method because more sophisticated approximations of the dependence of the optimal h(y) on x are difficult to calculate. Furthermore, even in simple cases when the optimal h(y) can be expressed analytically as a function of x, it is not necessarily convex (although in our example, it is). Finally, more sophisticated approximations would require us to use a different, less efficient algorithm than the simplex method for solving the water distribution problem.

Instead, we restricted the range over which the water distribution x was permitted to vary. We did this by imposing more limiting upper and lower bounds on key flows and water levels than are required in reality. For example, the flow in the Hollansche IJssel (link 61 HOLIJSEL) can actually vary from -35 m^3/s to as high a flow in the positive direction as desired--certainly as high as +50 m^3/s. In a particular run of MSDM we might require that it never become higher than 0 m^3/s. If oscillations still occurred, we might restrict it still further, or we might restrict a second variable. Alterations of these upper and lower bounds were carried out manually, and not by a computerized algorithm.

In some cases, it was evident that a particular variable must be in a given range in the optimal solution, and we could simply require it to remain there. In other cases, it was not clear what ranges were proper for all the key variables, and several different ranges had to be investigated. In the case of the Hollandsche IJssel, for example, it was not always clear whether reducing the water supply to the midwest was more costly than meeting the demand with saline water. In such a case we would make two MSDM runs, one restricting the flow in the Hollansche IJssel to be zero or negative (into the midwest), and the other restricting it to be zero or positive (out of the midwest).

In addition to providing a means to force convergence of MSDM, this method offered a means to explore different managerial strategies, in order to generate several good strategies that might be used in the same circumstances.

8.8. SOME DETAILS AND STATISTICS ON THE USE OF MSDM

In this section we have collected a number of facts that a person should know if he plans to use MSDM. First, we implemented MSDM on an IBM 370 model 3032, where it occupied 420K bytes of core. The model is programmed entirely in FORTRAN IV.

The problem we were solving had the following dimensions. The water distribution subproblem had 59 explicit constraints and 220 variables. (There were numerous implicit constraints, which corresponded to upper and lower bounds on each of the variables.) The BOD, phosphate, and heavy metal subproblems had 37 explicit constraints and 42 variables each. The salt subproblem had 38 explicit constraints and 51 variables, while the thermal subproblem had 44 explicit constraints and 91 variables. In each subproblem, 37 of the constraints correspond to the 37 nodes in the MSDM network. In the water distribution subproblem, the remaining explicit constraints consist of the upper and lower rivers hydrologic constraints, and numerous equations that relate shipping depths to flows at critical points in the network. In the thermal subproblem, two of the remaining explicit constraints relate to the two generating units that have optional cooling towers, while the other five constraints define the demands for electricity in the five generating regions.

We solved many particularizations of this problem, each representing a different set of conditions that might be encountered in a given decade (i.e., a ten-day period). We call each particularization a case. To define a case, we had to provide the appropriate boundary conditions (river flows and pollutant discharges into various nodes) and initial conditions (water levels and pollutant concentrations at nodes with storage). Generally, we solved several cases during each MSDM run, but because we specified boundary conditions and initial conditions exogenously for each case, the several solutions were mutually independent. The alternative would have been to initialize each case with water levels and pollutant concentrations from the solution to the previous case. In this event, the several cases in a run would be considered as a sequence of consecutive decades.

A typical MSDM run requires about five CPU seconds to read data, initialize arrays, etc., at a cost at our installation of two or three U.S. dollars. Then, each case involved in the run requires an additional 25 CPU seconds, at a cost of seven or eight dollars. Our runs typically consisted of from one to four cases, so the average cost per case was approximately ten dollars. This figure is misleading, however, since there were few cases for which MSDM converged without manual intervention. Typically, cases had to be modified manually and rerun four, five, or six times before a suitable solution was obtained. Our true cost per case, therefore, was closer to 60 dollars than to 10 dollars.

It is difficult to suggest a general strategy for successful manual intervention. (If it were easy, we could have automated it, making manual intervention unnecessary.) One must remember that MSDM only fails to converge because the influence of the water quality subproblems renders the overall problem nonconvex. Thus, oscillations in the solution from one iteration to the next will always involve flows on links which have an influence on the concentration of some pollutant at one or more nodes. Further, the more nodal concentrations that are influenced by a particular link flow, and the greater that influence on each, the more likely constraining that flow will reduce or eliminate the oscillations. This observation is not very helpful in cases where high costs are imposed for violations of BOD, phosphate, or heavy metal standards, since many links have a major effect on the concentration of at least one of these pollutants at one or more nodes. But fewer links have a major effect on salt or thermal pollution, so if BOD, phosphate, and heavy metal violations are ignored or treated lightly, the observation can help us.

The single most important salt-related phenomenon is the Rotterdam salt wedge. In MSDM, its effect is a function of the flows on links 9 NIEWMAAS and 61 HOLIJSEL (see Chap. 4). Thus, constraining the flows on either of these links, or on any link whose flow affects these links (e.g., link 6 NEDRIJN2), strongly tends to damp oscillations relating to the salt subproblem.

Oscillations due to the thermal subproblem most often concern the power plants on the Noordzeekanaal, and sometimes also the Bergum power plant. On one iteration MSDM will provide little cooling water to these plants, and deduce that providing more cooling water would reduce the total cost. On the next iteration MSDM will provide too much cooling water, and deduce that it can cut back at no cost. Accordingly, on the third iteration MSDM returns to the first solution. These oscillations can be stopped by constraining the flows on links 34 NZKANLSL and 45 MARGKAN2, although some experimentation may be needed to determine at what values the flows should be constrained.

If constraining some combination of these links fails to produce a satisfactory solution, the user must try to isolate some other aspect of a subproblem that is responsible for the oscillations. The output from MSDM includes certain quantities called shadow prices that may help the user in this task. In every subproblem, each constraint, whether explicit or implicit, has an associated shadow price, which measures the amount by which the total cost (i.e., the value of the objective function) would change if that constraint were relaxed by a unit amount. For example, the shadow price associated with the water balance constraint at node 6 UPRIVER, where the Rijn enters the Netherlands, measures the value of increasing the Rijn flow by one cubic meter per second. For the implicit constraints--the upper and lower bounds on the variables--one knows *a priori* what sign the associated shadow price should have. Under the conventions of MSDM, they should all be positive or zero; and if they are, the solution is (locally) optimal and no oscillations will take place.

The MSDM solution algorithm ensures that all shadow prices for constraints in the water quality subproblems will have the proper sign. If there are oscillations, therefore, the shadow prices associated with upper or lower bounds on some of the water distribution variables--i.e., flows on network links--must be negative, and these flows must be suspected of involvement in the causes of the oscillation. Unfortunately, the flows are all interrelated due to the explicit constraints of MSDM, so some flows with positive shadow prices may also be involved. However, we have found that the negative shadow prices are usually all associated with link flows in the same geographical part of the network, e.g., the North, or the Southern Highlands. (In some cases two or more parts may be represented, but never the entire network.) Hard thinking about what features of which subproblems might be strongly influenced by flows in the indicated part of the network, plus judicious trial and error, should lead to a satisfactory solution.

REFERENCES

8.1. Kuhn, H. W., and A. W. Tucker, "Nonlinear Programming," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, J. Neyman (ed.), University of California Press, Berkeley, California, 1950.
8.2. Hillier, F. S., and G. J. Lieberman, *Operations Research*, Holden-Day, Inc., San Francisco, California, 1974.
8.3. Dantzig, G. B., *Linear Programming and Extensions*, Princeton University Press, Princeton, New Jersey, 1963.

Chapter 9

AN "OPTIMAL" MANAGERIAL STRATEGY FOR PRESENT CONDITIONS

9.1. DESCRIPTION OF A MSDM CASE

In this section, we describe a case for MSDM to approximate conditions as they will be in the near future if the only changes from the present are those already decided upon. Thus:

- Infrastructure as in 1976, except that the Zoommeer has been made a fresh water lake (see App. A of Vol. VA).
- Electric power plant inventory of 1985, since most of the changes from the 1976 inventory had been made by 1979 (see App. E of Vol. VA).
- Sprinkling scenario of 1976 (see App. F of Vol. VA).
- Shipping scenario of 1976 (see App. G of Vol. VA).

We choose as external supplies of water and pollutants those observed during the first decade in July 1976 (decade number 19). The Rijn and the Maas flows were 855 m^3/s and 26.6 m^3/s, respectively. The Maas flow of 26.6 m^3/s includes 22.4 m^3/s which treaties require be sent to Belgium, so only 4.2 m^3/s is available for use by the Netherlands. Flows in the smaller rivers were: the Overijsslesche Vecht, 1.8 m^3/s; the Roer plus the Swalm, 11.0 m^3/s; and the Niers, 2.8 m^3/s. The concentrations of chromium, BOD, and phosphate in the rivers can be found in App. C of Vol. VA, salt in App. D of Vol. VA, and excess temperature (for thermal pollution) in Chap. 5 of this volume.

Finally, we use the rainfall and open water evaporation rate observed during the first decade in July 1976. This was an extremely dry decade, having no rainfall and 66 mm of open water evaporation.

We assumed all nodes with storage were filled to capacity with water at the start of the decade, and that the concentrations of pollutants were as shown in Table 9.1 below. Except for the initial excess temperatures, the initial pollutant concentrations came from data described in Chap. 3 of this volume, and App. C of Vol. VA, and they reflect typical summer concentrations from 1976. The analysis in Chap. 5 suggests that of the nodes with storage, HAL+IJMU and LINNE are likely to be at the limiting temperature of three degrees, so we assumed that water stored there had already reached that temperature. We took the initial excess temperatures at MAASLOO and ROER+BEL to be the steady state temperatures, that would be achieved if the flow conditions described above persisted indefinitely.

Initial conditions for the aggregate plots in MSDM were as described in App. F of Vol. VA. We assumed that groundwater levels had

Table 9.1

INITIAL POLLUTANT CONCENTRATIONS IN STORAGE

Node No.	Name	Salt	Excess Temp.	Chromium	BOD	Phosphate
1.	FRIELAND	236.0	0.0	0.0	15.0	0.8
10.	IJSLAKES	219.0	0.0	5.1	5.5	0.51
11.	NORHOLL	288.0	0.0	0.0	10.7	0.8
15.	AMSTEDAM	720.0	0.0	2.7	4.7	0.4
16.	HAL+IJMU	2000.0	3.0	4.1	3.8	0.3
17.	LOPIKWAR	206.0	0.0	0.0	5.5	2.2
18.	VECHT	114.0	0.0	4.4	6.4	3.1
21.	MIDWEST	300.0	0.0	3.8	7.0	1.0
25.	LOWRIVER	246.0	0.0	10.0	4.5	0.5
26.	DLTALAKE	300.0	0.0	0.0	0.0	0.0
27.	MAASLOO	50.0	0.41	5.5	5.7	0.5
28.	BORN+PAN	90.0	0.0	0.0	3.7	0.5
29.	LINNE	90.0	3.0	0.0	4.5	0.6
30.	ROER+BEL	90.0	0.83	0.0	5.0	0.8
31.	SAMGRALI	90.0	0.0	0.0	5.4	0.7

dropped one meter below their springtime highs, and that all root zones were dry enough that a failure to sprinkle in the current decade would result in drought damage to crops (but wet enough that damage could be avoided by sprinkling). Initial root zone and subsoil salt concentrations were calculated as appropriately weighted averages of the concentrations used in PAWN's primary model of agriculture, DISTAG. These concentrations were low enough that virtually no salt damage occurred in the MSDM runs to crops grown in the open air.

For our initial case, we relaxed all water quality standards to a point at which they would have no influence on the choice of a managerial strategy. Initially we wanted to compare model results with present practices. With only two exceptions, present practice does not adjust the flows of water in the network to improve water quality. Present practice does provide minimum flows past the Velsen power plant (at node 16 HAL+IJMU) and the Bergum power plant (at node 2 B'GUMMER) to control the excess temperatures. However, the thermal standard at these nodes is seven degrees, rather than the three degrees we have chosen as our nominal thermal standard for PAWN.

These inputs completely define the case to be run. Running the case involves the procedure outlined in Chap. 8. This procedure seeks a water managerial strategy which minimizes an objective function consisting of the following terms:

- Pumping energy cost, but not for back pumping at fresh-fresh locks (see Chap. 2).

- Salt damage to crops grown under glass (see Chap. 4).
- Salt damage to crops due to locally confined salt intrusion through locks (see Chap. 4).
- Marginal costs of salt intrusion abatement tactics at locks, including delay costs to shipping of managerial tactics at locks (see Chap. 4).
- Fuel cost for generating electric power (see Chap. 5).
- Marginal cost of sprinkling surface water on crops grown in the open air (see Chap. 6).
- Present and expected future drought damage to crops grown in the open air and sprinkled from surface water (see Chap. 6).
- Salt damage to crops grown in the open air and sprinkled from surface water (see Chap. 6).
- Shipping low-water losses (see Chap. 7).
- Costs of storing goods that cannot be shipped immediately (see Chap. 7).
- Costs due to sedimentation in the Waal downstream of Tiel and/or St. Andries. The costs may either be dredging costs, or expected future losses to shipping due to the reduction in depth occasioned by the sandbar (see Chap. 7). In the case described here, we took this cost to be the dredging cost.
- Marginal costs of water conservation at fresh-fresh locks, including shipping delays, lock operating cost, and cost of energy for back pumping (see Chap. 7).

We introduced two additional terms into the MSDM objective function. One term placed a small cost on reducing the amount of water stored in the IJssel lakes (node 10 IJSLAKES). We arbitrarily chose a cost of one Dflm for a 40-cm reduction in the IJssel lake level, which is equivalent to a value of stored water of 0.00122 Dfl/m^3. The other term placed a similar small cost on lowering the levels in the weir ponds of the Maas (nodes 27 MAASLOO, 28 BORN+PAN, 29 LINNE, 30 ROER+BEL, and 31 SAMGRALI). Here we arbitrarily chose a cost of 0.1 Dflm for a 40-decimeter reduction in the levels of all the weir ponds, or 0.00101 Dfl/m^3. The particular values chosen are not important, since both are quite small. Their only effect is to discourage the wasting of water in storage when there is nothing else to do but save it for a future use. Without such cost terms, MSDM would often find itself indifferent between retaining large amounts of water in storage on the one hand, or discharging much of the stored water into the North Sea on the other.

But retaining water has potential value, as insurance against possible water shortages in future decades. The terms we have added to the objective function, which we have referred to in PAWN as insurance functions, can therefore be considered as representing the expected future value of saving water for future uses. Indeed, we have used MSDM in an attempt to estimate the expected future value of water stored in the IJssel lakes, which we report in Sec. 9.5 below. No such attempt has been made for water stored in the Maas, but we feel it would be worthwhile to estimate this value.

9.2. OPTIMAL MANAGERIAL STRATEGY FOR THE INITIAL CASE

MSDM has found the optimal managerial strategy for the initial case described above. Tables 9.2 through 9.8 show this strategy, along with a strategy that approximates the present Dutch water management practice. In these tables, the optimal strategy found by MSDM is labeled the MSDM strategy, while the strategy that approximates present Dutch practice is labeled the RWS strategy. (RWS stands for Rijkswaterstaat, which is the Dutch government ministry ultimately responsible for water management decisions. Our client for PAWN, and the agency with the direct responsibility, is an agency within the RWS called WW, or Waterhuishouding en Waterbeweging.) The reader should keep in mind that the RWS strategy is only an approximation to Dutch practice, because the Dutch do not have a well-defined water management strategy. But they do have rough rules of thumb, which they apply most (but not all) of the time, and which serve as the basis for the RWS strategy presented here.

Table 9.2

COMPARISON OF THE MSDM STRATEGY WITH THE RWS STRATEGY:
OVERALL STRATEGY COSTS

Case Description			
Pollutant Penalties:	None		
Rijn Flow:	855.00 m^3/s		
Maas Flow:	26.60 m^3/s		
Net Evaporation:	66.00 mm		
Initial Lake Levels:	Maximum		
Infrastructure:	1976 + Fresh Zoommeer		
Power Plant Inventory:	1985		
Sprinkling Scenario:	1976		
Shipping Scenario:	1976		
Cost Components (Dflm/dec)		**MSDM Strategy**	**RWS Strategy**
Shipping		15.02	14.19
Agriculture			
Damage from Salt		186.26	192.63
Drought Damage Plus Sprinkling Cost		37.52	36.79
Thermal		71.58	71.58
Other Costs		1.09	1.13
Total Costs		310.83	316.32

Table 9.3

COMPARISON OF THE MSDM STRATEGY WITH THE RWS STRATEGY:
STRATEGIES FOR THE NATIONAL WATERWAYS

Case Description

Pollutant Penalties: None
Rijn Flow: 855.00 m^3/s
Maas Flow: 26.60 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum
Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Strategy Description	MSDM Strategy	RWS Strategy
Upper Rivers Flows (m^3/s)		
Waal	657.09	670.30
Neder-Rijn	47.77	25.00
IJssel	141.06	150.62
Lower Rivers Flows (m^3/s)		
Nieuwe Maas	207.00	163.39
Oude Maas	388.30	396.77
Nieuwe Waterweg	595.30	560.16
Volkerak + Haringvlietsluizen	40.11	40.11
Amsterdam-Rijnkanaal Flows (m^3/s)		
Withdrawal at Tiel	51.39	37.24
North from A-R Mond	9.49	32.10
Lek	64.53	5.00
Noordzeekanaal Flows (m^3/s)		
Input from IJssel Lakes	6.36	11.36
Discharge to North Sea	20.00	40.00
IJssel Lakes Level Change (cm)	-6.63	-6.30
Maas Weir Pond Level Changes (cm)		
MAASLOO	-87.36	-55.01
BORN+PAN	0.0	0.0
LINNE	0.0	0.0
ROER+BEL	0.0	0.0
SAMGRALI	0.0	0.0
Maas Flows (m^3/s)		
JULCANL1	3.58	2.5
MAAS3	11.12	10.04
AMER1 (Discharge into Lower Rivers)	15.00	13.92

Table 9.4

COMPARISON OF THE MSDM STRATEGY WITH THE RWS STRATEGY:
SOUTHEAST HIGHLANDS AND DELTA REGION

Case Description		
Pollutant Penalties:	None	
Rijn Flow:	855.00 m^3/s	
Maas Flow:	26.60 m^3/s	
Net Evaporation:	66.00 mm	
Initial Lake Levels:	Maximum	
Infrastructure:	1976 + Fresh Zoommeer	
Power Plant Inventory:	1985	
Sprinkling Scenario:	1976	
Shipping Scenario:	1976	
Strategy Description (m^3/s)	**MSDM Strategy**	**RWS Strategy**
Water for Sprinkling	3.63	3.63
Other Consumption	0.33	0.33
Extraction from Upper Maas	4.84	4.84
Extraction from Delta Lake	1.49	1.49
Discharge to Lower Maas	2.38	2.38
Target Sprinkling Water	3.63	3.63
Cutbacks in Sprinkling	0.0	0.0
Regional Agricultural Damage (Dflm)		
Damage from Salt	0.0	0.0
Drought Damage Plus Sprinkling Cost	2.64	2.64

Table 9.5

COMPARISON OF THE MSDM STRATEGY WITH THE RWS STRATEGY:
NORTHEAST HIGHLANDS REGION

Case Description

Pollutant Penalties: None
Rijn Flow: 855.00 m^3/s
Maas Flow: 26.60 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum
Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Strategy Description (m^3/s)	MSDM Strategy	RWS Strategy
Water for Sprinkling	4.69	4.69
Other Consumption	-4.56	-4.56
Overijsselsche Vecht	1.8	1.8
Extractions from IJssel	5.89	5.89
Discharge to IJssel	7.36	7.36
Discharge to North	0.2	0.2
Target Sprinkling Water	5.09	5.09
Cutbacks in Sprinkling	0.40	0.40
Regional Agricultural Damage (Dflm)		
Damage from Salt	0.12	0.12
Drought Damage Plus Sprinkling Cost	2.37	2.37

Table 9.6

COMPARISON OF THE MSDM STRATEGY WITH THE RWS STRATEGY:
NORTHERN REGION

Case Description		
Pollutant Penalties:	None	
Rijn Flow:	855.00 m^3/s	
Maas Flow:	26.60 m^3/s	
Net Evaporation:	66.00 mm	
Initial Lake Levels:	Maximum	
Infrastructure:	1976 + Fresh Zoommeer	
Power Plant Inventory:	1985	
Sprinkling Scenario:	1976	
Shipping Scenario:	1976	

Strategy Description (m^3/s)	MSDM Strategy	RWS Strategy
Water for Sprinkling	15.13	15.13
Other Consumption	41.57	41.57
Extraction from IJssel Lakes	62.49	66.99
From Northeast Highlands	0.20	0.20
Discharges to Waddenzee	6.00	10.50
Target Sprinkling Water	15.13	15.13
Cutbacks in Sprinkling	0.0	0.0
Regional Agricultural Damage (Dflm)		
Damage from Salt	1.02	1.02
Drought Damage Plus Sprinkling Cost	4.53	4.53

Table 9.7

COMPARISON OF THE MSDM STRATEGY WITH THE RWS STRATEGY:
NORTH HOLLAND REGION

Case Description		
Pollutant Penalties:	None	
Rijn Flow:	855.00 m^3/s	
Maas Flow:	26.60 m^3/s	
Net Evaporation:	66.00 mm	
Initial Lake Levels:	Maximum	
Infrastructure:	1976 + Fresh Zoommeer	
Power Plant Inventory:	1985	
Sprinkling Scenario:	1976	
Shipping Scenario:	1976	

Strategy Description (m^3/s)	MSDM Strategy	RWS Strategy
Water for Sprinkling	6.33	6.33
Other Consumption	13.80	13.80
Extractions from IJssel Lakes	30.00	30.00
Discharges to Waddenzee	2.87	2.87
Discharges to Noordzeekanaal	7.00	7.00
Target Sprinkling Water	6.33	6.33
Cutbacks in Sprinkling	0.0	0.0
Regional Agricultural Damage (Dflm)		
Damage from Salt	1.31	1.31
Drought Damage Plus Sprinkling Cost	2.67	2.67

Table 9.8

COMPARISON OF THE MSDM STRATEGY WITH THE RWS STRATEGY: MIDWEST AND UTRECHT REGION

Case Description		
Pollutant Penalties:	None	
Rijn Flow:	855.00 m^3/s	
Maas Flow:	26.60 m^3/s	
Net Evaporation:	66.00 mm	
Initial Lake Levels:	Maximum	
Infrastructure:	1976 + Fresh Zoommeer	
Power Plant Inventory:	1985	
Sprinkling Scenario:	1976	
Shipping Scenario:	1976	

Strategy Description (m^3/s)	MSDM Strategy	RWS Strategy
Water for Sprinkling	4.40	7.98
Other Consumption	32.72	32.72
Extraction from Hollandsche IJssel	16.84	20.42
Extraction from Amsterdam-Rijnkanaal	24.10	24.10
Extractions from IJssel Lakes	10.00	2.40
Discharges to Amsterdam-Rijnkanaal	7.60	0.0
Discharges to Noordzeekanaal	3.00	3.00
Discharges to North Sea	3.23	3.23
Target Sprinkling Water	7.98	7.98
Cutbacks in Sprinkling	3.58	0.0
Regional Agricultural Damage (Dflm)		
Damage from Salt	161.97	168.26
Drought Damage Plus Sprinkling Cost	4.06	3.34

We have organized the discussion of these two strategies into separate sections dealing with national and regional parts of the MSDM network. The five regions are:

- Southeast Highlands and Delta Region, which contains nodes
  - 34. WEER+MEY
  - 35. HELMOND
  - 36. OSTRHOUT
  - 37. FIJNAART

- Northeast Highlands Region, which contains nodes
  - 4. NEHIGH
  - 5. OVIJVECH
  - 8. TWENTEND

- Northern Region, which contains nodes
  - 1. FRIELAND
  - 2. B'GUMMER
  - 3. GRONETAL

- North Holland Region, which contains the single node
  - 11. NORHOLL

- Midwest and Utrecht Region, which contains nodes
  - 17. LOPIKWAR
  - 18. VECHT
  - 20. GOUDA
  - 21. MIDWEST

All nodes which are not in one of the five regions are national nodes. Figure 9.1 shows the MSDM network, with the nodes in the different regions delineated.

The regions we have defined here are closely related to the sections of the network defined in App. A of Vol. VA. The sections differ from the regions in that they contain both regional and national nodes. But if the national nodes (as defined here) are discarded, then the following correspondence exists between the regions and the earlier sections. The Southeast Highlands and Delta region corresponds to the Upper Rijn and Southeast Highlands section, plus the Lower Rijn and Delta section. We combined the two sections because, after the national nodes were removed, only one node (37 FIJNAART) remained in the Lower Rijn and Delta section. The remaining four regions correspond to the four sections of the same names--i.e., the Northeast Highlands region to the Northeast Highlands section, the Northern region to the Northern section, etc.

9.2.1. Overall Strategy Costs

Table 9.2 shows the overall costs of both the MSDM strategy and the RWS strategy, broken down into their major components. The largest

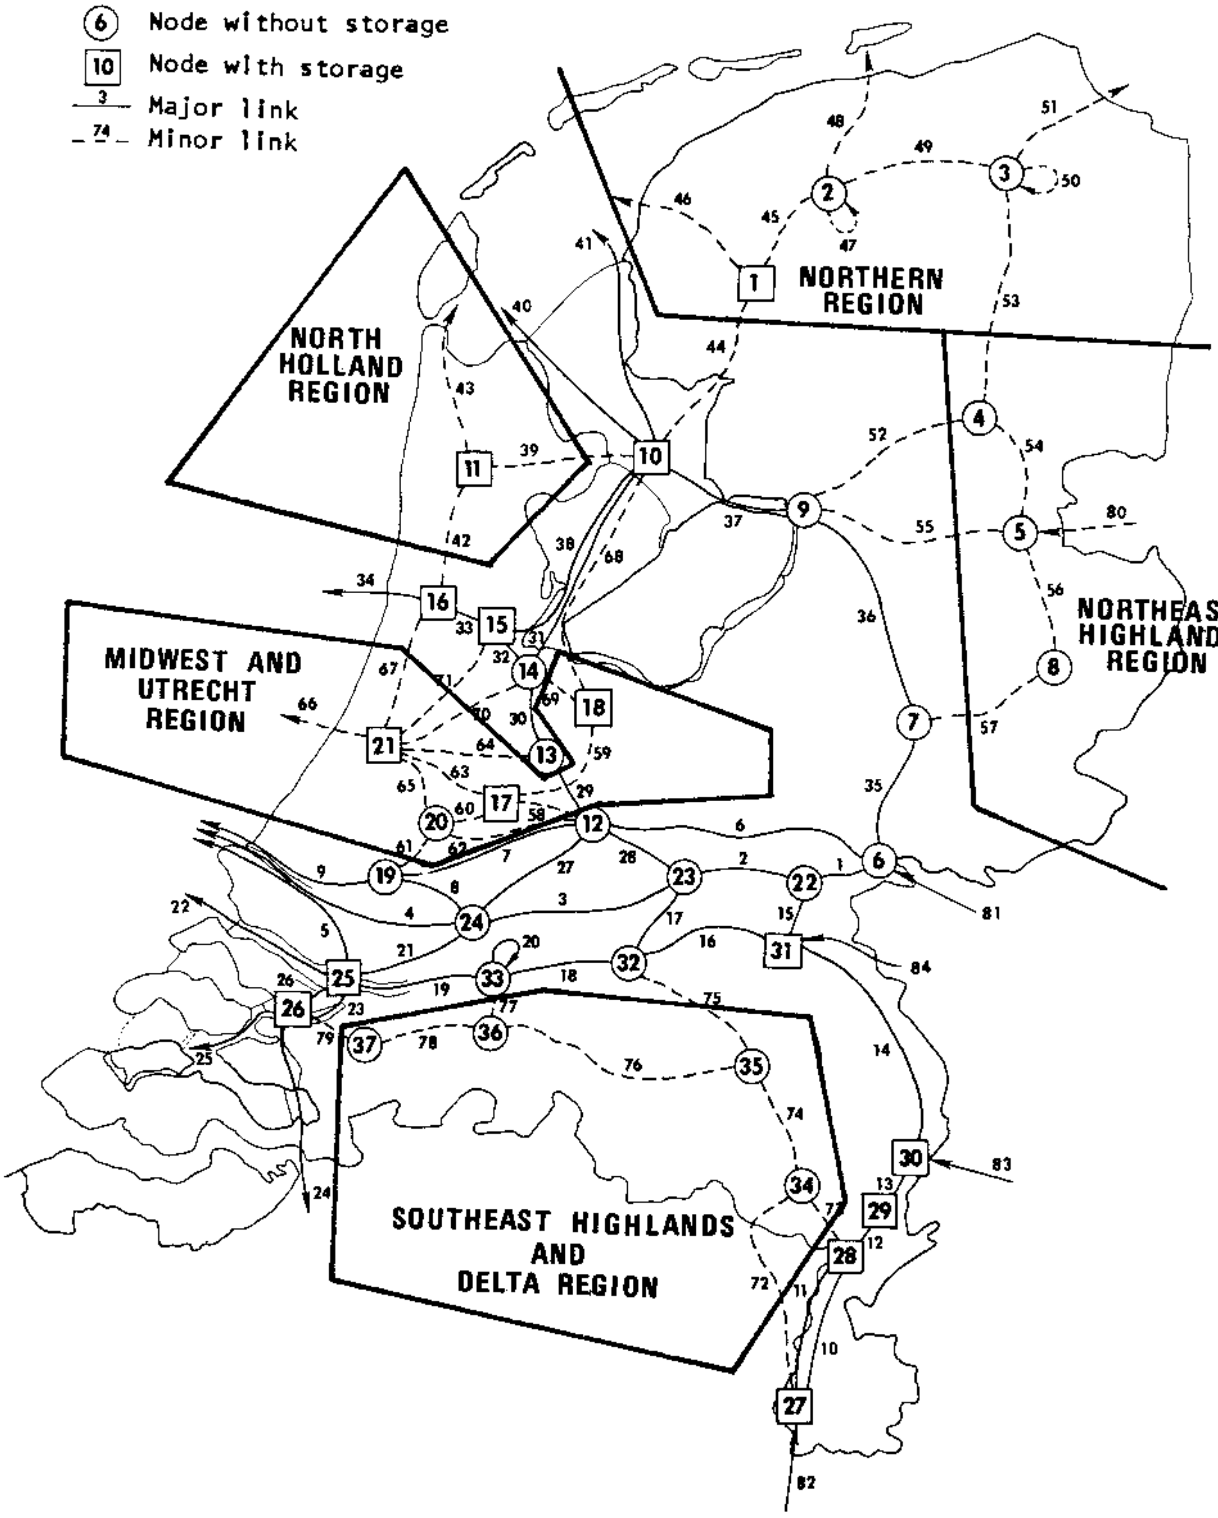

Fig. 9.1--The five regional parts of the MSDM network
(nodes not in a region are called national nodes)

single component is agricultural losses, which we have further broken down into damage from salt and drought damage plus sprinkling costs. Next in size are fuel costs for electric power generation, labeled thermal costs in the table. Shipping costs are fourth largest, and all other costs (e.g., pumping, lock operating costs) make up the very small remainder.

A mere comparison of the relative sizes of the components is misleading, however, since different fractions of each costs component are unavoidable--i.e., cannot be affected by the managerial strategy. For example, most of the agricultural damage from salt is damage to crops grown under glass in the Midwest. But these crops are supplied from node 21 MIDWEST, a node whose stored water has a salt concentration of 300 mg/l at the start of the decade. Using the information from Chap. 4, the salt damage in the Midwest alone due to a salt concentration of 300 mg/l should exceed 150 Dflm. Because the volume of water stored at this node is so large (119 million m^3), no managerial strategy can much affect this concentration in a single decade. Thus, 150 Dflm is approximately a lower bound on the amount of salt damage that must occur in the Midwest. Unavoidable salt damage will occur elsewhere too, so almost all of the salt damage shown in Table 9.2 for both strategies is unavoidable.

Similarly, as discussed at length in Chap. 5, the thermal cost shown for both strategies is the smallest this cost can possibly be. Much of the shipping cost is also unavoidable, as can be seen by the fact that the shipping cost functions (see App. G of Vol. VA) have minima considerably above zero. In all, probably less than 20 Dflm of the costs are affected by managerial tactics.

Thus, it is wiser to examine the differences between the cost components from one strategy to another, rather than to compare different components for the same strategy. In Table 9.2, the greatest difference between the strategies occurs in the salt damage component, and the next greatest in the shipping component. The other components are hardly different. The major cost differences are due to differences in the national aspects of the two strategies, and will be discussed in Sec. 9.2.2. The smaller, almost incidental differences, are mostly due to small differences in the regional aspects of the strategies, and will be discussed later, in Sec. 9.2.3.

9.2.2. Strategies for the National Waterways

In this section we discuss the MSDM and RWS strategies in the national part of the network. The reader will find it helpful to refer frequently to Fig. 9.1.

9.2.2.1. The Waal, Neder-Rijn, and Delta. Table 9.3 shows both the MSDM and RWS strategies in the national waterways. In the MSDM strategy, the weir at Driel should be opened far enough to permit 47.77 m^3/s to flow down the Neder-Rijn. Another 9.08 m^3/s will be extracted by areas served by the Linge river, and areas served directly from the Pannerdensch Kanaal. According to the hydrological equation for upper rivers, the remainder of the Rijn flow entering the Netherlands would be divided, 657.09 m^3/s flowing down the Waal and 141.06 m^3/s flowing north along the IJssel.

Farther west, the rather large amount of 51.39 m^3/s is withdrawn from the Waal at Tiel. Most of this water flows west along the Lek to combat the Rotterdam salt wedge. Only 9.49 m^3/s is sent north along the Amsterdam-Rijnkanaal.

In the lower rivers area, the flow in the Neder-Rijn and the withdrawal at Tiel discussed above result in a flow in the Nieuwe Maas of 207 m^3/s and a flow in the Oude Maas, just before it joins with the Nieuwe Maas, of 388.3 m^3/s. The Nieuwe Waterweg is formed by the junction of the Nieuwe Maas and the Oude Maas; the flow here is 595.3 m^3/s.

It is no accident that MSDM chooses a Nieuwe Maas flow of exactly 207 m^3/s. Looking back to App. D of Vol. VA, one finds that the damage to crops under glass due to the Rotterdam salt wedge is virtually eliminated at flows in the Nieuwe Maas of 207 m^3/s and above. Only the few crops under glass supplied from node 19 IJSLMOND will be affected. The transportable component of the salt wedge representation in MSDM, which is the only component that affects the salinity at nodes 20 GOUDA and 21 MIDWEST (see Chap. 4), first becomes zero at this flow. At smaller flows, the transportable component is positive, and the damage due to the salt wedge can be considerable.

The RWS strategy under the same circumstances closes the weir at Driel almost completely, reducing the Neder-Rijn flow to its minimum of 25 m^3/s. This increases the Waal flow to 670.3 m^3/s and the IJssel flow to 150.62 m^3/s, benefiting shipping on both rivers as compared with the MSDM strategy. In addition, the RWS strategy reduces the withdrawal at Tiel to the minimum necessary to maintain a nominal flow of 5 m^3/s on the Lek. Since 32.1 m^3/s is sent north on the Amsterdam-Rijnkanaal, the minimum withdrawal from Tiel is 37.24 m^3/s. (We explain later why the RWS strategy sends 32.1 m^3/s up the Amsterdam-Rijnkanaal, while the MSDM strategy sends only 9.49 m^3/s.) As explained in Chap. 7, a large withdrawal at Tiel lowers the water level at the Waal critical point for shipping, and also causes a sandbar to form. These phenomena cause the MSDM strategy to have higher shipping costs than the RWS strategy (see Table 9.2).

In the RWS strategy, the decisions concerning the weir at Driel and withdrawals at Tiel imply that the flow on the Nieuwe Maas is only 163.39 m^3/s. This gives rise to a present and expected future damage to crops under glass (grown mostly in the Midwest) 6.37 Dflm larger than that suffered under the MSDM strategy (see Table 9.2), which more than offsets the 0.83 Dflm benefit to shipping.

Shipping and agricultural salt damage account for nearly all of the difference of 5.49 Dflm between the total cost of the RWS strategy and the minimum cost found by MSDM. This is a considerable amount of unnecessary damage to assert is caused by the RWS strategy.[1] Indeed, it is probably an overestimate, because measures not represented in the model are available in reality to avoid some of the damage. The extra damage is entirely due to salt damage in the midwest, caused by the additional salt intrusion from the Rotterdam salt wedge that is permitted by the RWS strategy. In reality, it requires more than a week for the effect of low flows in the Nieuwe Maas to be reflected by an increase in the salt concentration at Gouda. The effect can be further delayed by closing the storm-surge barrier at the mouth of the Hollandsche IJssel, and drawing down the impounded water to supply the midwest. This option is not available in MSDM, where the effect is assumed to be instantaneous. Finally, waterboards in the midwest can probably meet their demands for a few days by drawing down the level of water in the boezems and ditches a few centimeters, instead of admitting saline water at Gouda. MSDM demands that the boezems and ditches be maintained at constant levels.

It is also important to remember that all of this damage would not be experienced in the single decade. Salt damage in the midwest consists almost entirely of damage to crops grown in glasshouses. Our calculation of damage to these crops includes not only the damage in the present decade, but damage which is expected to occur in future decades due to the persistence of salt admitted in the present decade. In the case of the midwest, future damage accounts for almost all of the total salt damage.[2]

The trade-off between shipping losses and damage to glasshouse crops in the midwest is an important consideration in choosing an optimal managerial strategy. This trade-off is governed by two variables, the flow in the Neder-Rijn (which is an index of the positioning of the weir at Driel), and the withdrawal from the Waal at Tiel. Figure 9.2 shows how the total cost associated with the strategy depends on these two variables. The star denotes the MSDM strategy, and corresponds to the minimum possible cost of 310.83 Dflm. The RWS strategy cannot be represented in this figure because there are more differences between the two strategies than those that follow necessarily from the flow in the Neder-Rijn and the withdrawals at Tiel. The figure was constructed assuming that no such differences exist.

Note that the total strategy cost is hardly affected by variations that increase the Neder-Rijn flow by 0.77 m^3/s for each one m^3/s decrease in the withdrawal at Tiel. From the lower rivers hydrological equations (see App. A of Vol. VA), one can calculate that changes in these two flows in exactly the stated proportion will maintain the flow in the Nieuwe Maas at a constant level, and hence the Rotterdam salt wedge at a constant position. Therefore, these variations will not affect the losses due to the salt wedge. In addition, an increase in the

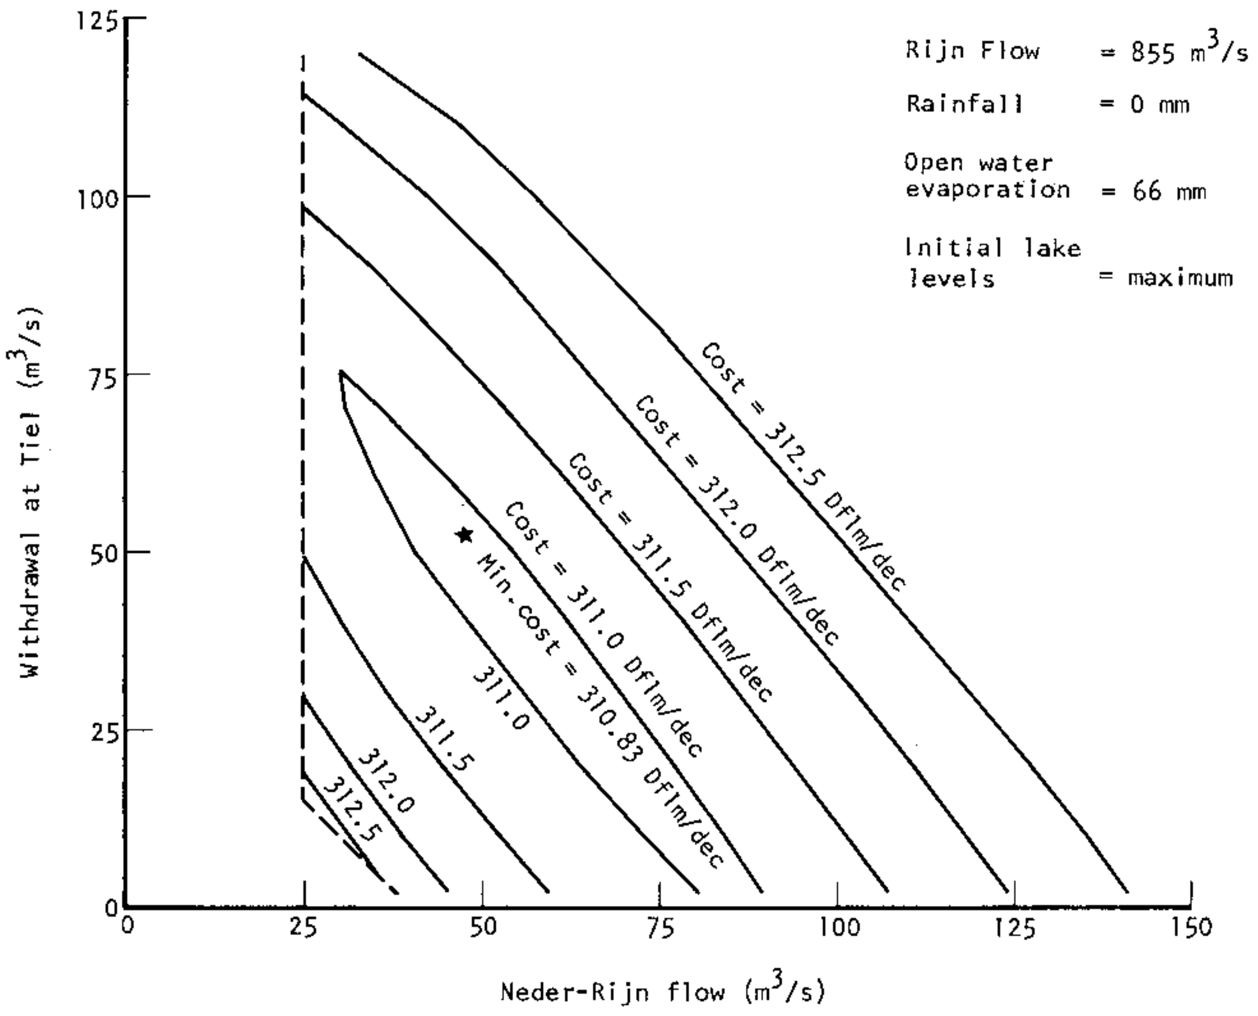

Fig. 9.2--Contours of equal total cost for variations in Neder-Rijn flow and withdrawals at Tiel

Neder-Rijn flow causes shipping costs to increase on both the Waal and the IJssel, while a reduction in the withdrawal at Tiel causes shipping losses on the Waal to decrease. Changing the flows in the proportions given above, therefore, results in partially offsetting changes in shipping low water losses.

9.2.2.2. The Amsterdam-Rijnkanaal and the Noordzeekanaal. There are extractions from and discharges into the Amsterdam-Rijnkanaal all along its length, but under the conditions that define this case, the extractions exceed the discharges. In order to maintain a minimum flow of 5 m^3/s on all sections of the canal, the MSDM strategy sends 9.49 m^3/s north from the Neder-Rijn, and routes 7.6 m^3/s from the IJssel lakes via the Vecht into the canal at Diemen (links 68 VECHT2 and 69 ARKVECHT).

In the MSDM strategy, the flow into the North Sea from the Noordzeekanaal is only 20 m^3/s. This flow is made up of a flow of 5 m^3/s from the Amsterdam-Rijnkanaal, plus a flow from the IJssel lakes through the Oranjesluis of 6.36 m^3/s, plus additions from North Holland at Zaandam and from the Midwest at Spaarndam and Halfweg totaling 10 m^3/s. Evaporation from the Noordzeekanaal reduces the total to 20 m^3/s by the time it is discharged into the North Sea.

The RWS strategy sends 40 m^3/s along the Noordzeekanaal, in order to provide cooling water for the Velsen power plant at IJmuiden. (MSDM does not do so because, in this case, we have ignored the thermal standards entirely, and hence prefer to save the water in the IJssel lakes for possible future uses, instead of dumping it into the North Sea.) Half of this water is extracted from the Neder-Rijn and brought north on the Amsterdam-Rijnkanaal. The rule observed by the strategy is that 20 m^3/s must flow from the Amsterdam-Rijnkanaal into the Noordzeekanaal; to accomplish this, 32.1 m^3/s must be taken from the Neder-Rijn. To this flow is added 10 m^3/s from North Holland and the Midwest, and 11.36 m^3/s from the IJssel lakes via the Oranjesluis. Evaporation reduces the total to 40 m^3/s by the time it is discharged into the North Sea.

9.2.2.3. The IJssel Lakes. The RWS strategy's extra demand for water from the IJssel lakes is more than offset by the increased supply from the IJssel River. Thus the IJssel lake level drops by 6.30 cm under the RWS strategy, and 6.63 cm under the MSDM strategy. Note that most of this drop in lake level is due to the extremely high evaporative demand (66 mm) during the decade.

9.2.2.4. The Maas. The MSDM strategy drops the water level in the first weir pond of the Maas (node 27 MAASLOO) by 87.36 cm, reducing the shipping depth to the same as that found in the Waal. No additional low water losses are incurred by shipping, but the flow in the Julianakanaal (link 10 JULCANL1) is augmented by 2.92 m^3/s, to a total of 3.58 m^3/s. The higher flow reduces the need for, and cost of, water conservation measures (see Chap. 7).

The RWS strategy, by contrast, seeks to maintain the weir ponds at their maximum levels. By using all available water conservation techniques at the Julianakanaal locks, the flow can be reduced to 2.5 m^3/s. This raises the cost due to shipping delays by 0.03 Dflm over the cost in the MSDM strategy, but it makes it possible to maintain a higher water level at MAASLOO. In the case under discussion, its level need drop by only 54.99 cm. This is not necessarily the wrong decision, since we have included only a nominal, very small term in the objective function of MSDM to reflect the future value of water saved in the Maas during the present decade.

9.2.3. Strategies in the Regions

In this section, we describe the MSDM and RWS strategies in the regional parts of the MSDM network. Again, it will help the reader to refer frequently to Fig. 9.1.

9.2.3.1. The Southeast Highlands and Delta Region. As shown in Table 9.4, the MSDM and RWS strategies are identical in this region. Both provide all the water demanded for sprinkling, plus extra water to help shipping on the canals in the region (the Zuidwillems Vaart and the Wilhelminakanaal). Thus, the Southeast Highlands and Delta region extracts a total of 4.84 m^3/s from nodes 27 MAASLOO and 28 BORN+PAN. Much of this water is consumed at node 34 WEER+MEY; only 1.79 m^3/s is allowed to flow farther along the Zuid-Willemsvaart. The Zoommeer (which, the reader will recall, is a freshwater lake in this case) provides an additional 1.49 m^3/s. Not all of this supply is consumed in the region; 2.68 m^3/s flows back to the lower reaches of the Maas.

9.2.3.1. The Northeast Highlands Region. Table 9.5 shows that MSDM and RWS strategies are identical in the Northeast Highlands region. However, not all demands are met; farmers must reduce their sprinkling by 0.40 m^3/s from their demand of 5.09 m^3/s. The entire reduction occurs at node 4 NEHIGH, where the demand for sprinkling is 4.13 m^3/s. The present infrastructure can deliver only 4.05 m^3/s to NEHIGH, and groundwater-fed springs supply only 0.08 m^3/s. Of this amount, 0.2 m^3/s leaks unavoidably through locks on link 53 NOWILKAN, and another 0.2 m^3/s is lost via link 54 OMMERKAN.

In spite of the need to cut back sprinkling, the region as a whole supplies water to the national network. Water flows into the Overijsselsche Vecht from groundwater and from leakage through locks on the link 54 OMMERKAN. The Overijsselsche Vecht then delivers the water to the IJssel River at node 9 IJSLEND, more water than the region extracts from the IJssel. This is clear evidence that the water supplies in this region are not well distributed.

Better distribution of this water could reduce drought damage plus sprinkling costs, and eliminate salt damage in the region. Eliminating sprinkling cutbacks would reduce drought damage plus sprinkling cost. Essentially the entire damage from salt occurs at node 4 NEHIGH, because the water delivered there is highly saline Rijn water (from the IJssel). This damage would be eliminated if the Rijn water were replaced by the much less saline Overijsselsche Vecht water or the groundwater that presently drains from the region. The total potential gain in this decade is no more than 0.5 Dflm. Because this decade is extremely dry, and hence gives rise to unusually large demands for water, the potential gain in other decades is smaller.

9.2.3.3. The Northern Region. As shown in Table 9.6, the two strategies differ slightly in the Northern region. The MSDM strategy takes 62.49 m^3/s from the IJssel lakes (node 10 IJSLAKES), and uses it to supply all of the water demanded in the region. In addition, a total of six m^3/s is discharged into the Waddenzee via links 46 FRIEHARL and 51 EEMSKAN. The discharge via EEMSKAN is at its minimum. The discharge via FRIEHARL is larger than the minimum necessary, in order to reduce local salt damage (see Chap. 4).

The RWS strategy differs from the MSDM strategy only by taking an extra 4.5 m^3/s from the IJssel lakes and sending it past node 2 B'GUMMER into the Waddenzee. Links 44 MARGKAN1, 45 MARGKAN2, and 48 B'GMSINK comprise the route used. This extra flow cools B'GUMMER, where the Bergum power plant discharges its heat, to a maximum excess temperature of seven degrees.

9.2.3.4. The North Holland Region. Table 9.7 shows the two strategies to be identical in the North Holland region. Both extract 30 m^3/s from the IJssel lakes, and supply all the water demanded in the region. Of the excess water, the minimum of 7 m^3/s is discharged via link 42 ZAAN into the Noordzeekanaal, and 2.87 m^3/s is discharged via link 43 NOHOLKAN into the North Sea. The extraction of 30 m^3/s from the IJssel lakes is the maximum allowed; if possible, MSDM would extract more to send out link 43 NOHOLKAN, and further reduce local salt intrusion (see Chap. 4).

9.2.3.5. The Midwest and Utrecht Region. Table 9.8 compares the MSDM and RWS strategies in the Midwest and Utrecht region. Of all the regions, this is the only one in which the two strategies differ. The MSDM strategy finds it optimal to cut back sprinkling below its demand by 3.58 m^3/s; the RWS strategy supplies the entire demand. The extra water needed to do so is supplied from the Hollandsche IJssel, 20.42 m^3/s being needed under the RWS strategy, and only 16.84 m^3/s under the MSDM strategy. MSDM reduces the extraction from the Hollandsche IJssel, even at the cost of sprinkling cutbacks, in order to reduce salt damage to glasshouse crops. Water from this source is more saline than water already stored at node 21 MIDWEST, so the more is extracted, the more saline will node 21 MIDWEST become. Of course, the difference of 3.58 m^3/s in extractions from the Hollandsche IJssel does not fully explain the difference in salt damage. Most of the difference is due to the fact that the MSDM strategy devotes more water to combating the Rotterdam salt wedge than does the RWS strategy (see Sec. 9.2.2), so that the water extracted from the Hollandsche IJssel is much less saline under the MSDM strategy.

The other difference between the two strategies in this region has been mentioned already in Sec. 9.2.2. The MSDM strategy extracts 10 m^3/s from the IJssel lakes via the Vecht (link 68 VECHT2) and discharges 7.6 m^3/s of it into the Amsterdam-Rijnkanaal at Diemen (via link 69 ARKVECHT). Under the RWS strategy, both flows are reduced by 7.6 m^3/s. But the reader will recall that this reduction necessitated a compensating increase of 7.6 m^3/s in the flow north from the Neder-Rijn, in order to maintain the flows in all links of the Amsterdam-Rijnkanaal at or above their minima. This additional extraction from the Neder-Rijn is costly to both shipping and to glasshouse cultivation in the Midwest, since it diverts water that could otherwise be used to increase the depth of the Waal and to combat the salt wedge.

9.3. MANAGERIAL STRATEGIES TO IMPROVE WATER QUALITY

Next, we defined a case using the nominal water quality standards. These are 200 mg/l for salt, 3 degrees Celsius for excess temperature, 50 microgram/l for chromium, 5 mg/l for BOD, and 0.3 mg/l for phosphate. In addition, we used less stringent external supply conditions, raising the flow in the Rijn to 1000 m^3/s and that in the Maas to 31 m^3/s. Again, 22.4 m^3/s of Maas flow was required to be sent to Belgium, so only 8.6 m^3/s remained for use in the Netherlands.

It should come as no surprise that the case described above has no feasible solution. The water quality standards for salt, chromium, BOD, and phosphate are all violated in the Rijn water that enters the Netherlands. It is impossible to meet the water quality standards everywhere if they are violated at the border. Moreover, as can be seen in Table 9.1, the standards for salt, BOD, and phosphate are also violated initially at most of the nodes with storage. In fact, the only standard it is always possible to meet is the 3-degree thermal standard.

To permit MSDM to find any solution at all, therefore, we had either to relax the standards, or to allow them to be violated. We chose the latter course, and allowed the standards to be violated, but we imposed a penalty at each node at which the standard was not met. The penalty was proportional to the amount by which the pollutant exceeded the standard. The mathematical formulation of this penalty was discussed in Chap. 3.

Replacing an inviolable standard by a violable one with a penalty on violations raises the question: How large should the penalty be? This was precisely the question we tried to avoid by imposing standards on water quality. The proper penalty should take into account the equivalent monetary harm done by the pollutant, as well as the cost of reducing its concentration; and the equivalent monetary harm is virtually impossible to determine.

But it is interesting to ask by how much can managerial tactics improve water quality. Accordingly, we defined six new cases for MSDM, one in which no penalties were assessed for violations of the water quality standards, and five other cases, each of which imposed a gigantic penalty on violations of a different water quality standard. In each case, the penalty was large enough that reducing the concentration of the penalized pollutant took precedence over all other goals. One case penalized violations of the thermal standard only, and ignored violations of the salt, chromium, BOD, and phosphate standards. A second case penalized salt violations only, a third case chromium violations only, a fourth case BOD violations only, and the fifth case phosphate violations only.

The case that penalized violations of the 50 ug/l chromium standard proved not to be interesting, since this standard is only violated in the Rijn where it enters the Netherlands. Thus managerial tactics have absolutely no influence if the chromium penalty is defined relative to this standard. Accordingly, we replaced this case with one in which the chromium standard was set to zero. This had the effect of requiring that the chromium concentrations be minimized. The reader will recall from Chap. 3 that a chromium standard of zero was in fact suggested by RIN (the Dutch Institute for Nature Management).

By observing the differences in pollutant concentrations from one of these cases to another, we were able to determine the scope for water quality improvement offered by managerial tactics.

9.3.1. Overall Strategy Costs

We used MSDM to derive strategies for each of the six cases described above. In Table 9.9 we present the overall costs for the six strategies. The Nominal MSDM Strategy has the lowest total cost, as should be expected, since no consideration was paid by this strategy to improving water quality. The cost of the Thermal case is hardly different. The thermal cost component has risen slightly, as has the "other" cost component, but overall the change is small.

The remaining four strategies, however, show significant increases in total cost[3] when compared to the nominal strategy. The Chloride strategy results in a substantial increase in the shipping cost and in drought damage plus sprinkling cost. We shall see later that the increase in drought damage plus sprinkling cost is due to the diversion of water from sprinkling to flushing. Surprisingly, the agricultural damage from salt also rises, even though the strategy is attempting to reduce chloride (i.e., salt) concentrations. The explanation for this is that the strategy has found that by allowing a small increase in the salt concentration in the Midwest, where most

Table 9.9

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: OVERALL STRATEGY COSTS

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Cost Components (Dflm/dec)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Shipping	11.57	11.57	17.10	18.24	14.09	11.57
Agriculture						
Damage from Salt	181.06	181.03	182.71	181.05	181.01	183.02
Drought Damage Plus Sprinkling Cost	36.88	36.88	45.20	55.71	41.67	39.43
Thermal	71.58	71.94	71.58	71.58	71.58	71.58
Other Costs	1.00	1.17	1.40	1.16	1.26	1.14
Total Costs	302.09	302.59	317.99	327.74	309.61	306.74

of the damage from salt occurs, it can reduce the salt concentration elsewhere by a larger amount. That is, if one tries to minimize salt damage to agriculture, one will apply different penalty weights at different nodes. The Chloride strategy assumes the same penalty weight at every node.

The Chromium and BOD strategies both increase the shipping cost and drought damage plus sprinkling cost components, leaving the other cost components essentially unchanged from the nominal strategy. The BOD strategy has a considerably smaller effect on both of these costs than the Chromium strategy. Finally, the Phosphate strategy affects only the drought damage plus sprinkling cost component. The shipping cost is the same as under the nominal strategy.

9.3.2. The Strategies in the National Network

9.3.2.1. The Waal, Neder-Rijn, and Delta. Table 9.10 shows the six strategies for the national waterways, and Tables 9.11 through 9.15 show the resulting concentrations of the different pollutants at the national nodes. Looking first at the nominal strategy in Table 9.10, we see that the weir at Driel has been set to permit only 37.65 m^3/s to flow down the Neder-Rijn. According to the Upper Rivers Hydrologic Equation (see App. A of Vol. VA), this implies that 767.91 m^3/s will flow down the Waal, and 185.36 m^3/s will flow north along the IJssel. The remaining 9.08 m^3/s is extracted from node 6 UPRIVER for local consumption.

Farther west, the minimum of 2 m^3/s is extracted from the Waal at Tiel, and 9.49 m^3/s is sent north from the Neder-Rijn along the Amsterdam-Rijnkanaal. After accounting for other extractions along the Neder-Rijn and the Lek, this implies that the flow in the Lek reaches a minimum of 5 m^3/s.

The decisions concerning the flow in the Neder-Rijn and the withdrawal at Tiel imply, according to the Lower Rivers Hydrologic Equations (see App. A of Vol. VA), that the flow in the Nieuwe Mass will be 224.19 m^3/s, that in the Oude Maas will be 474.85 m^3/s, and hence the flow in the Nieuwe Waterweg will be 699.04 m^3/s. Note that the flow in the Nieuwe Maas is well above the minimum of 207 m^3/s necessary to prevent the Rotterdam salt wedge from penetrating to Gouda.

Among the other five cases, three show essentially the same pattern as the nominal strategy. These are the Thermal, BOD, and Phosphate strategies. The remaining two cases (Chloride and Chromium), however, are quite different. In both, the weir at Driel is fully open, allowing the maximum of 187 m^3/s to flow along the Neder-Rijn. This results in 681.29 m^3/s flowing on the Waal, and only 122.63 m^3/s flowing north on the IJssel.

Farther west, at the mouth of the Amsterdam-Rijnkanaal (node 12 A-R.MOND), even these two cases are different. The Chloride strategy

Table 9.10

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: STRATEGIES FOR THE NATIONAL WATERWAYS

Case Description

Rijn Flow: 1000.00 m^3/s	Infrastructure: 1976 + Fresh Zoommeer
Maas Flow: 31.00 m^3/s	Power Plant Inventory: 1985
Net Evaporation: 66.00 mm	Sprinkling Scenario: 1976
Initial Lake Levels: Maximum	Shipping Scenario: 1976

Strategy Description	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Upper Rivers Flows (m^3/s)						
Waal	767.91	767.91	681.29	681.29	767.91	768.34
Neder-Rijn	37.65	37.65	187.00	187.00	37.65	36.90
IJssel	185.36	185.36	122.63	122.63	185.36	185.68
Lower Rivers Flows (m^3/s)						
Nieuwe Maas	224.19	228.46	318.88	400.55	237.98	210.81
Oude Maas	474.85	482.42	373.84	405.73	499.27	473.84
Nieuwe Waterweg	699.04	710.88	692.72	806.28	737.25	684.55
Volkerak + Haringvlietsluizen	40.11	40.11	38.61	40.11	40.11	40.11
Amsterdam-Rijnkanaal Flows (m^3/s)						
Withdrawal at Tiel	2.00	2.00	120.00	120.00	2.00	2.00
North from A-R Mond	9.49	9.49	81.86	9.49	9.49	8.74
Lek	5.00	5.00	209.25	272.35	5.00	5.00
Noordzeekanaal Flows (m^3/s)						
Input from IJssel Lakes	6.36	111.36	136.00	5.89	6.36	0.0
Discharge to North Sea	20.00	125.00	230.00	20.00	20.00	29.14
IJssel Lakes Level Change (cm)	-4.74	-10.24	-14.26	-6.48	-6.14	-5.93
Maas Weir Pond Level Changes (cm)						
MAASLOO	-47.51	-47.51	-102.47	-130.00	-130.00	0.0
BORN+PAN	0.0	-47.51	0.0	-130.00	-130.00	0.0
LINNE	0.0	-47.51	0.0	-130.00	-130.00	0.0
ROER+BEL	0.0	-47.51	0.0	-130.00	-130.00	0.0
SAMGRALI	0.0	-47.51	0.0	-130.00	-130.00	0.0
Maas Flows (m^3/s)						
JULCANL1	6.65	6.65	8.49	6.25	6.25	6.15
MAAS3	14.19	16.02	16.03	22.67	22.67	13.29
AMER1 (Discharge into Lower Rivers)	18.06	29.90	18.41	56.27	56.27	19.51

Table 9.11

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: EXCESS TEMPERATURES AT SELECTED MSDM NODES

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Excess Temperatures (Deg Celsius)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
2. B'GUMMER	5.13	3.00	4.61	5.59	5.13	4.61
13. UTR+MAAR	10.95	10.95(a)	2.62	11.60	10.95	12.00
15. AMSTEDAM	13.68	1.25	1.01	13.68	13.68	7.33
16. HAL+IJMU	17.10	3.48(a)	2.25	17.10	17.10	12.13
29. LINNE	12.10	11.90(a)	11.90	10.00	10.00	12.18
33. GERTRUID	7.86	6.30(a)	7.80	4.59	4.59	7.70

(a) These excess temperatures were calculated, as were all the others, assuming that the optimal thermally unrestricted generating schedule were being used (see Chap. 5). To reduce them to three degrees, as is required by the standards, the generating schedule was changed, using the linear program formulated in Chap. 5. The increased cost of the new generating schedule is reflected in the thermal cost component shown in Table 9.9.

Table 9.12

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: CHLORIDE CONCENTRATIONS IN NATIONAL WATERWAYS

Case Description

Rijn Flow: 1000.00 m^3/s	Infrastructure: 1976 + Fresh Zoommeer
Maas Flow: 31.00 m^3/s	Power Plant Inventory: 1985
Net Evaporation: 66.00 mm	Sprinkling Scenario: 1976
Initial Lake Levels: Maximum	Shipping Scenario: 1976

Chloride Concentrations (mg/l)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Waal, Neder-Rijn, and Delta						
6. UPRIVER	311.12	311.12	311.12	311.12	311.12	311.12
12. A-R.MOND	313.77	313.77	313.77	313.77	313.77	313.77
19. IJSLMOND	556.58	544.69	376.60	333.78	518.17	593.86
22. NIJMEGEN	311.07	311.07	311.07	311.07	311.07	311.07
23. TIEL	311.64	311.64	311.64	311.64	311.64	311.64
24. GOR+DOR	312.48	312.48	312.48	312.48	312.48	312.48
25. LOWRIVER	257.51	256.25	254.35	251.77	254.73	257.41
26. DLTALAKE	329.07	329.07	328.91	329.06	329.07	329.07
Amsterdam-Rijnkanaal and Noordzeekanaal						
13. UTR+MAAR	324.50	324.50	324.50	324.50	324.50	324.50
14. DIEMEN	240.56	240.56	311.48	233.15	240.56	232.38
15. AMSTEDAM	1356.04	684.84	499.49	1355.73	1356.04	1329.46
16. HAL+IJMU	2210.33	2002.33	1468.35	2210.54	2210.33	2201.45
IJssel and IJssel Lakes						
7. TWENMOND	309.24	309.24	308.23	308.23	309.24	309.24
9. IJSLEND	301.58	301.58	296.68	296.68	301.58	301.01
10. IJSLAKES	222.54	222.54	222.54	222.54	222.54	222.54
Maas						
27. MAASLOO	67.79	67.79	68.25	68.47	68.47	67.37
28. BORN+PAN	73.71	73.13	73.79	71.85	71.85	73.76
29. LINNE	80.30	79.35	79.56	76.84	76.84	80.65
30. ROER+BEL	105.15	104.95	104.08	102.28	102.28	106.06
31. SAMGRALI	91.84	92.47	91.85	93.00	93.00	91.91
32. DBOS+BOX	95.45	95.05	95.35	94.73	94.73	95.83
33. GERTRUID	94.35	94.31	92.93	93.45	93.45	92.76

Table 9.13

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: CHROMIUM CONCENTRATIONS IN NATIONAL WATERWAYS

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Chromium Concentrations (ug/l)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Waal, Neder-Rijn, and Delta						
6. UPRIVER	55.69	55.69	55.69	55.69	55.69	55.69
12. A-R.MOND	46.36	46.36	46.16	46.16	46.77	46.16
19. IJSLMOND	38.82	38.82	38.75	38.73	38.84	38.75
22. NIJMEGEN	52.60	52.60	52.60	52.60	52.60	52.60
23. TIEL	48.16	48.16	48.16	48.16	48.16	48.16
24. GOR+DOR	43.64	43.64	43.64	43.64	43.64	43.64
25. LOWRIVER	12.98	12.91	11.78	11.56	12.75	11.96
26. DLTALAKE	0.17	0.17	0.17	0.15	0.16	0.15
Amsterdam-Rijnkanaal and Noordzeekanaal						
13. UTR+MAAR	17.82	17.82	41.63	17.74	17.82	17.69
14. DIEMEN	3.25	3.25	34.65	3.02	3.25	8.07
15. AMSTEDAM	2.22	3.09	10.18	2.23	2.23	2.46
16. HAL+IJMU	2.94	3.26	6.85	2.92	2.95	2.77
IJssel and IJssel Lakes						
7. TWENMOND	50.86	50.86	50.86	50.68	50.98	50.70
9. IJSLEND	43.93	43.93	43.34	43.34	44.25	43.35
10. IJSLAKES	4.34	4.34	4.34	4.27	4.42	4.27
Maas						
27. MAASLOO	2.40	2.40	2.34	2.32	2.32	2.40
28. BORN+PAN	0.54	0.56	0.50	0.68	0.68	0.54
29. LINNE	0.30	0.33	0.29	0.46	0.46	0.30
30. ROER+BEL	0.06	0.07	0.07	0.13	0.13	0.06
31. SAMGRALI	0.01	0.01	0.01	0.03	0.03	0.01
32. DBOS+BOX	1.20	0.88	1.13	0.60	1.60	1.20
33. GERTRUID	2.20	1.64	2.23	1.27	1.27	1.75

Table 9.14

MSDM STRATEGIES FOR IMPROVING WATER QUALITY:
BOD CONCENTRATIONS IN NATIONAL WATERWAYS

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

BOD Concentrations (mg/l)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Waal, Neder-Rijn, and Delta						
6. UPRIVER	14.66	14.66	14.66	14.66	14.66	14.66
12. A-R.MOND	11.84	11.84	11.96	11.96	11.84	11.84
19. IJSLMOND	14.29	14.20	13.53	13.16	14.29	15.71
22. NIJMEGEN	13.75	13.75	13.67	13.67	13.75	13.75
23. TIEL	11.36	11.36	11.05	11.05	11.36	11.36
24. GOR+DOR	8.17	8.17	7.06	7.06	8.17	8.17
25. LOWRIVER	3.42	3.42	3.32	3.32	3.42	3.42
26. DLTALAKE	0.04	0.04	0.04	0.04	0.04	0.04
Amsterdam-Rijnkanaal and Noordzeekanaal						
13. UTR+MAAR	5.62	5.62	10.70	5.45	5.62	5.40
14. DIEMEN	5.54	5.54	9.61	5.67	5.54	6.60
15. AMSTFDAM	3.47	3.42	4.64	3.47	3.47	3.50
16. HAL+IJMU	4.48	3.83	4.36	4.50	4.48	4.55
IJssel and IJssel Lakes						
7. TWENMOND	13.79	13.79	13.70	13.70	13.79	13.79
9. IJSLEND	11.03	11.03	10.27	10.27	11.03	11.03
10. IJSLAKES	3.58	3.58	3.58	3.58	3.58	3.58
Maas						
27. MAASLOO	7.94	7.94	8.07	8.13	8.13	7.94
28. BORN+PAN	6.84	6.77	6.81	7.31	7.31	6.84
29. LINNE	12.09	12.14	11.72	11.87	11.87	12.09
30. ROER+BEL	7.80	7.83	7.44	7.58	7.58	7.80
31. SAMGRALI	5.80	5.81	5.55	5.60	5.60	5.80
32. DBOS+BOX	28.61	22.51	27.01	16.48	16.48	28.61
33. GERTRUID	17.50	16.15	16.29	13.30	13.30	17.50

Table 9.15

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: PHOSPHATE CONCENTRATIONS IN NATIONAL WATERWAYS

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Phosphate Concentrations (mg/l)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l 5 mg/l Standard	0.3 mg/l Phosphate Standard
Waal, Neder-Rijn, and Delta						
6. UPRIVER	1.04	1.04	1.04	1.04	1.04	1.04
12. A-R.MOND	0.60	0.60	0.83	0.83	0.60	0.60
19. IJSLMOND	0.78	0.77	0.71	0.74	0.77	0.78
22. NIJMEGEN	0.95	0.95	0.99	0.99	0.95	0.95
23. TIEL	0.85	0.85	0.88	0.88	0.85	0.85
24. GOR+DOR	0.72	0.72	0.68	0.68	0.72	0.72
25. LOWRIVER	0.36	0.37	0.35	0.35	0.37	0.36
26. DLTALAKE	0.0	0.0	0.0	0.0	0.0	0.0
Amsterdam-Rijnkanaal and Noordzeekanaal						
13. UTR+MAAR	0.50	0.50	0.77	0.59	0.50	0.50
14. DIEMEN	1.40	1.40	0.83	1.50	1.40	1.48
15. AMSTEDAM	0.28	0.29	0.40	0.28	0.28	0.28
16. HAL+IJMU	0.36	0.31	0.36	0.36	0.36	0.36
IJssel and IJssel Lakes						
7. TWENMOND	1.07	1.07	1.10	1.10	1.07	1.07
9. IJSLEND	1.11	1.11	1.23	1.23	1.11	1.11
10. IJSLAKES	0.33	0.33	0.33	0.33	0.33	0.33
Maas						
27. MAASLOO	1.49	1.49	1.48	1.47	1.47	1.50
28. BORN+PAN	0.89	0.92	0.90	1.00	1.00	0.89
29. LINNE	1.02	1.05	1.02	1.10	1.10	1.02
30. ROER+BEL	1.10	1.07	1.10	1.18	1.18	1.10
31. SAMGRALI	0.66	0.70	0.66	0.78	0.78	0.66
32. DBOS+BOX	2.42	1.97	2.30	1.58	1.58	1.58
33. GERTRUID	1.59	1.48	1.53	1.25	1.25	1.25

diverts 81.86 m^3/s north from the Neder-Rijn, while the Chromium strategy diverts only the minimum flow of 9.49 m^3/s. In both cases, however, the diversion leaves large amounts of water to flow west in the Lek, 209.25 m^3/s in the Chloride case, and 272.35 m^3/s in the Chromium case.

The extra water in the Lek in the Chloride and Chromium cases results in much larger flows in the Nieuwe Maas than under the nominal strategy. In the Chloride case, the Nieuwe Maas flow is 318.88 m^3/s, while in the Chromium case it is 400.55 m^3/s. Flows this large have no effect on the effect of the Rotterdam salt wedge at Gouda, but they do reduce its effect at IJsselmonde (node 19 IJSLMOND; see App. D of Vol. VA). Indeed, therein lies part of the explanation for setting the weir at Driel fully open in the Chloride case (see Table 9.12). The explanation in the Chromium case is different--it is to prevent chromium, whose major source is the Rijn, from flowing along the IJssel and into the IJssel lakes.

<u>9.3.2.2. The Amsterdam-Rijnkanaal and the Noordzeekanaal</u>. The six strategies all manage the Amsterdam-Rijnkanaal and the Noordzeekanaal quite differently. The nominal strategy sends only 20 m^3/s into the North Sea from the Noordzeekanaal, made up from the same sources as described for the MSDM strategy in Sec. 9.2.2.2 above.

The Thermal strategy for the Amsterdam-Rijnkanaal and the Noordzeekanaal differs from the nominal strategy only in that it sends additional water through the Oranjesluis (link 38 ORANJESL) and along the Noordzeekanaal. The total discharge into the North Sea is 125 m^3/s, which we determined in Chap. 5 was the maximum amount that could be profitably used to cool the power plants at Amsterdam and IJmuiden. This does not mean that no power plants on the Amsterdam-Rijnkanaal and the Noordzeekanaal are constrained by the thermal standard. To the contrary, if the heat discharge at node 13 UTR+MAAR and either node 15 AMSTEDAM or 16 HAL+IJMU is not cut back, Table 9.11 shows that the three-degree standard will be violated at both nodes 13 UTR+MAAR and 16 HAL+IJMU. See Chap. 5 for further details.

The Chloride strategy is perhaps the most different from the nominal strategy. Remember that 81.86 m^3/s was sent north from the Neder-Rijn along the Amsterdam-Rijnkanaal. That amount was determined to be the amount necessary to increase the flow in the Noordzeekanaal to its maximum of 230 m^3/s, when added to the maximum discharge of link 38 ORANJESL of 136 m^3/s, plus the various smaller discharges from North Holland and the Midwest. The Chloride strategy maximizes the flow in the Noordzeekanaal in order to reduce salt intrusion through the locks at IJmuiden to a minimum. As can be seen in Table 9.12, no other strategy reduces the chloride concentrations at nodes 15 AMSTEDAM and 16 HAL+IJMU as much as the Chloride strategy.

The Chromium strategy is like the nominal strategy, in that it minimizes the water discharged via the Noordzeekanaal into the North

Sea. The reasons are quite different, however. The nominal strategy has a mild preference for conserving water, and no reason for sending it along the Noordzeekanaal (remember that the nominal case has no thermal penalty). By contrast, the Chromium strategy wants to keep the flows in the Amsterdam-Rijnkanaal and the Noordzeekanaal at their minimum levels, to give the chromium the maximum possible opportunity to "decay" (see Chap. 3, and App. C of Vol. VA). (For chromium, we interpret "decay" to be sedimentation.) From Table 9.13 we see the difference in the chromium concentrations on the Amsterdam-Rijnkanaal and the Noordzeekanaal between a case with high canal flows (the Chloride strategy) and one with low flows (the Chromium strategy).

The BOD strategy treats the Amsterdam-Rijnkanaal and the Noordzeekanaal in the same manner as does the nominal strategy. By flushing the Noordzeekanaal with water from the IJssel lakes, it would be possible to reduce the BOD concentrations at nodes 15 AMSTEDAM and 16 HAL+IJMU below those shown in Table 9.14. But because the concentrations are already within the 5 mg/l standard for BOD (two of the very few nodes at which this standard is not violated), the BOD strategy finds no advantage in doing so.

The different flow in the Noordzeekanaal between the Phosphate and nominal strategies is dictated by differences in the strategies in the Midwest and Utrecht region, and will be discussed there. Here we will only say that the Phosphate strategy would prefer to minimize the flow in the Noordzeekanaal and the Amsterdam-Rijnkanaal in order to allow the phosphate as much time as possible to "decay."

9.3.2.3. The IJssel Lakes. The six strategies deliver different amounts of water to the IJssel lakes via the IJssel River, and extract different amounts from it. In consequence, the lake level drops by different amounts. The nominal strategy best conserves the lake level, allowing a drop of only 4.74 cm. The Chloride strategy does worst in this respect, resulting in a drop of 14.26 cm. The other strategies fall between these extremes, as shown in Table 9.10.

9.3.2.4. The Maas. The strategies are almost all quite different on the Maas. The nominal strategy allows the water level in the first weir pond (node 27 MAASLOO) to drop until the minimum depth on the Maas equals the depth on the Waal. This increases the flow on the Julianakanaal to 6.65 m^3/s, and reduces the cost to shipping below the value it would be if the water level had been maintained (as is the present practice).

The Thermal strategy is not content with depleting only the first weir pond; it depletes all five, so that the depth of each becomes equal to the Waal depth. As explained in Chap. 5, this provides extra cooling water for the Amer power plant, located at node 33 GERTRUID, and the Maasbracht power plant, located at node 29 LINNE. Even so, as is shown in Table 9.11, the heat discharges at both nodes must be cut back if the 3-degree thermal standard is to be met. Incidentally, there is some question whether depleting the weir ponds to provide extra cooling is more valuable than saving the water for possible future needs. The value we have placed on water stored in the weir ponds is very small; it corresponds to the assumption that current uses are all more valuable than possible future uses, on the principle that "a bird in the hand is worth any number in the bush." Earlier we suggested that an investigation of this question might prove worthwhile.

As can be seen in Table 9.12, the chloride concentrations at nodes on the Maas are all below the standard of 200 mg/l. Thus, like the nominal strategy, the Chloride strategy allows the depth of the first weir pond to decrease to the Waal depth. The difference between the two cases arises because, owing to the different strategies regarding the weir at Driel and withdrawals at Tiel, the Waal is shallower under the Chloride strategy than under the nominal strategy.

Both the Chromium and BOD strategies deplete all the weir ponds to their minimum levels. In the case of the Chromium strategy, the purpose of this tactic is to provide as much Maas water as possible to dilute the water stored at node 25 LOWRIVER (see Table 9.13). The BOD strategy does so in order to provide a maximum amount of water to dilute the relatively large BOD discharges into node 32 DBOS+BOX (see Table 9.14). This also reduces the BOD concentration at all nodes downstream of DBOS+BOX, notably node 33 GERTRUID.

Finally, the Phosphate strategy chooses to maintain all the weir ponds at their maximum levels. The reason here is that the Maas has a higher phosphate concentration than node 25 LOWRIVER, and MSDM prefers to minimize the contamination of the latter node by the former source of water. Note that the reasoning is exactly the reverse of that used in describing the Chromium strategy, where a maximum amount of Maas water was used to dilute water at node 25 LOWRIVER, because the chromium concentration in the Maas was lower than at LOWRIVER.

9.3.3. Water Quality in the National Waterways

We now wish to ask how effectively the six strategies just described improve the quality of the water in the national waterways. This question is answered by referring to Tables 9.11 through 9.15.

9.3.3.1. Thermal Pollution. Table 9.11 deals with thermal pollution. The nodes shown in this table are those we identified as "hot" nodes in Chap. 5. Only at these nodes will the excess temperature ever exceed three degrees. One of these nodes, 2 B'GUMMER, is actually a regional node, and will be discussed in a later section. Of the remaining nodes, three (13 UTR+MAAR, 15 AMSTEDAM, and 16 HAL+IJMU) are on the Amsterdam-Rijnkanaal or the Noordzeekanaal. At these three nodes, the large range of excess temperatures demonstrates that managerial tactics can be very effective at reducing thermal pollution. However, the fact that the Thermal strategy does not have the lowest excess temperatures suggests that the alternative method for reducing thermal pollution--i.e., shifting the electric generating load to power plants at other nodes--is sometimes more economical than employing managerial tactics.

At the two "hot" nodes on the Maas, 29 LINNE and 33 GERTRUID, a similar story can be told. By depleting the weir ponds, extra cooling water can be provided at these nodes. However, the Thermal strategy only depletes the weir ponds to the point that their depths are equal to the depth of the Waal. Further depletion is possible, as in the Chromium and BOD strategies, but it results in a higher cost than shifting the load. The situation on the Maas does differ from that on the Amsterdam-Rijnkanaal and Noordzeekanaal in that the maximum possible effect of managerial tactics on excess temperatures is much smaller.

9.3.3.2. Chloride. Table 9.12 shows chloride concentrations at the national nodes. The only nodes where chloride concentrations differ significantly among strategies are 19 IJSLMOND, 14 DIEMEN, 15 AMSTEDAM, and 16 HAL+IJMU. Different Nieuwe Maas flows, on which the position of the Rotterdam salt wedge depends, accounts for the differences in chloride concentration at IJSLMOND. Different chloride concentrations at AMSTEDAM and HAL+IJMU are caused by different flows in the Noordzeekanaal, which governs the amount of salt intrusion at IJmuiden.

The differences at DIEMEN are due to the fact that some of the flow past DIEMEN is from node 18 VECHT in the Midwest and Utrecht region. Water stored there has a relatively low chloride concentration. The remainder of the flow past DIEMEN is composed of more saline Rijn water sent north on the Amsterdam-Rijnkanaal from the Neder-Rijn. In the Chloride strategy, in which a large amount of water is sent north from the Neder-Rijn, the diluting effect of water from VECHT is overwhelmed, and the chloride concentration at DIEMEN is high. Thus the high chloride concentration at DIEMEN that results from the Chloride strategy is an undesirable side effect of the attempt to reduce the concentrations at AMSTEDAM and HAL+IJMU.

We may legitimately ask whether the high chloride concentrations at AMSTEDAM and HAL+IJMU are really worth reducing. The nominal chloride standard for PAWN requires that the chloride concentration be less than 200 mg/l, but as explained in Chap. 4, alternative standards exist. Standards suggested for the preservation of nature (by RIN, the Dutch Institute for Nature Management) would require the chloride concentrations in the Noordzeekanaal (i.e., at AMSTEDAM and HAL+IJMU) to exceed 500 mg/l.

9.3.3.3. Chromium. Table 9.13 shows the chromium concentrations at national nodes. It appears that managerial tactics have a significant effect on chromium concentrations only at nodes on the Amsterdam-Rijnkanaal and Noordzeekanaal. The large flow of water in the canals under the Chloride strategy reduces the travel time of the water, and hence reduces the amount of sedimentation of chromium that occurs.

9.3.3.4. BOD. Table 9.14 shows the concentrations of BOD at the national nodes. Once again managerial tactics have a significant effect on BOD concentrations at some of the nodes on the Amsterdam-Rijkanaal and the Noordzeekanaal. The reason is the same as that given for chromium; the longer the transit time, the more BOD will decay. The Chloride strategy, with its high flows in these canals and consequent short transit time, has correspondingly higher concentrations. There is also significant dilution of BOD discharges DBOS+BOX and 33 GERTRUID, both on the Maas, under the strategies that release the most water from the weir ponds.

9.3.3.5. Phosphate. Table 9.15 presents phosphate concentrations at the national nodes. There are many nodes at which managerial tactics appear to have a significant effect, once again including the nodes on the Amsterdam-Rijnkanaal and Noordzeekanaal. The increased flow and decreased transit time on these canals under the Chloride strategy again has its detrimental side effect, that of reducing phosphate losses due to "decay." It has a beneficial effect at DIEMEN, however, since it provides extra Rijn water (which is low in phosphate) to mix with water from node 18 VECHT (which is higher in phosphate). Recall that the chloride concentrations in these two water sources are in the reverse relationship, so that this same phenomenon has a detrimental result in the case of chloride.

9.3.4. Strategies in the Regions

9.3.4.1. The Southeast Highlands and Delta Region. Table 9.16 shows the six strategies in the Southeast Highlands and Delta region, as well as their effect on the four pollutants chloride, chromium, BOD, and phosphate. Thermal pollution does not play a role in this region, so the nominal and Thermal strategies are identical. They are, in fact, the same as the MSDM strategy described in Sec. 9.2.3.1 above. The other four strategies all cut back sprinkling, the Chloride strategy entirely and the Chromium, BOD, and Phosphate strategies almost entirely. In every case, water that would have been used for sprinkling is instead used for diluting pollutants. The Chloride strategy also differs from the others in that it extracts no water from the Delta Lake, which has a high chloride concentration. The Delta Lake is the source of water for node 37 FIJNAART in all other strategies; in the Chloride strategy, the water comes instead from node 36 OSTRHOUT via link 78 MARK1.

The effect of the strategies on pollutant concentrations in this region is relatively coherent. The Chloride strategy does indeed reduce chloride concentrations, especially at node 37 FIJNAART. This

Table 9.16

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: SOUTHEAST HIGHLANDS AND DELTA REGION

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Strategy Description (m^3/s)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Water for Sprinkling	3.63	3.63	0.0	0.60	0.60	0.60
Other Consumption	0.33	0.33	0.33	0.33	0.33	0.33
Extraction from Upper Maas	4.84	4.84	4.15	4.15	4.15	4.15
Extraction from Delta Lake	1.49	1.49	0.0	1.49	1.49	1.49
Discharge to Lower Maas	2.38	2.38	3.82	4.72	4.72	4.72
Target Sprinkling Water	3.63	3.63	3.63	3.63	3.63	3.63
Cutbacks in Sprinkling	0.0	0.0	3.63	3.03	3.03	3.03
Regional Agricultural Damage (Dflm)						
Damage from Salt	0.0	0.0	0.0	0.0	0.0	0.0
Drought Damage Plus Sprinkling Cost	2.64	2.64	5.53	5.19	5.19	5.19
Pollutant Concentrations						
Chloride (mg/l)						
34. WEER+MEY	72.48	72.24	70.00	73.07	73.07	73.07
35. HELMOND	79.59	79.33	72.92	76.11	76.11	76.11
36. OSTRHOUT	66.47	66.35	69.29	71.83	71.83	71.83
37. FIJNAART	213.05	213.05	65.20	213.05	213.05	213.05
Chromium (ug/l)						
34. WEER+MEY	0.92	0.92	1.72	0.65	0.65	0.65
35. HELMOND	0.50	0.50	1.29	0.49	0.49	0.49
36. OSTRHOUT	27.99	27.99	11.32	10.88	10.88	10.88
37. FIJNAART	0.06	0.06	0.48	0.06	0.06	0.06
BOD (mg/l)						
34. WEER+MEY	7.72	7.70	8.28	7.94	7.94	7.94
35. HELMOND	8.49	8.46	8.63	8.27	8.27	8.27
36. OSTRHOUT	27.48	27.48	15.80	15.51	15.51	15.51
37. FIJNAART	1.77	1.77	9.27	1.77	1.77	1.77
Phosphate (mg/l)						
34. WEER+MEY	1.28	1.29	1.52	1.24	1.24	1.24
35. HELMOND	1.41	1.42	1.58	1.29	1.29	1.29
36. OSTRHOUT	4.30	4.30	2.64	2.41	2.41	2.41
37. FIJNAART	1.21	1.21	2.73	1.21	1.21	1.21

was the only node at which the standard is violated in the nominal case. It also has a beneficial effect on chromium, BOD, and phosphate at OSTRHOUT. But it increases the same three concentrations at FIJNAART.

The other three strategies are identical in this region, and hence have identical effects on pollutant concentrations. Chloride concentrations are essentially the same as in the nominal case. The chromium and BOD concentrations are reduced at OSTRHOUT, but hardly affected elsewhere. The phosphate concentration is reduced significantly at OSTRHOUT, and its reduction at HELMOND might also be barely significant. In this region, therefore, the Chromium/BOD/Phosphate strategy simultaneously improves all water quality parameters.

However, we must interject a word of caution. The MSDM network represents the actual water management infrastructure of this region in a very aggregate way. Only four nodes and eight links are used to represent a region with a multitude of streams and rivers in addition to its canals. Even the canals, which are the waterways represented in MSDM, are more complex than their representation indicates. Accordingly, the strategies in this region can only be defined, and their effects on water quality assessed, in very aggregate terms. Details that do not appear in the MSDM representation of the region might force a change in the estimate of a strategy's effectiveness. Opportunities unseen by MSDM might exist in reality. The same comments apply equally to the other regions. For these reasons, any conclusions we draw about the existence and description of strategies to improve regional water quality must be regarded as highly tentative and indicative only.

9.3.4.2. The Northeast Highlands Region. Table 9.17 shows the six strategies, and their effects on water quality, in the Northeast Highlands region. Again, thermal pollution does not occur in this region, so the Thermal strategy is the same as the nominal strategy, which is in turn the same as the MSDM strategy described in Sec. 9.2.3.2. The Chloride and Chromium strategies are the same as each other, but quite different from the nominal strategy. Instead of supplying the maximum possible amount of water to nodes 4 NEHIGH and 8 TWENTEND from the IJssel River, they send the minimum possible amount of water. The reason for this is that most of the chloride and chromium in the region has its origin in IJssel River water, which in turn originated as Rijn water. The other sources of water in the Northeast Highlands are groundwater and the Overijsselsche Vecht, both of which contain low pollutant concentrations.

The BOD strategy differs from the nominal strategy in only one respect. While it sends the maximum possible amount of water to NEHIGH from the IJssel River, it does not use it for sprinkling. Instead, the water is allowed to flow into the Northern region via link 53 NOWILKAN. Used in either way the water would dilute the BOD discharged into NEHIGH, but when it is sent on to the Northern

Table 9.17

MSDM STRATEGIES FOR IMPROVING WATER QUALITY:
NORTHEAST HIGHLANDS REGION

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Strategy Description (m^3/s)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Water for Sprinkling	4.69	4.69	0.20	0.20	0.96	4.69
Other Consumption	-4.56	-4.56	-4.56	-4.56	-4.56	-4.56
Overijsselsche Vecht	1.80	1.80	1.80	1.80	1.80	1.80
Extractions from IJssel	5.89	5.89	1.40	1.40	5.89	10.89
Discharge to IJssel	7.36	7.36	7.36	7.36	7.36	12.36
Discharge to North	0.20	0.20	0.20	0.20	3.93	0.20
Target Sprinkling Water	5.09	5.09	5.09	5.09	5.09	5.09
Cutbacks in Sprinkling	0.40	0.40	4.89	4.89	4.13	0.40
Regional Agricultural Damage (Dflm)						
Damage from Salt	0.10	0.10	0.01	0.01	0.10	0.10
Drought Damage Plus Sprinkling Cost	2.37	2.37	4.61	4.61	4.27	2.37
Pollutant Concentrations						
Chloride (mg/l)						
4. NEHIGH	244.94	244.94	107.10	107.10	244.94	244.15
5. OVIJVECH	80.46	80.46	76.79	76.79	80.46	164.06
8. TWENTEND	251.68	251.68	224.01	224.01	251.68	290.68
Chromium (ug/l)						
4. NEHIGH	6.28	6.28	1.69	1.69	6.33	6.33
5. OVIJVECH	5.69	5.69	5.02	5.02	5.70	11.25
8. TWENTEND	7.96	7.96	7.40	7.40	7.44	30.22
BOD (mg/l)						
4. NEHIGH	35.14	35.14	117.03	117.03	35.14	35.14
5. OVIJVECH	15.69	15.69	15.69	15.69	15.69	12.98
8. TWENTEND	6.27	6.27	8.48	8.48	6.27	9.85
Phosphate (mg/l)						
4. NEHIGH	6.73	6.73	59.18	59.18	6.73	6.73
5. OVIJVECH	3.51	3.51	4.89	4.89	3.51	2.62
8. TWENTEND	1.62	1.62	1.90	1.90	1.62	1.25

region, it can be used to dilute BOD discharges into node 3 GRONETAL as well.

The Phosphate strategy extracts 5 m^3/s more from the IJssel River via link 57 TWENKAN1 than does the nominal strategy. This water flows past node 8 TWENTEND, then along link 56 OVIJKAN1 to node 5 OVIJVECH, where it augments the flow in the Overijsselsche Vecht (link 55 OVIJVEC2) and thus returns to the IJssel River at node 9 IJSLEND. This water is used to dilute the phosphate discharges at TWENTEND and OVIJVECH.

The effects of these strategies on pollutant concentrations are inconsistent. By reducing the intake of IJssel River water, the Chloride and Chromium strategies do indeed reduce the chloride and chromium concentrations throughout the region. However, the BOD and phosphate concentrations rise tremendously at two of the three nodes (NEHIGH and TWENTEND). The BOD strategy has no effect on any pollutant concentration in the region (but an effect will be seen in the Northern region). The Phosphate strategy increases the chloride and chromium concentrations, especially at OVIJVECH, since more IJssel River water is delivered there. It has a mixed effect on BOD, increasing it at TWENTEND but reducing it at OVIJVECH. And it has the desired, and expected, effect on phosphate concentrations.

9.3.4.3. The Northern Region. Table 9.18 shows the six strategies and their effects on water quality in the Northern region. As in the other regions, the nominal strategy is the same as the MSDM strategy described in Sec. 9.2.3.3. Unlike the other regions, however, thermal pollution does play a role here. The Thermal strategy differs from the nominal by extracting 23.41 m^3/s more from the IJssel lakes, sending it via links 44 MARGKAN1 and 45 MARGKAN2 past node 2 B'GUMMER (site of the Bergum power plant), and thence into the North Sea via link 48 B'GMSINK. As we determined in Chap. 5, this provides just enough cooling capacity to allow the Bergum plant to operate at full capacity, without the 3-degree standard being violated at B'GUMMER.

The Chloride and Phosphate strategies are similar to the Thermal strategy, except that they increase extractions from the IJssel lakes to the maximum possible extent. They use it to dilute the water stored at node 1 FRIELAND, and then discharge it to the North Sea. It would be preferable to send more water from FRIELAND to node 3 GRONETAL (via links 45 MARGKAN2 and 49 STABOKAN), but the capacity of link 49 STABOKAN is limited to 16 m^3/s.

The BOD strategy is identical to the Chloride and Phosphate strategies except for an additional flow of 3.73 m^3/s from the Northeast Highlands region, which reduces the BOD concentration at node 3 GRONETAL.

Finally, the Chromium strategy reduces extractions from the IJssel lakes to a minimum, since that is the main source of what little

Table 9.18

MSDM STRATEGIES FOR IMPROVING WATER QUALITY:
NORTHERN REGION

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Strategy Description (m^3/s)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Water for Sprinkling	15.13	15.13	15.13	0.0	15.13	15.13
Other Consumption	41.57	41.57	41.57	41.57	41.57	41.57
Extraction from IJssel Lakes	62.49	85.90	97.00	43.37	97.00	97.00
From Northeast Highlands	0.20	0.20	0.20	0.20	3.93	0.20
Discharges to Waddenzee	6.00	29.41	40.50	3.00	44.23	40.50
Target Sprinkling Water	15.13	15.13	15.13	15.13	15.13	15.13
Cutbacks in Sprinkling	0.0	0.0	0.0	15.13	0.0	0.0
Regional Agricultural Damage (Dflm)						
Damage from Salt	1.03	1.02	1.01	1.04	1.02	1.01
Drought Damage Plus Sprinkling Cost	4.53	4.53	4.53	10.94	4.53	4.53
Pollutant Concentrations						
Excess Temperature (Degrees Celsius)						
2. B'GUMMER	5.13	3.00	4.61	5.59	5.13	4.61
Chloride (mg/l)						
1. FRIELAND	239.30	238.64	238.28	240.14	238.28	238.28
2. B'GUMMER	253.41	244.27	242.41	257.23	251.98	242.41
3. GRONETAL	259.87	250.46	246.86	262.84	255.69	248.58
Chromium (ug/l)						
1. FRIELAND	0.56	0.67	0.77	0.41	0.77	0.77
2. B'GUMMER	0.60	0.70	0.78	0.45	0.83	0.78
3. GRONETAL	4.24	4.24	4.24	4.60	3.29	3.85
BOD (mg/l)						
1. FRIELAND	9.59	9.17	9.05	9.66	9.05	9.05
2. B'GUMMER	10.15	9.38	9.21	10.35	9.57	9.21
3. GRONETAL	51.49	49.67	49.75	60.30	45.19	49.75
Phosphate (mg/l)						
1. FRIELAND	0.56	0.55	0.55	0.57	0.55	0.55
2. B'GUMMER	0.59	0.56	0.56	0.61	0.58	0.56
3. GRONETAL	4.45	4.30	4.99	6.13	4.78	4.30

chromium there is in the Northern region. This strategy even eliminates the supplies of water for sprinkling in order to reduce the demands for water.

A glance at the pollutant concentrations shown in the table shows that none of the strategies has a very large effect on any concentration. In this region, it would appear that managerial tactics offer no substantial possibilities to improve water quality. We again stress, however, the tentative nature of this result, which is due to the very aggregate nature of the MSDM representation of the water management infrastructure in this region.

9.3.4.4. The North Holland Region. Table 9.19 shows the six strategies and their effects on water quality in the North Holland region. The nominal strategy is the same as the MSDM strategy discussed in Sec. 9.2.3.4. Indeed, only one of the strategies, the Chromium strategy, is at all different. That strategy reduces extractions from the IJssel lakes to a minimum, by eliminating all deliveries of water for sprinkling and by reducing discharges to the North Sea (via link 43 NOHOLKAN) to their minimum of 2 m^3/s. This has only a small effect on the pollutant concentrations at the single node in this region, node 11 NORHOLL. Once again we must stress that the aggregate nature of the MSDM representation of the infrastructure in this region may conceal some opportunities for improving water quality by managerial tactics.

9.3.4.5. The Midwest and Utrecht Region. Table 9.20 shows the six strategies and their effects on water quality in the Midwest and Utrecht region. The nominal strategy is the same as the MSDM strategy discussed in the earlier Sec. 9.2.3.5. Since none of the nodes in this region are affected by thermal pollution, the Thermal strategy is the same as the nominal strategy. The BOD strategy also happens to be the same.

The other three strategies are quite different, however. The Chloride strategy cuts back sprinkling, and replaces extractions from the Amsterdam-Rijnkanaal by increases in the extraction from the Hollandsche IJssel. A large part of the extractions from the Amsterdam-Rijnkanaal is supplied via link 58 MERWKAN2, which delivers the water to node 17 LOPIKWAR. At the start of the decade, the chloride concentration of the water stored at LOPIKWAR is much lower than that of the water supplied from the canal. Thus the strategy minimizes contamination of the former by the latter. Similarly, the strategy minimizes the contamination of water stored at node 21 MIDWEST by reducing the demand for water there. This is the explanation for eliminating sprinkling.

The Chromium strategy is similar to the Chloride strategy in that it reduces the intake of water into LOPIKWAR via link 58 MERWKAN2. This avoids the contamination of water stored at LOPIKWAR by water from the canal, which contains a higher chromium concentration. Unlike the Chloride strategy, however, the Chromium strategy does not reduce

Table 9.19

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: NORTH HOLLAND REGION

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Strategy Description (m^3/s)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Water for Sprinkling	6.33	6.33	6.33	0.0	6.33	6.33
Other Consumption	13.80	13.80	13.80	13.80	13.80	13.80
Extractions from IJssel Lakes	30.00	30.00	30.00	22.80	30.00	30.00
Discharges to Waddenzee	2.87	2.87	2.87	2.00	2.87	2.87
Discharges to Noordzeekanaal	7.00	7.00	7.00	7.00	7.00	7.00
Target Sprinkling Water	6.33	6.33	6.33	6.33	6.33	6.33
Cutbacks in Sprinkling	0.0	0.0	0.0	6.33	0.0	0.0
Regional Agricultural Damage (Dflm)						
Damage from Salt	1.26	1.26	1.26	1.29	1.26	1.26
Drought Damage Plus Sprinkling Cost	2.67	2.67	2.67	7.02	2.67	2.67
Pollutant Concentrations						
Chloride (mg/l)						
11. NORHOLL	270.94	270.94	270.94	274.82	270.94	270.94
Chromium (ug/l)						
11. NORHOLL	1.08	1.09	1.08	0.86	1.11	1.07
BOD (mg/l)						
11. NORHOLL	8.64	8.64	8.64	9.05	8.64	8.64
Phosphate (mg/l)						
11. NORHOLL	0.67	0.67	0.67	0.69	0.67	0.67

Table 9.20

MSDM STRATEGIES FOR IMPROVING WATER QUALITY: MIDWEST AND UTRECHT REGION

Case Description

Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum

Infrastructure: 1976 + Fresh Zoommeer
Power Plant Inventory: 1985
Sprinkling Scenario: 1976
Shipping Scenario: 1976

Strategy Description (m^3/s)	Nominal MSDM Strategy	3-Deg Thermal Standard	200 mg/l Chloride Standard	Minimize Chromium Level	5 mg/l BOD Standard	0.3 mg/l Phosphate Standard
Water for Sprinkling	6.29	6.29	0.47	0.0	6.29	6.29
Other Consumption	32.72	32.72	32.72	32.72	32.72	32.72
Extraction from Hollandsche IJssel	18.73	18.73	30.16	22.16	18.73	35.00
Extraction from Amsterdam-Rijnkanaal	24.10	24.10	6.85	14.85	24.10	16.66
Extraction from IJssel Lakes	10.00	10.00	10.00	10.00	10.00	10.00
Discharges to Amsterdam-Rijnkanaal	7.60	7.60	7.60	8.07	7.60	16.42
Discharges to Noordzeekanaal	3.00	3.00	3.00	3.00	3.00	3.00
Discharges to North Sea	3.23	3.23	3.23	3.23	3.23	3.23
Target Sprinkling Water	7.98	7.98	7.98	7.98	7.98	7.98
Cutbacks in Sprinkling	1.69	1.69	7.51	7.98	1.69	1.69
Regional Agricultural Damage (Dflm)						
Damage from Salt	157.09	157.09	158.94	157.26	157.09	159.06
Drought Damage Plus Sprinkling Cost	3.42	3.42	6.61	6.70	3.42	3.42
Pollutant Concentrations						
Chloride (mg/l)						
17. LOPIKWAR	293.97	293.97	255.02	255.02	293.97	293.97
18. VECHT	165.55	165.55	162.87	162.87	165.55	199.23
20. GOUDA	311.29	311.29	317.44	318.72	311.29	316.96
21. MIDWEST	306.14	306.14	308.58	306.34	306.14	308.57
Chromium (ug/l)						
17. LOPIKWAR	34.10	34.10	19.84	19.84	34.41	34.10
18. VECHT	4.81	4.81	4.81	4.81	4.81	10.60
20. GOUDA	32.69	32.69	34.40	34.68	32.70	34.55
21. MIDWEST	5.13	5.13	6.01	4.84	5.26	6.04
BOD (mg/l)						
17. LOPIKWAR	9.86	9.86	7.94	7.94	9.86	9.86
18. VECHT	8.35	8.35	8.30	8.30	8.35	8.63
20. GOUDA	12.63	12.63	13.57	13.28	12.63	15.46
21. MIDWEST	5.77	5.77	5.96	5.75	5.77	6.14
Phosphate (mg/l)						
17. LOPIKWAR	0.86	0.86	1.33	1.33	0.86	0.86
18. VECHT	1.94	1.94	1.97	1.97	1.94	1.70
20. GOUDA	1.05	1.05	0.93	1.04	1.05	0.97
21. MIDWEST	0.78	0.78	0.76	0.77	0.78	0.76

extractions elsewhere on the Amsterdam-Rijnkanaal (links 64 LEIDRIJN and 70 ANGSTEL) in favor of increasing extractions from the Hollandsche IJssel. As the water flows north along the Amsterdam-Rijnkanaal, the chromium it contains is mostly lost through sedimentation, so that water from this source has less chromium than Hollandsche IJssel water. Because salt does not "decay," this phenomenon has no effect on the Chloride strategy.

A main feature of the Phosphate strategy is the increase in discharges to the Amsterdam-Rijnkanaal, which occurs via link 69 ARKVECHT. The purpose of the discharge is to flush the heavy phosphate discharges out of node 18 VECHT. The extra water to support this tactic is provided from LOPIKWAR via link 59 MERWKAN3. In the nominal strategy, this water is sent instead to node 21 MIDWEST; in the Phosphate strategy, this reduction in the supply to MIDWEST is compensated by an increase in extractions from the Hollansche IJssel. In addition, extractions from the Amsterdam-Rijnkanaal via link 70 ANGSTEL from node 14 DIEMEN are minimized, and the supply replaced by further increases in extractions from the Hollandsche IJssel (to the maximum of 35 m^3/s) because the phosphate concentration at DIEMEN is higher than that in the Hollandsche IJssel.

If we examine the effects of these strategies on the pollutant concentrations in the Midwest and Utrecht region, we see that in only a few instances are concentrations affected significantly. All the pollutants can be affected at LOPIKWAR; chloride, chromium, and BOD are all reduced by the Chloride and Chromium strategies, while phosphate is increased. All four pollutants are affected at the node VECHT by the Phosphate strategy, but are unaffected by the other strategies; chloride, chromium, and BOD are increased, while phosphate is reduced. It seems, therefore, that no single strategy can be defined for this region that causes a general, comprehensive improvement in water quality. Once again, of course, we must caution that the MSDM representation of this region is highly aggregated, and we may well have missed significant opportunities for improving water quality that a more detailed examination would reveal.

9.3.5. Observations and Conclusions

We can make several general observations about the effect of managerial tactics on water quality.

- The effects in the national waterways are generally very small, except on the Amsterdam-Rijnkanaal and the Noordzeekanaal, and in some instances on the Maas.
- Tactics which cause a significant reduction in the concentration of one pollutant at one national node frequently cause a significant increase in the

concentration of another pollutant at the same node, or of the same pollutant at a different node.
- With the exception of thermal pollution, managerial tactics never enabled a water quality standard to be met at any node for any pollutant, unless it was already met in the nominal case.
- Because of the highly aggregated nature of the MSDM regional networks, we cannot draw definitive conclusions regarding the effects of regional strategies on water quality. We feel, however, that the effect of aggregation is to conceal opportunities for improving water quality, rather than to suggest false hopes. Accordingly, we suggest that managerial strategies may be effective in improving water quality in the Southeast Highlands and Delta region. But more detailed study is needed for all regions.

It is clear, therefore, that managerial tactics have little value as devices for improving water quality in the MSDM network. What little value they have is related to thermal pollution. In our remaining results, therefore, we shall develop strategies without regard for their effect on chloride, chromium, BOD, or phosphate concentrations. Our nominal strategy will ignore thermal pollution as well, but as an excursion case we will develop strategies that include the effect of a three-degree thermal standard.

9.4. VARIATIONS OF RIVER FLOWS, RAINFALL, AND LAKE LEVELS

In this section we will consider the effect on the nominal MSDM strategy of varying the Rijn and Maas flows, the rainfall, and the starting amounts of water in the IJssel lakes and the weir ponds of the Maas. Except for these variations, the cases considered will be identical with the original case defined in Sec. 9.1.

We will consider variations in the Rijn flow at Lobith between 700 m^3/s and 1500 m^3/s. Flows as low as 700 m^3/s are quite rare, and so we think it unnecessary to consider lower flows. Flows higher than 1500 m^3/s are not interesting because additional water does not help to satisfy additional needs. That is, any needs unmet at a Rijn flow of 1500 m^3/s are unmet because of limitations of the infrastructure, and not because of any shortage of water.

We will consider variations in the flow on the Maas in perfect correlation with the Rijn flow, from 25 m^3/s when the Rijn flow is 700 m^3/s, to 41 m^3/s when the Rijn flow is 1500 m^3/s. In all cases, 22.4 m^3/s of Maas flow is reserved for Belgium, and only the remainder is available to the Netherlands. As in the case of the Rijn, this range of Maas flows includes flows near the lowest observed values (once the 22.4 m^3/s is subtracted) to values high enough to be uninteresting.

9.4.1. Net Evaporation = +66 mm/dec, Initial Lake Levels = Maximum

We first consider a series of cases which all start with the lake levels at their maxima, and which all have 66 mm of net open water evaporation. Only the Rijn and Maas flows will differ among the cases, as described above. Looking at Fig. 9.3, we see what effect variations in the river flows has on the optimal strategy. For Rijn flows above approximately 1000 m^3/s, the weir at Driel is set to provide a flow of 37.5 m^3/s in the Neder-Rijn. According to the upper rivers hydrologic equation, this dictates the flows in both the Waal and the IJssel. Farther west, the withdrawal at Tiel remains at its minimum level of 2 m^3/s until the Rijn flow drops to 900 m^3/s, while a constant 9.49 m^3/s is sent north from the Neder-Rijn on the Amsterdam-Rijnkanaal, regardless of the Rijn flow.

The flows described above are sufficient to determine the flows in the lower rivers part of the network, according to the lower rivers hydrological equation. For Rijn flows above 1000 m^3/s, the flows in the Nieuwe Maas, Oude Maas, and Nieuwe Waterweg all decrease steadily and uneventfully with decreasing Rijn flows. At a Rijn flow of 1000 m^3/s, however, the Nieuwe Maas flow reaches 207 m^3/s, the flow at which the Rotterdam salt wedge first begins to affect nodes 20 GOUDA and 21 MIDWEST. At this point the optimal strategy recommends that the weir at Driel be opened, and additional water be allowed to flow west to combat the salt wedge. Through coordinated manipulations of the Neder-Rijn flow (an indicator of the setting of the weir at Driel) and withdrawals at Tiel, the flow in the Nieuwe Maas is maintained at 207 m^3/s for all Rijn flows below 1000 m^3/s.

The coordination between the Neder-Rijn flow and the withdrawal at Tiel is interesting. As the Rijn flow drops from 1000 m^3/s to 900 m^3/s, the withdrawal at Tiel remains at its minimum level, 2 m^3/s. The strategy prefers to combat the salt wedge by increasing the Neder-Rijn flow instead. Between Rijn flows of 900 m^3/s and 800 m^3/s, the strategy is first to reduce the Neder-Rijn flow and compensate by increasing the withdrawal at Tiel, and then to reverse the process. The purpose of this maneuver is to prevent excessive losses to shipping on the IJssel. Indeed, during the first phase of the maneuver, between Rijn flows of 900 m^3/s and 860 m^3/s, the weir at Driel is set in such a way as to maintain the IJssel flow at a constant 141 m^3/s, which corresponds to a shipping depth of 16 decimeters.

An IJssel depth of 16 decimeters has no great significance in reality, but takes on an artificial significance in the model because of the manner in which we have represented the shipping loss functions. As described in App. G of Vol. VA, each shipping loss function is represented as a piecewise linear function, with abrupt changes in slope where successive linear pieces are joined. The slope of the function is one of the factors that determines the value of using water to reduce shipping losses (this is discussed in

Net evaporation = 66 mm/dec
Initial lake levels = maximum

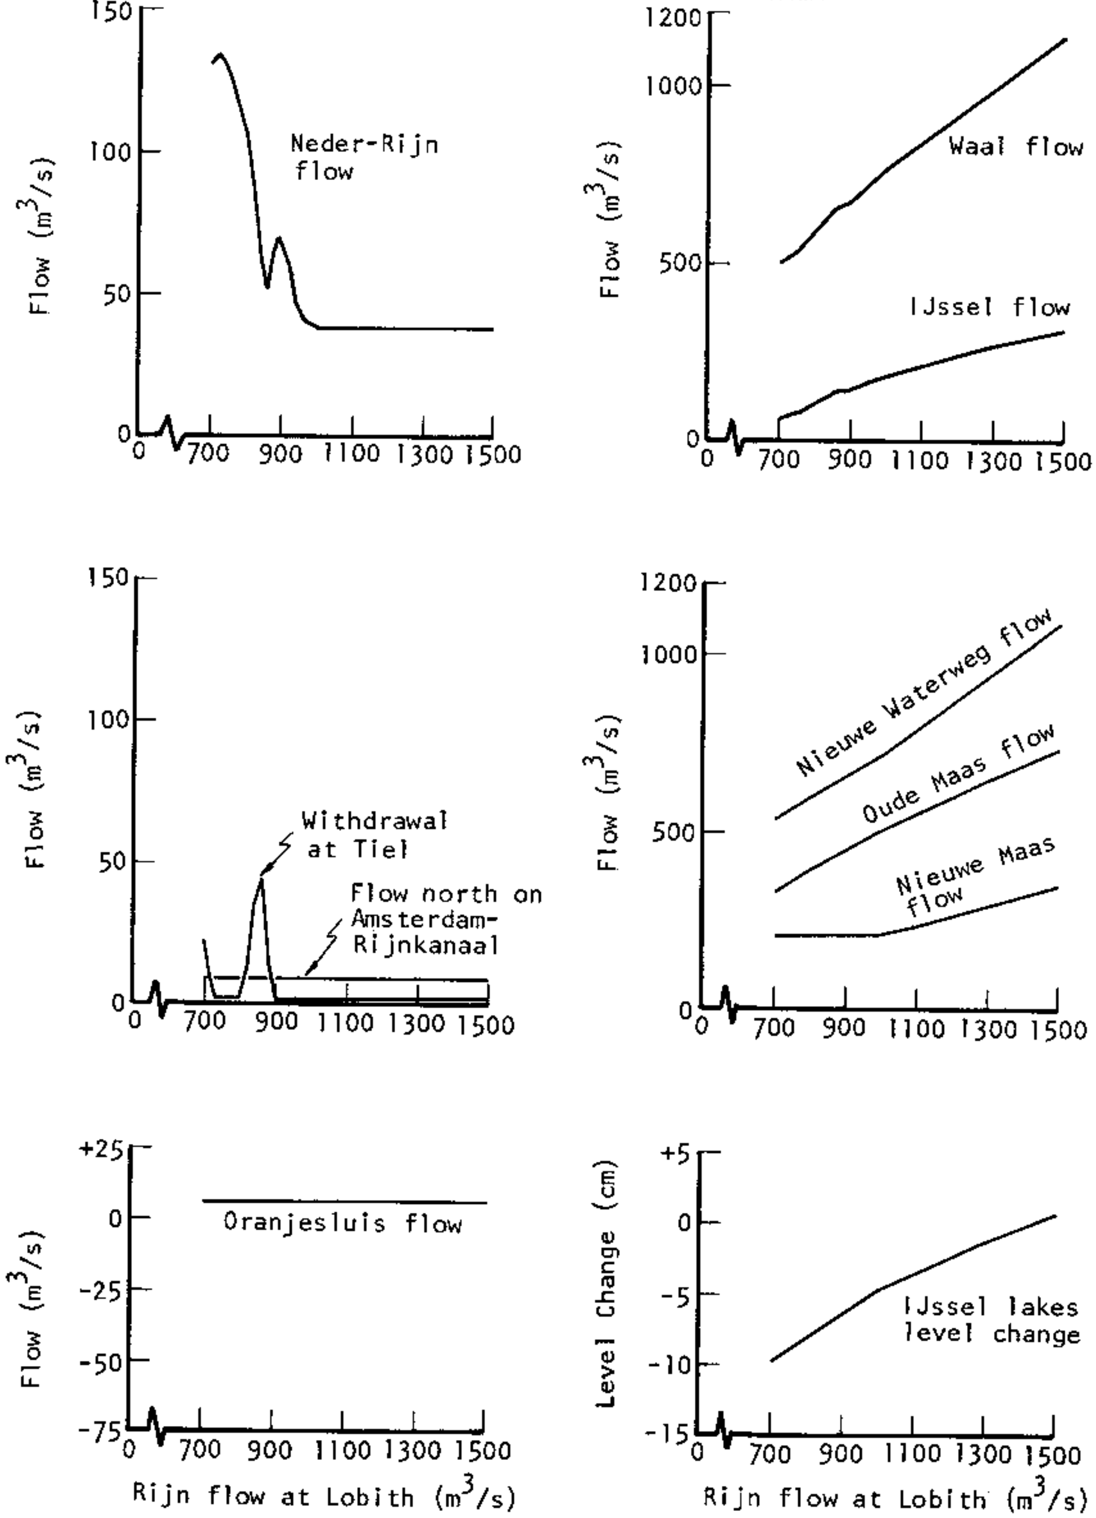

Fig. 9.3--MSDM strategy versus Rijn flow, case 1

Chap. 7). The point at which the slope changes abruptly is therefore a point where the value of water changes abruptly, and hence is a point at which the strategy may "stick" while the value of water for other purposes "catches up." An IJssel depth of 16 km is one such "sticking point." By suitably smoothing the shipping loss functions, this sort of "sticking" behavior could be eliminated.

By the time the Rijn flow drops to 800 m^3/s, the withdrawal at Tiel has once again reached its minimum value of 2 m^3/s, and it remains at this value until the Rijn flow drops to about 730 m^3/s. At this point, the weir at Driel is fully open. To maintain the flow in the Nieuwe Maas at 207 m^3/s for Rijn flows lower than 730 m^3/s requires more than a fully open weir; it requires an above-minimum withdrawal at Tiel as well.

In Fig. 9.3, we have also presented curves showing the flow of water out of the IJssel lakes via the Oranjesluis (link 38 ORANJESL), and the change in the level of the IJssel lakes. In this series of cases, the Oranjesluis flow is constant, and the IJssel lake levels drop steadily, as the Rijn flow decreases.

Because it is so easily described, Fig. 9.3 shows nothing about the strategy for the Maas. The level of the first weir pond (node 27 MAASLOO) is allowed to drop as far as necessary to provide a total flow of 2.5 m^3/s in the Julianakanaal (link 10 JULCANL1), or to the level of the Waal, whichever is the larger reduction in level. In the cases described in this section, it is never necessary or desirable to reduce the levels of other weir ponds, or to reduce extractions from the Maas for use in the Southeast Highlands and Delta region.

Figure 9.4 shows the cost components and the total costs of the strategies described in this section. As we pointed out in Sec. 9.2.1, the magnitudes of the cost components are not important. Rather, one should consider the relative changes in the different cost components. From the figure, we see that the thermal cost is constant (since we have ignored the three-degree standard in these cases) and the drought damage plus sprinkling cost is very nearly constant. The latter fact means that cutbacks in sprinkling are not required or, indeed, desirable, even for Rijn flows as low as 700 m^3/s.

The other two components, the shipping cost and the salt damage to agriculture, do show substantial increases as the Rijn flow decreases. It is easy to accept the increase in the shipping cost, in view of the foregoing discussion of the strategy. But the reason for the increase in salt damage to agriculture is perhaps not so clear. After all, the strategy never allows the Rotterdam salt wedge to affect nodes 20 GOUDA or 21 MIDWEST. But we have assumed in these cases that the chloride content of the Rijn remains constant, even as the Rijn flow decreases. Thus, the chloride concentration, which is the cause of the salt damage, increases as the flow decreases. In

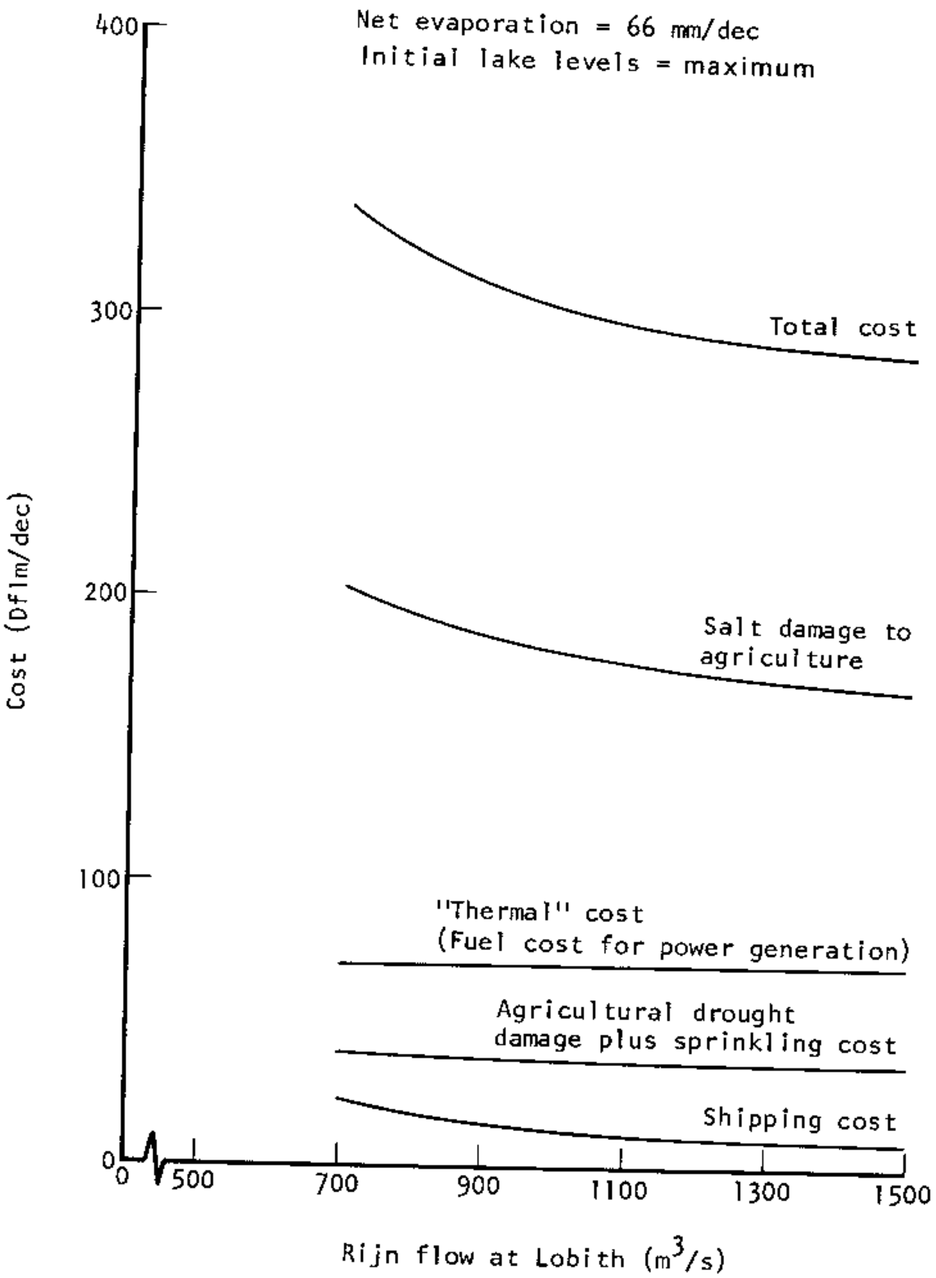

Fig. 9.4--Cost components and total cost of MSDM strategy versus Rijn flow, case 1

fact, virtually all of the increase in salt damage to agriculture that is seen in the figure is due to the increase in the chloride concentration in the Rijn.

9.4.2. Net Evaporation = +66 mm/dec, Initial Lake Levels = Minimum

Next we consider a series of cases which differ from those of the previous section only by starting the decade with all lakes at their minimum levels. As before, the cases have 66 mm of net open water evaporation, and only the Rijn and Maas flows differ among the cases. Looking at Fig. 9.5, we see what effect the variation in Rijn flow has on the optimal strategy. For Rijn flows larger than 1430 m^3/s, the strategy is identical to the strategy of the previous section, because the MSDM optimal strategy actually <u>increases</u> the lake level. For Rijn flows below 1430 m^3/s, strategies from the previous section draw down the IJssel lake level; in the present section, this is not allowed.

The fashion in which the strategies of this section maintain the IJssel lakes at their minimum levels surprised us. As the Rijn flow declines below 1430 m^3/s, the strategy chooses to <u>open</u> the weir at Driel. This increases the flow in the Neder-Rijn at the expense of the flows in the IJssel and the Waal. But the IJssel transports Rijn water to the IJssel lakes; thus a decrease in the IJssel flow makes it necessary to supply still more water to the IJssel lakes from some other source.

That other source is the Amsterdam-Rijnkanaal. The water in the Neder-Rijn is sent north on the Amsterdam-Rijnkanaal and into the IJssel lakes via the Oranjesluis. It is possible to operate the Oranjesluis in this reverse direction by means of the pumping stations Zeeburg and Schellingwoude, whose combined capacity is 73 m^3/s.

For Rijn flows between 1430 m^3/s and 1300 m^3/s, the water required on the Amsterdam-Rijnkanaal is provided entirely by opening the weir at Driel wider and wider. The other option would be to leave the weir closed, and provide water to the Amsterdam-Rijnkanaal by withdrawing at Tiel. MSDM has determined, however, that this would result in a higher shipping cost than the strategy of opening the weir.

When the Rijn flow drops below 1300 m^3/s, it finally becomes more cost-effective to begin closing the weir at Driel, and to provide water to the Amsterdam-Rijnkanaal by withdrawing at Tiel. The strategy undergoes various gyrations as the Rijn flow decreases still further, gyrations that reflect the "sticky" behavior of MSDM that we discussed in the previous section. Below a Rijn flow of 1100 m^3/s, the IJssel lake level can no longer be maintained. The weir at Driel is fully closed, so that the flow in the IJssel is maximized. The

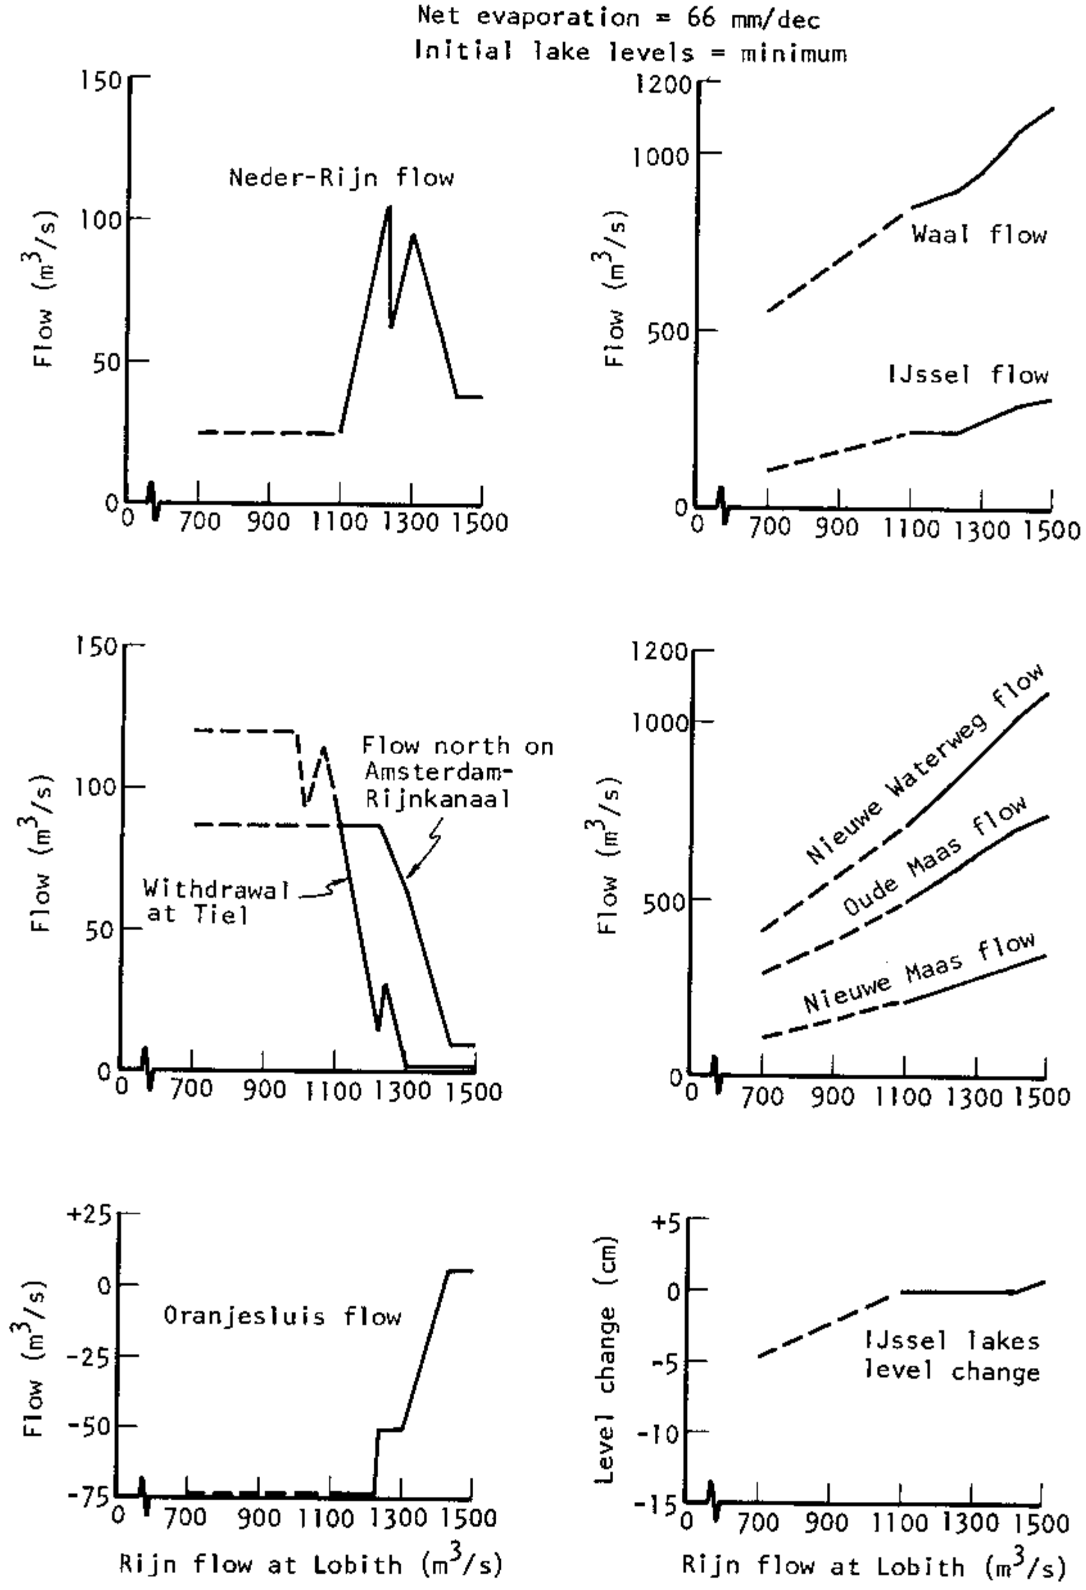

Net evaporation = 66 mm/dec
Initial lake levels = minimum

Fig. 9.5--MSDM strategy versus Rijn flow, case 2

pumps on the Oranjesluis are working at full capacity. In this range of Rijn flows, all thought of combating the Rotterdam salt wedge has been abandoned; the Nieuwe Maas flow is allowed to drop well below the magic value of 207 m^3/s.

In fact, the cases described in this section become infeasible for Rijn flows of 1100 m^3/s and below. The standard formulation of MSDM has an absolutely rigid requirement that the IJssel lake level be maintained at or above its minimum. This is why the strategy is presented in the figure as dashed lines for Rijn flows below 1100 m^3/s. To obtain strategies for these infeasible cases, we had to reformulate MSDM, to allow the lakes to drop below their minimum levels. To prevent this from occurring when unnecessary, we imposed a large penalty on violations of the minimum level condition. In effect, then, strategies for Rijn flows below 1100 m^3/s maximize the lake levels.

Figure 9.5 shows nothing about the strategy for the Maas. Because the weir ponds start the decade at their minimum levels, they cannot be drawn down to augment the Maas flow when it is otherwise insufficient. There are minimum flow requirements of 2.5, 1, and 0.35 m^3/s, respectively, on links 10 JULCANL1, 11 MAAS1, and 72 ZUIDWLM2, plus a treaty requirement to supply 22.4 m^3/s to Belgium. For Maas flows below approximately 26 m^3/s, therefore, MSDM is infeasible. We therefore reformulated the problem to permit MSDM to reduce the Belgian requirement, but imposed a large penalty on each cubic meter of reduction. Before MSDM reduces the Belgian requirement, therefore, it will first make all other possible reductions in extractions from node 27 MAASLOO.

Figure 9.6 shows the cost components and total costs for the strategies discussed in this section. Once again, the thermal cost component is constant, reflecting the fact that these strategies ignore the three-degree thermal standard. However, all other components increase significantly as the Rijn flow decreases. The increase in shipping cost is not much greater than we observed in Fig. 9.4. However, where we observed almost no increase in drought damage and sprinkling cost before, we now observe a large increase. And we observe a much larger increase in salt damage to agriculture, reflecting the fact that at low Rijn flows we have no water to spare from maintaining the lake level, and cannot therefore combat the salt wedge.

The increase in drought damage plus sprinkling cost occurs rather sharply at a Rijn flow of 1100-1200 m^3/s. The cause of the increase is cutbacks in sprinkling. Almost no cutbacks occur until the need to maintain the lake levels forces them to be made; at that point, every cutback is made that will increase the water supplies to the lake.

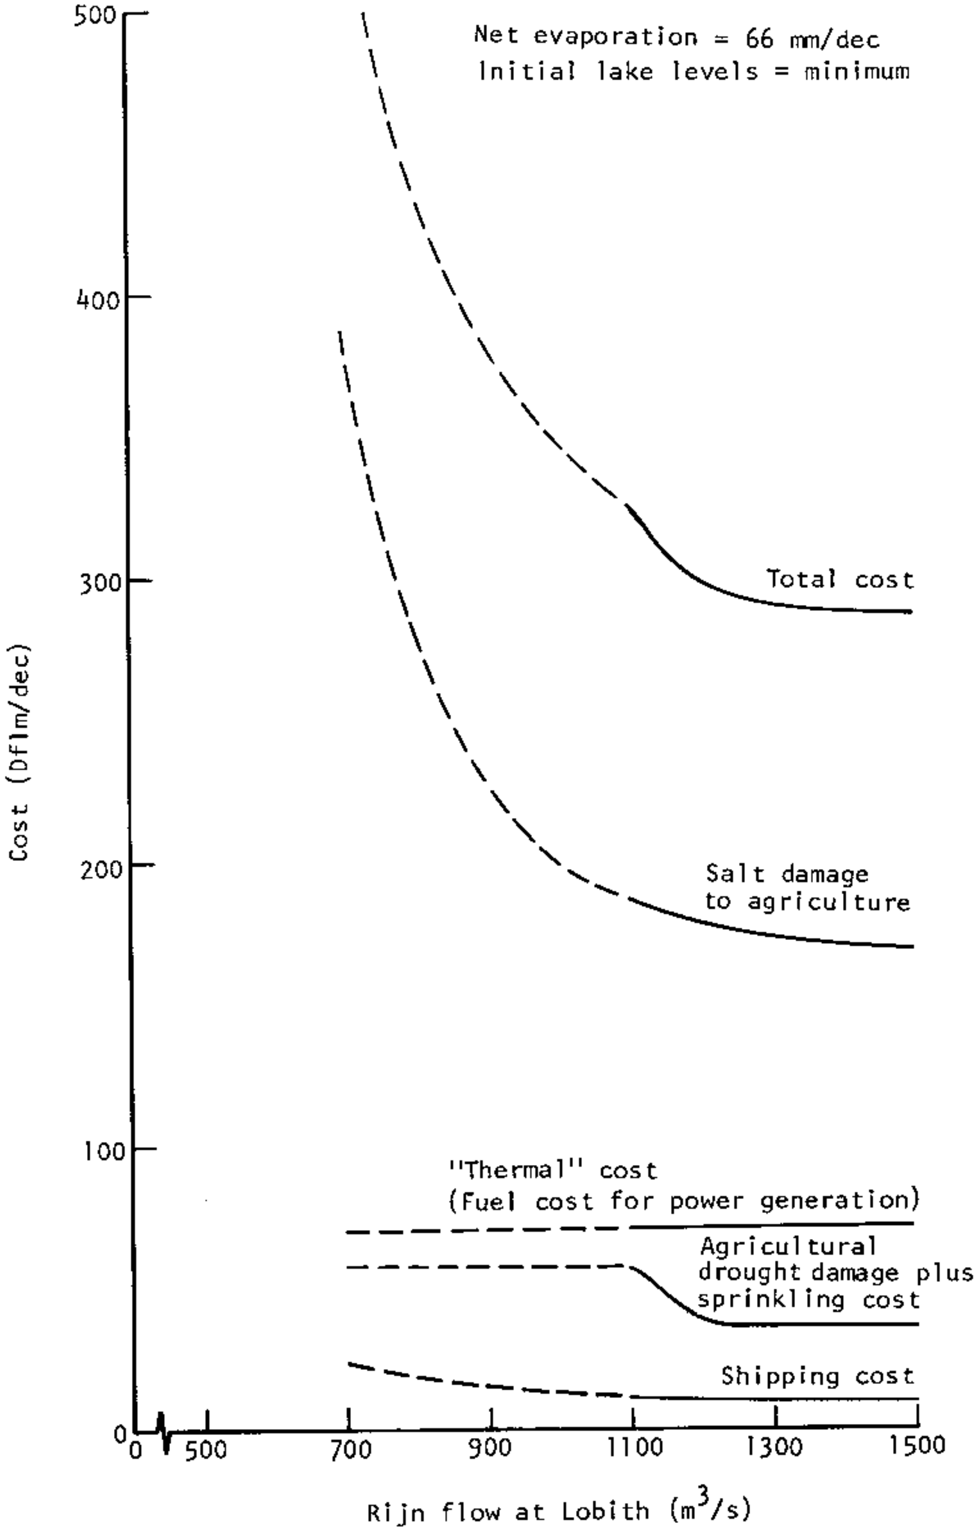

Fig. 9.6--Cost components and total cost of MSDM strategy versus Rijn flow, case 2

9.4.3. Net Evaporation = -11 mm/dec, Initial Lake Levels = Any

Finally, we consider a series of cases which have a net open water evaporation of -11 mm--i.e., which have more rain than evaporation. Again, only the Rijn and Maas flows differ among the cases. For these cases the starting lake levels are irrelevant, since at no values of river flows is it necessary to extract water from the lakes. In fact, in Fig. 9.7 it can be seen that at all Rijn flows, the strategy supplies enougn water to the IJssel lakes to raise their level substantially. (In practice, if some of this water were not needed in the IJssel lakes, it could be discharged instead through the Noordzeekanaal into the North Sea.)

The strategies for these cases are easy to describe. The weir at Driel is kept fully closed, allowing only the minimum flow of 25 m^3/s in the Neder-Rijn. The withdrawal at Tiel and the flow north on the Amsterdam-Rijnkanaal are also at their minima. This strategy gives rise to Waal and IJssel flows, and flows in the lower river branches, as shown in the figure. On the Maas, the strategy lowers the water level at node 27 MAASLOO until its depth is the same as the Waal depth. This reduces the cost associated with conserving water on link 10 JULCANL1 (see Chap. 7).

Figure 9.8 shows the cost components and total cost for these cases. Again, the thermal cost is constant, since we have ignored the three-degree thermal standard. In these cases, drought damage plus sprinkling cost is constant (and very small) also. Because the decade is so wet (net evaporation = -11 mm), sprinkling demands are zero, and drought damage is suffered only because we assumed that the farmers' soil was quite dry at the start of the decade.

The salt damage to agriculture rises rather little as the Rijn flow decreases. There is an interesting break in this cost component at a Rijn flow of approximately 1040 m^3/s. Above this flow, the chloride concentration of the Rijn is less than or equal to 300 mg/l, which is the initial chloride concentration at node 21 MIDWEST. Accordingly, the optimal strategy is to flush MIDWEST with as much Rijn water as possible. For Rijn flows less than 1040 m^3/s, the chloride concentration of the Rijn exceeds 300 mg/l. In the cases in Sec. 9.4.1, in which the net evaporation is 66 mm/decade, there is a substantial demand at MIDWEST for water for both sprinkling and level control; thus Rijn water must be supplied to MIDWEST in spite of the fact that it will raise the chloride concentration there. In the cases discussed in the present section, by contrast, there is no demand at MIDWEST. The negative net evaporation (i.e., the surplus of rain over evaporation) implies that there is excess water at MIDWEST that must be discharged. Thus, for Rijn flows less than 1040 m^3/s, no Rijn water at all is allowed to contaminate the water stored at MIDWEST, and salt damage to agriculture is correspondingly lower.

The shipping cost also increases as the Rijn flow decreases. However, the increase is much smaller than in Figs. 9.4 and 9.6. Because there is no intake of Rijn water into MIDWEST at low Rijn

Net evaporation = -11 mm/dec
Initial lake levels = any

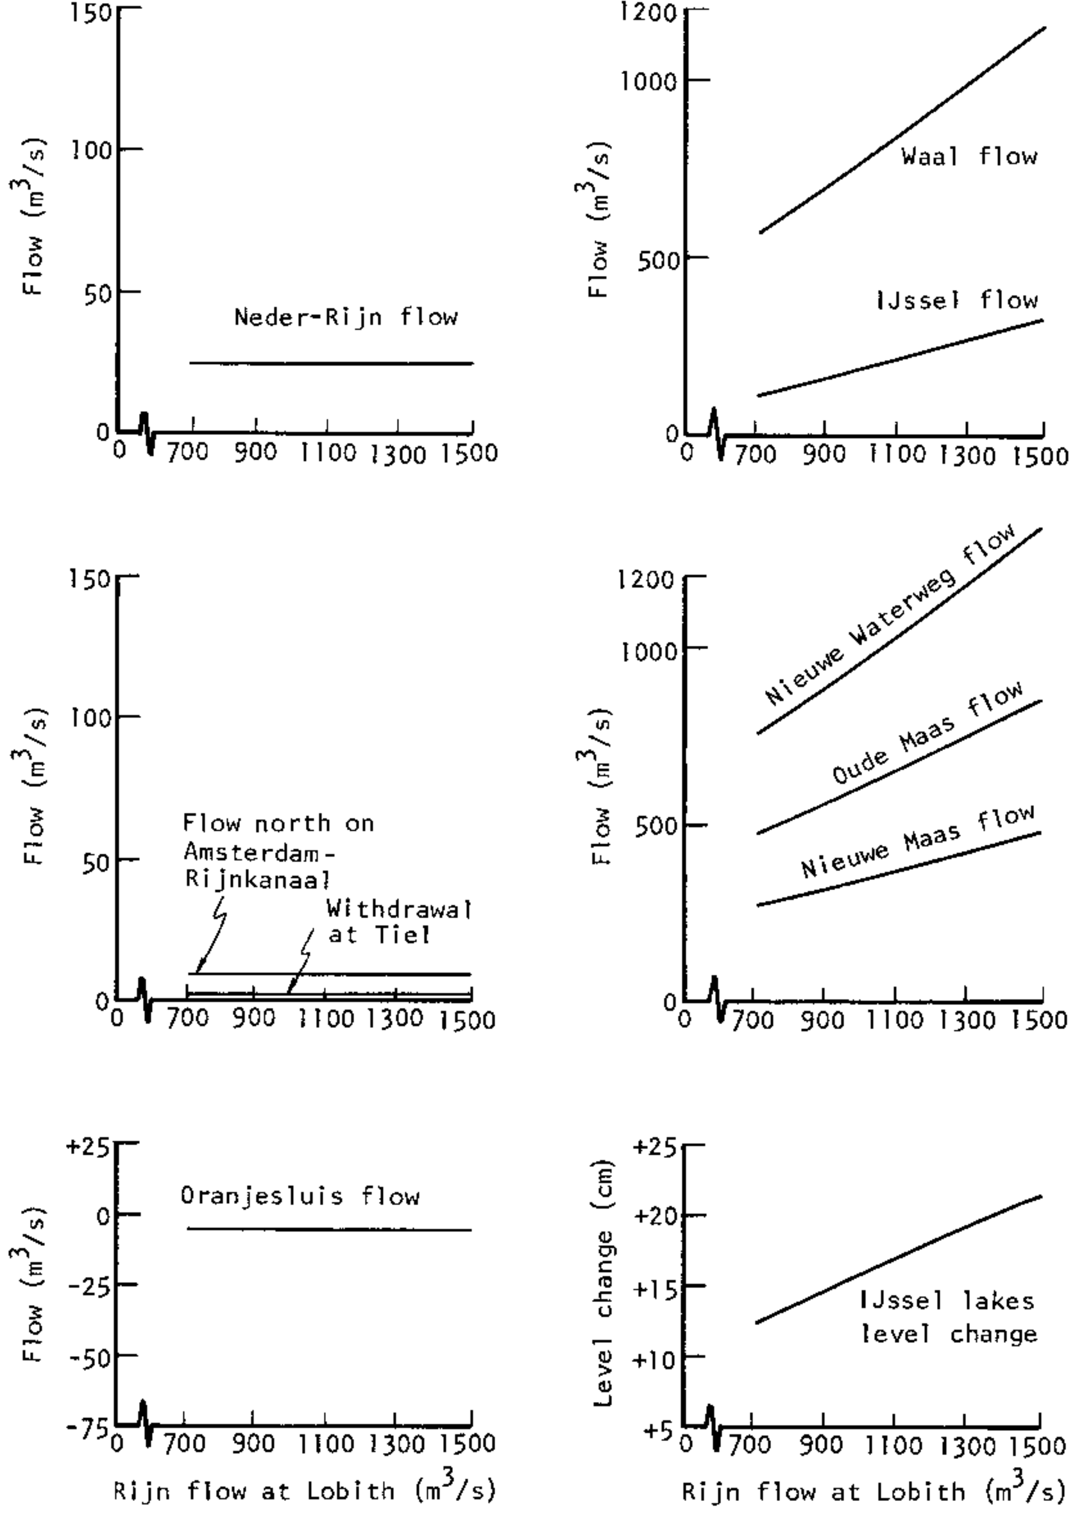

Fig. 9.7--MSDM strategy versus Rijn flow, case 3

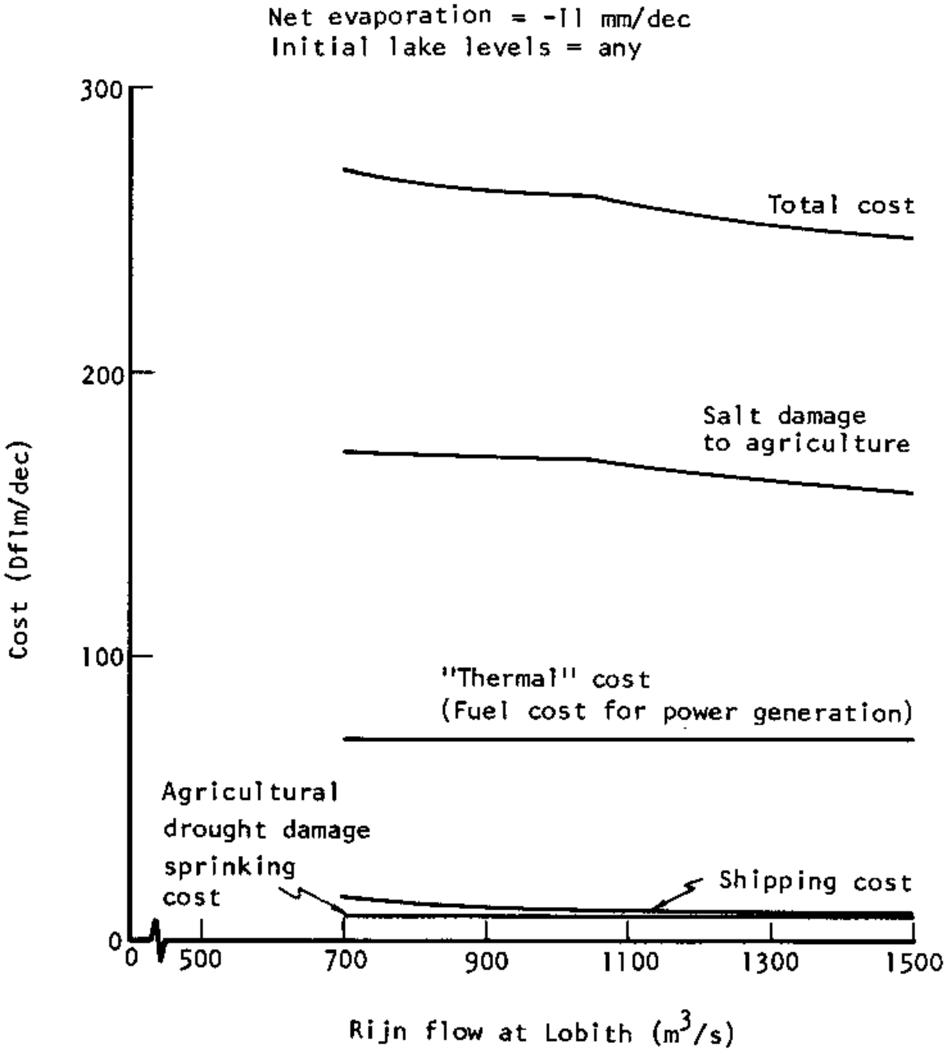

Fig. 9.8--Cost components and total cost of MSDM strategy versus Rijn flow, case 3

flows, there is no need to combat the Rotterdam salt wedge. The flow out of MIDWEST on the Hollandche IJssel (link 61 HOLIJSEL) prevents the salt wedge from affecting GOUDA and MIDWEST regardless of how far it intrudes in the Nieuwe Maas. Accordingly, the effect of the strategies for these cases is to minimize shipping costs for each Rijn flow.

9.5. THE EFFECT OF THE THREE-DEGREE THERMAL STANDARD

In this section, we consider the effect of imposing the three-degree thermal standard on every node of the network. We will compare two managerial strategies which can meet the standard. The first combines the nominal strategy with shifting the generating load among power plants in such a way as to meet the standards. It would increase the generating cost by approximately 0.9 Dflm/dec to shift the load away from power plants at nodes 15 AMSTEDAM and 16 HAL+IJMU, and 0.28 Dflm/dec to shift the load away from the power plant at B'GUMMER. There would also be a cost that, depending on the Maas flow, could vary between zero and one Dflm/dec for shifting the load away from the power plants at nodes 29 LINNE and 33 GERTRUID (see Chap. 5).

The second managerial strategy we consider is the Thermal strategy discussed in Sec. 9.3. The Thermal strategy releases more water from the weir ponds on the Maas, to provide additional cooling capacity for the power plants Amer (at node 33 GERTRUID) and Maasbracht (at node 29 LINNE). The Thermal strategy also extracts 128 m^3/s more water from the IJssel lakes (if available), and sends 105 m^3/s out the Noordzeekanaal to cool power plants at nodes 15 AMSTEDAM and 16 HAL+IJMU, and 23 m^3/s past node 2 B'GUMMER to cool the Bergum power plant. All of these differences involve diverting water from storage in order to cool power plants. In Sec. 9.3, we asked whether the use of water in these three ways was more valuable than retaining the water in storage for possible future uses. In this section, we investigate whether using IJssel lakes water for cooling is preferable to storing it for the future. We have not carried out a similar exercise for the weir ponds of the Maas, but we consider that such an exercise would be worthwhile.

9.5.1. Average Annual Cost Difference for Power Generation

We first calculate the average annual benefits of changing from the nominal strategy to the Thermal strategy. This benefit arises because the Thermal strategy extracts IJssel lakes water to cool power plants more often than the nominal strategy. If the IJssel lakes are filled to capacity, and water must be discharged to prevent the level from rising further, the nominal strategy should use it to cool the power plants. Conversely, the nominal strategy reduces the flow of cooling water by 128 m^3/s whenever the IJssel lake levels are below their maxima. But the Thermal strategy will only reduce the flow of cooling water when the IJssel lakes are at their minimum levels.

When a strategy must cut back the flow of cooling water from the IJssel lakes by 128 m^3/s, the fuel cost for generating electric power increases by 1.18 Dfl/dec. To estimate the frequency with which the nominal strategy suffers this cost, we have carried out a simulation using the Distribution Model applied to 44 years of river flow and rainfall data. The simulation assumed the RWS managerial strategy was used. We described this strategy in Sec. 9.2, and pointed out that its effect on the IJssel lake level was nearly the same as the nominal MSDM strategy. Accordingly, the simulation should provide a reasonable guide to how frequently the IJssel lakes are not filled to capacity under the nominal strategy, and hence how frequently it will fail to provide cooling water. In the 44 simulated years, the IJssel lakes dropped below capacity in only 14 decades.

But the Thermal strategy must also cut back the flow of cooling water the IJssel lake levels reach their minima. To estimate the frequency of this occurrence we modified the results of the 44-year simulation as follows. First we note that an extra extraction of 128 m^3/s from the IJssel lakes is sufficient to reduce the lake levels by 5.48 cm per decade. For each year of the simulation, we examine the lake levels decade by decade, considering only the summer decades (there is always ample water available, and only small demands, during the winter half-year). If we come across a sequence of decades in the simulation in which the lake levels remain below their maxima, we subtract 5.48 cm from the level in the first decade of that sequence, twice 5.48 cm in the second decade, and so forth.

Undoubtedly, this procedure will fail to discover some decades in which an extra extraction of 128 m^3/s would lower the lake levels below their maxima, but in which the simulation itself shows the lake levels to be at their maxima. Undoubtedly also, there are occasions in which the decade preceding a sequence of below-maximum decades would be of this kind, and that the lake levels estimated by this procedure are therefore somewhat too high. Nevertheless, we believe that this procedure will give reasonable estimates of the lowest levels that would have been achieved had the simulation been carried out using the Thermal strategy in place of the RWS strategy.

In the present infrastructure, the minimum lake level is NAP - 40 cm. Using this procedure, we estimate that there would have been five decades during the 44 years of simulation with the Thermal strategy during which the lakes would have been at their minimum levels, and hence when the extractions of cooling water from the IJssel lakes could not be continued. Thus, adopting the Thermal strategy reduces the frequency with which the fuel cost of power generation rises by 1.18 Dflm/dec, from 14 to 5 decades per 44 years. The Thermal strategy thus conveys an average annual benefit from this source of 0.241 Dflm/yr, when compared to the nominal strategy.

9.5.2. Average Annual Cost Difference from Lake Depletion

But there is another source of cost that we have only considered cavalierly in MSDM. This is the cost that may be incurred in the event that the IJssel lakes reach their minimum levels. Recall that we assumed that water stored in the IJssel lakes had a small value, 0.00122 Dfl/m^3. We introduced this value in order to prevent water from being discharged into the North Sea for no reason at all, and we explained it by saying that we might find a profitable use for the water in future decades. A comparison of Figs. 9.4 and 9.6 shows the source of this possible future profit. The strategy cost is much higher if there is no water available in storage (and if the Rijn flow is low and the decade has little rain as well). Because the Thermal strategy extracts more from the IJssel lakes than the nominal strategy, it increases the probability that the IJssel lake level will reach its minimum, and hence that the cost difference illustrated in Figs. 9.4 and 9.6 will be realized.

Figure 9.9 shows the annual minimum summer levels that we estimate would have been achieved under the nominal and Thermal strategies. We derived these minimum levels from the 44-year simulation discussed earlier. The levels are plotted against frequency; the lowest level at a frequency of 1/44, the next lowest at 2/44, etc. To estimate how frequently the lakes might reach their minimum level of NAP - 40 cm under the nominal strategy, it is necessary to extrapolate the upper curve of the figure. A rough guess is that the lakes might reach their minimum level once in 200 years--i.e, with an annual probability of 0.005. From the lower curve of Fig. 9.9, we can estimate the analogous probability for the Thermal strategy to be 0.054.

When the lakes are at their minimum levels, and the Rijn flow and rainfall are sufficiently small, a cost will be incurred. We estimate from Figs. 9.4 and 9.6 that the total cost differs by about 25 Dflm/dec when the Rijn flow is 1100 m^3/s, and 200 Dflm/dec when the Rijn flow is only 700 m^3/s. When we consider that the decade from which these costs are obtained has no rain and very large evaporative and other demands, we can justifiably say that the difference in cost between having water stored in the IJssel lakes and having none is between 25 and 100 Dflm/dec.

But this cost is only borne if the decade has little rain and low Rijn flow. In fact there is no incremental cost at all if the Rijn flow exceeds 1200 m^3/s, even when the rainfall is zero and the evaporation is as high as 66 mm/decade. For Rijn flows smaller than 1200 m^3/s, it should be possible to determine rates of net evaporation (i.e., evaporation minus rain) which should also result in zero incremental cost. At worst, if all demands for water except evaporation from the IJssel lakes remains as high as when net evaporation is 66 mm/decade, we should be able to compensate for a

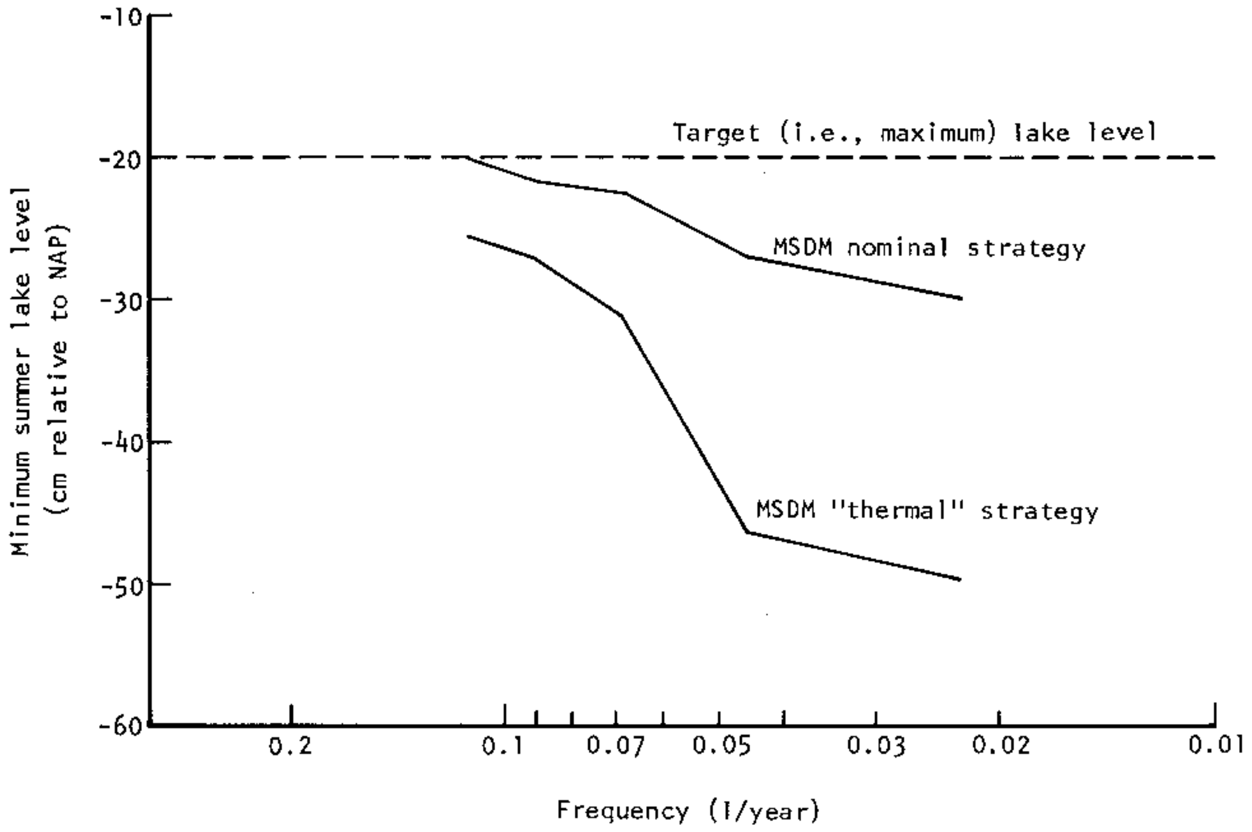

Fig. 9.9--Frequencies of low summer lake levels: present infrastructure and present sprinkling scenario

reduction in the IJssel River flow by an equal reduction in evaporation from the lakes. Thus, at a Rijn flow of 1200 m^3/s, the IJssel flow is 239 m^3/s and the evaporation from the IJssel lakes is 154 m^3/s, for a difference of 85 m^3/s. At a Rijn flow of 1000 m^3/s, the IJssel flow is 185 m^3/s, so if we reduce evaporative losses to 100 m^3/s, the incremental cost of having no available water stored in the lake should remain zero. An evaporative loss of 100 m^3/s from an area the size of the IJssel lakes (2045 square kilometers) is equivalent to 43 mm/decade of evaporation.

By examining the 44 years of river flow, rainfall, and evaporation data used in the simulation, we discovered that there is a probability of 0.021 that a summer decade would have a nonzero incremental cost. From our earlier discussion, we know that the Thermal strategy increases the probability of depleting the IJssel lakes from 0.005 to 0.054 per year, or an increase in the annual probability of 0.049. We have estimated that the incremental cost that will be suffered if the lakes reach their minimum levels, and the following decade is dry enough, will be between 25 and 100 Dflm. Multiplying these three figures together provides an estimate of the average annual difference in the cost due to lake depletion of the two strategies. That estimate is between 0.026 and 0.103 Dflm/yr.

9.5.3. The Thermal Strategy Is Preferable

In Sec. 9.5.1 we estimated that the Thermal strategy would save an average of 0.241 Dflm/yr in fuel costs for power generation, as compared with the nominal strategy. In Sec. 9.5.2 we estimated that this savings would be partially offset by no more than 0.103 Dflm/yr due to the increased probability of depleting the IJssel lakes. On balance, then, it appears that the Thermal strategy should be preferred to the nominal strategy.

However, the average annual amounts are small, so the motive to implement the Thermal strategy is not strong. Further, the benefits of the Thermal strategy (relative to the nominal strategy) will occur in perhaps one year out of nine or ten, whereas the costs will occur in only one year out of 1000--that is, when the lake reaches its minimum level <u>and</u> this event is followed by a dry enough decade. And when the costs occur, they are likely to be very large--between 25 and 100 Dflm. It is quite reasonable that one might be willing to pay 0.241 Dflm/yr in power generation costs in order to avoid a very rare, but very costly event. This, after all, is the principle behind insurance.

Moreover, there are strategies intermediate between the Thermal strategy and the nominal strategy. For example, we could use the Thermal strategy as long as the levels of the IJssel lakes remained at least five cm above their minima, but switch to the nominal strategy for lower levels. By cutting back the use of water for cooling in this way, we would reduce the probability of depleting the lakes. Indeed, the reduction in this probability might so reduce the average annual cost due to lake depletion that such an intermediate strategy would be preferable to both the nominal and Thermal strategies. In our opinion, it would be worthwhile investigating such intermediate strategies.

As a by-product of the comparison we have carried out in this section, we can form a better estimate of the value of water stored in the IJssel lakes (the insurance functions mentioned in Sec. 9.1 above). There were nine decades in 44 years during which the Thermal strategy extracted 128 m^3/s more from the IJssel lakes than did the nominal strategy. This amounts to a difference of 22.94 million cubic meters per year left in storage, at a difference in cost between 0.026 and 0.103 Dflm/yr. The average value of this water is therefore between 0.0011 and 0.0045 Dfl/m^3, as compared with our original guess of 0.00122 Dfl/m^3. Of course, we have estimated only an average value. Surely, the value of having an additional cubic meter in storage will be much smaller if the lakes are nearly full, and larger if they are nearly empty. This is another way of saying that it is probably worthwhile to use lake water for cooling power plants when the lake levels are high, but that it will be preferable to conserve the water once the lake levels have dropped sufficiently.

Finally, we note that all results in this section have assumed that power plants at AMSTEDAM, HAL+IJMU, and B'GUMMER will be subject to the three-degree thermal standard. This standard does not presently apply at these locations, and it appears unlikely that it will be soon applied. Instead, a seven-degree standard is applied at these locations. It is approximately true that the amount of water needed for optimum cooling to seven degrees excess temperature is three-sevenths as large as the amount needed for optimum cooling to a three-degree standard. Thus, a "Seven-Degree Thermal" strategy would have considerably less chance of depleting the IJssel lakes, and would look correspondingly even more preferable to the nominal strategy.

9.6. DESCRIBING STRATEGIES IN TERMS OF PRIORITIES

MSDM constructs its strategies by minimizing an objective function while satisfying constraints. (Actually, as we explained in Chap. 8, MSDM cannot truly minimize the objective. Instead, with the help of judicious human intervention, it finds strategies that yield nearly minimum values of the objective.) However, this description of how to construct a near-optimal strategy fails to suggest the form the strategy should take.

In this section, we will develop a more transparent description of the optimal strategy by comparing Figs. 9.3 through 9.8. This description takes the form of priorities attatched to various uses of water. High priority uses are to be satisfied first. Once no more can be done to satisfy a high priority use, a lower priority use can be considered. In effect, the priority list attempts to replace the objective function of MSDM, and to obviate the need for the complex trade-offs among uses implicit in the objective. (The attempt is largely but not completely successful. We discuss some of the problems with using the priority list below). The order, from high priority to low, is the following:

Priority 1: Supply level control requirements for boezems and lakes, and meet all other constraints on MSDM.

Priority 2: Supply water to farmers for irrigating their crops. Also, establish certain nominal flushing rates for locks at which salt intrusion causes damage to crops grown locally.

Priority 3: Trade off shipping losses due to low water on the Waal and the IJssel, the dredging cost (or other sedimentation cost) due to withdrawals at Tiel, and salt damage to agriculture due to the Rotterdam salt wedge, by simultaneously adjusting the Neder-Rijn flow (by adjusting the weir at Driel) and withdrawals at Tiel.

Priority 4: Use water from the IJssel lakes for cooling the power plants at AMSTEDAM, HAL+IJMU, and B'GUMMER to whatever thermal standard has been set for them.

Priority 5: Raise the IJssel lakes to their maximum levels to meet the possible future needs for water.

Priority 6: Use water for flushing boezems and ditches, and for raising the flushing rates of locks with local salt intrusion above the nominal rates established in Priority 2.

This ordering of water uses corresponds roughly to the relative economic values that water has in the various uses, which we discussed in earlier chapters. It cannot precisely reflect the different values of water because some of the values vary strongly with the amount of water devoted to that use. For example, the marginal value of water used for combating the Rotterdam salt wedge (Priority 3) is very high when the Rijn flow is low and the salt wedge has penetrated past the mouth of the Hollandsche IJssel. However, when the Rijn flow is high enough, it drives the salt wedge downstream of the Hollandsche IJssel, and additional water to combat the salt wedge is valueless.

9.6.1. How the Priorities Define the Strategy

Several complications make it difficult to translate the priority scheme given above into a workable managerial strategy. Sometimes, a particular cubic meter of water can be used for several different purposes in succession. For example, water flowing north along the IJssel River benefits shipping by raising the water level there (Priority 3). But once it reaches the IJssel lakes, it can be further used to sprinkle crops in Friesland or North Holland (Priority 2), or to raise the IJssel lakes level (Priority 5). As another example, some of the water used to flush the boezems of Rijnland (Priority 6) might be discharged into the Noordzeekanaal at Spaarndam or Halfweg, from where it could be used to cool the power plant at HAL+IJMU (Priority 4) before being discharged into the North Sea.

In addition, it is frequently true that a particular cubic meter of water is in a location to from which it can be used to satisfy only some demands. For example, once the weir at Driel has been entirely closed, no further means exist to send water north along the IJssel for the benefit of shipping. As another example, there may be ample water stored in the IJssel lakes to satisfy all demands for sprinkling water in the Southeast Highlands, but there is no means to transport the water to the point of demand.

Under such circumstances, the priority scheme requires that each cubic meter of water be used for the highest priority use for which it can be used. In our first example, the water that cannot help shipping on the IJssel may be used to flush the boezems of the midwest (a Priority 6 use). If all of it cannot be used for this low value purpose, some may simply be discharged into the North Sea, a disposition of water that offers no benefit to anyone (in this example, we have assumed that the Rijn flow is adequate to prevent the salt wedge from reaching the Hollandsche IJssel). In our second example, the water in the IJssel lakes might be used to cool the power plants at AMSTEDAM, HAL+IJMU, and B'GUMMER (priority 4), since it cannot be used for sprinkling in the Southeast Highlands (Priority 2).

We should note, however, that most of the time there is ample water available to provide for all users, of all priority classes. The salt wedge rarely intrudes as far as the mouth of the Hollandsche IJssel. The Rijn flow is usually large enough to minimize shipping costs. The lakes are ordinarily at their maximum levels, barring a few decades early in April, when they are being raised from their low winter levels. In fact, there is usually a need to discharge water from the lakes to prevent them from rising above their maximum levels. This excess water might as well be used to cool power plants or flush the the boezems of North Holland. Thus, the managerial strategy discussed here is only useful a fraction of the time.

9.6.2. Justification of Priority 1

The highest priority use of water is to meet the constraints imposed on MSDM. These constraints were discussed in Chap. 2. Many of them are natural laws, such as the constraints requiring conservation of water at each node, and the upper and lower rivers hydrologic equations. Others are limits on flows in canals due to pump capacities or canal dimensions and slopes. There is no question that these constraints must be obeyed.

Other constraints, however, may not be so rigid in practice as we consider them to be in MSDM. For example, water velocities in canals are often limited, say to 20 cm/s, in order to provide safety for shipping. We expect, however, that the velocity could be increased to 25 cm/s if the need were great enough, without substantial loss of safety. As another example, a minimum flow of 7 m^3/s is required in link 67 ZAAN, to maintain the water quality in the vicinity of Zaandam. We suspect that if there were sufficient reason, it could be at least temporarily reduced to 4 or 5 m^3/s without much damage.

In fact, we have ignored some limits on flows that the Dutch ordinarily observe, and even think of as requirements. For example, the Dutch managerial strategy is to maintain a minimum flow of 10 m^3/s in link 38 ORANJESL, in order to maintain the water quality in the city canals of Amsterdam. But if we imposed this as a constraint on MSDM, it could never have sent water along the Amsterdam-Rijnkanaal and ORANJESL to maintain the IJssel lakes at their minimum levels.

One sort of constraint we have imposed on MSDM is the level control requirement. There are a variety of reasons for the IJssel lakes level control requirements. The maximum level is determined by the needs of certain polders to drain excess water by gravity into the lakes, by the configurations of certain harbors (e.g., the heights of quays), and, for very high water levels--in excess of NAP + 0.7 m--by considerations of safety from flooding. The minimum level is determined by the elevations of various intake works for the water supply to a number of polders, the construction of certain locks and harbor works, and, for large reductions in depths, by the requirements of shipping. The present minimum and maximum levels are separated by 20 cm.

The requirement for level control of the IJssel lakes rarely plays a role in the managerial strategy. This is because there is rarely a long enough dry period to deplete the lakes. It requires a month or more during which the Rijn flow is abnormally low (less than approximately 1000 m^3/s) and the net evapotranspiration is unusually high (two or three millimeters water loss per day, at least) to drop the lakes from their maximum levels to the minimum.

When the lakes do reach their minimum levels, however, we have seen that the cost calculated by MSDM can be very large. In large part this extra cost is due to the effort to maintain the IJssel lakes level at the minimum by sending water along the Amsterdam-Rijnkanaal and ORANJESL. This tactic diverts the water from the Waal, where it would help shipping, and from the Nieuwe Maas, where it would combat the Rotterdam salt wedge (compare Figs. 9.4 and 9.6). Now we raise the question: is it realy worth this extra cost to maintain the IJssel lakes level? or might it be that the cost of letting the lakes drop below their minimum levels need be little more than the change in agricultural drought damage and sprinkling cost between Figs. 9.4 and 9.6? We think this question deserves some investigation.

We think the similar level control requirements for boezems and ditches should also be questioned. (Of course, the fact that a requirement is questioned does not imply that it will necessarily be changed. Upon investigation, one might discover that the costs of allowing the water level to drop as suggested here exceed the benefits. But we are not aware of any presently existing estimate of such costs.) Three reasons have been suggested to explain these requirements. For boezems, the main reason is that some industrial intake works would be exposed if the boezem level were allowed to drop more than a few centimeters. Presumably ships could have difficulty passing certain locks if water levels were not maintained.

The other two reasons pertain to ditches. We have heard it said that if the water level in the ditches were lowered (it cannot be raised without serious consequences, since it is barely below the ground level), the pressure difference between the saline ground water and the more-or-less fresh surface water would be increased. This would result in an increase in seepage, thus accelerating the salt contamination of the water supply. We question this reason, because in order to maintain the water levels in the midwest, for example, the Dutch are willing to admit highly saline water at Gouda. Again, in the part of North Brabant bordering the future fresh Zoommeer, we have been told that the Dutch will take in seawater, if necessary, to maintain ditch levels.

The third reason advanced to explain ditch level control has to do with preventing subsidence of the land. It is asserted that if the upper layer of the soil is allowed to dry out, it will settle. Once it has settled, it cannot be restored to its present level. If it does settle, the surface water level cannot be maintained at as high a level as formerly. This would cause the pressure difference between the saline ground water and the fresh surface water to increase permanently, resulting in an increased seepage rate.

In MSDM, level control requirements are specified as rigid constraints, so the position of level control at the top of the priority list is implicit in the model formulation. We have no evidence beyond the arguments given above that level control is really essential. We suspect that level control requirements, at least for boezems and ditches, could be temporarily relaxed at little cost in order to make water available for other uses. By temporary, we mean for one or two decades, or at most a month. We don't imagine that the land would settle much in such a short period. Moreover, we are not suggesting a very large decrease in the ditch level. Perhaps

10 cm would be enough to enable a region like the midwest to avoid admitting salty water at Gouda. (In fact, levels in the Friesland boezems varied by as much as 10 cm during the summer of 1976 [9.1].)

9.6.3. Justification of Priority 2

Priority 2 uses of water are sprinkling and nominal rates of flushing at locks with local salt intrusion. As we saw in Chap. 4, a small amount of flushing can markedly reduce the damage caused by local salt intrusion.

From Figs. 9.4 and 9.6 we see that agricultural drought damage plus sprinkling cost only rises when the lakes are at their minimum levels, and water must be diverted from sprinkling to satisfy the Priority 1 uses.

9.6.4. Justification of Priority 3

Priority 3 involves trading off water used to help shipping on the Waal and the IJssel against water used to combat the Rotterdam salt wedge. From the evidence of Figs. 9.4 and 9.6, the uses in this priority class are sacrificed to Priority 1 uses when necessary. In particular, water is sent north to the Amsterdam-Rijnkanaal to maintain the IJssel lakes at their minimum levels, even though doing so increases the cost to shipping and the damage to glasshouse agriculture in the midwest due to the Rotterdam salt wedge.

There are a few Priority 3 uses that are sacrificed to Priority 2 uses as well. For example, water is withdrawn from the IJssel at node 7 TWENMOND to supply sprinkling demands in the Northeast Highlands, when leaving it in the IJssel would reduce shipping costs. On the other hand, sprinkling demands in the Midwest and Utrecht region are sometimes not fully met (see Sec. 9.2), in order to limit the amount of saline water that must be taken from the Hollandsche IJssel, a sacrifice of a Priority 2 use in favor of a Priority 3 use. For the most part, however, Priorities 2 and 3 are not in substantial competition.

9.6.5. Justification of Priority 4

Priority 4 is the use of water to cool the power plants at AMSTEDAM, HAL+IJMU, and B'GUMMER using water from the IJssel lakes. MSDM could instead cool AMSTEDAM and HAL+IJMU using water taken from the Neder-Rijn and sent north along the Amsterdam-Rijnkanaal. The fact that it does not, and instead chooses to use IJssel lakes water, is our justification for placing this use of water below the Priority 3 uses. A comparison of the values of water for combating the Rotterdam salt wedge (Chap. 4) and for benefiting shipping (Chap. 7) with the value of water for cooling the power plants (Chap. 5) supports this ordering of the priorities.

9.6.6. Justification of Priority 5

Priority 5 is to fill the IJssel lakes to their maximum levels. In Sec. 9.4.4 we determined that water used for this purpose was less valuable, at least until the lakes dropped to levels near their minimum, than the Priority 4 use of cooling power plants.

9.6.7. Justification of Priority 6

Priority 6 uses of water include all uses not mentioned in the higher priority classes. One such use is flushing locks to prevent salt intrusion, either locally or globally, at rates in excess of the nominal rates. In Chap. 4 we established that, after a small amount of water has been devoted to this purpose, further increases in the flushing rate are all but worthless.

Other uses are flushing of boezems and ditches. The most important example of this is flushing the water stored at node 21 MIDWEST. Not unreasonably, MSDM chooses to do this only when the water to be used for flushing has a lower salt concentration than the water in storage. But this condition holds only when the Rijn flow is rather large, and hence plenty of water is available for uses of higher priorities. Clearly, improving the water quality in the midwest (especially reducing the salt concentration) would be worthwhile. This is demonstrated by the value of water used to combat the Rotterdam salt wedge. Thus, it is not the low value of flushing that relegates it to a low priority. Rather it is the coincidence that clean water for flushing is only available when water is abundant.

NOTES

1. One of our reviewers contends that not all of this cost difference is unnecesary. He argues that part of the extra damage under the RWS strategy is due to the extra diversion of water north along the Amsterdam-Rijnkanaal. This diversion reduces the flow in the Waal and Lek, thus increasing both the shipping losses and the salt damage to agriculture in the midwest. But it also provides additional water for cooling to power plants along the Noordzeekanaal. In effect, says this argument, the MSDM strategy appears to compare favorably with the RWS strategy largely because it ignores the costs of thermal pollution, and hence includes no measures for dealing with it.
   In fact, the reviewer is quite correct. If the MSDM strategy were modified to withdraw 20 m^3/s more water from the Waal at Tiel, or to reduce the flow in the Lek, in order to augment the flow in the Amsterdam-Rijnkanaal, the comparison between the strategies would be much more nearly equal. However, as will be

seen in MSDM runs described later, this is not the response of MSDM when thermal standards are introduced. Instead, MSDM responds by withdrawing more water from the IJssel lakes in order to provide cooling water in the Noordzeekanaal, and the flows in the Waal, Lek, and Amsterdam-Rijnkanaal remain as they are in the present MSDM strategy. Accordingly, the cost comparison is also (essentially) unchanged.

2. The future salt damage numbers we report are our rough estimates of the expected future damage. Depending on future circumstances (e.g., rainfall, flushing rates), the actual future damage might be much higher or much lower. While we have done no formal analysis of the amount of variation one should expect in the future salt damage, we feel it will be quite large.
3. Since we are unable to assign equivalent monetary costs to increases in pollutant concentrations, these are not, strictly speaking, total costs. They do include, however, all the monetizable components of cost that we have identified as being significantly related to water management issues.

REFERENCE

9.1. Rijkswaterstaat, Directie Friesland, Verslag Droge Zomer 1976 (Drought Summer 1976), Report No. ANP, April 1977.

Chapter 10

AN "OPTIMAL" MANAGERIAL STRATEGY FOR POSSIBLE FUTURE CONDITIONS

10.1. CASE DESCRIPTION

The cases we will consider in this chapter are intended to represent a possible future condition. They differ from the cases considered in Chap. 9 in three ways. First, we replace the 1976 shipping loss functions and cost functions at locks with the 1985 functions. These functions can be found in App. G of Vol. VA. Second, we replace the 1976 sprinkling scenario with the "maximum" sprinkling scenario, thus raising the area sprinkled from surface water from 1600 to 6000 square kilometers. Both the present and maximum sprinkling scenarios are described in Chap. 6 of this volume, with details to be found in App. F of Vol. VA.

Third, we add a number of the tactics described in Vol. XVI to the water management infrastructure. To expand the supply capacity to the Northern region, we increase the capacity of the Prinses Margrietkanaal (links 44 MARGKAN1 and 45 MARGKAN2) from 97 m^3/s to 136 m^3/s. We also increase the capacity of the Van Starkenborghkanaal (link 49 STABOKAN) from 16 m^3/s to 25 m^3/s. In Vol. XVI, the expansion of STABOKAN is tactic 4.4. However, the expansion of the Prinses Margrietkanaal does not appear as a tactic in Vol. XVI. At the time we constructed these cases, this was a tactic being considered in PAWN; but it has since been decided that the canal capacity is already 136 m^3/s, at least when the IJssel lakes level is high, and an expansion of the canal capacity is not considered to be a tactic in the final PAWN reports.

To expand the supply capacity to the Northeast Highlands region, we utilize the Twenthekanaal supply route described in Vol. XVI. In MSDM, this involves increasing the capacities of three links. First, we increase the upstream capacity of the Twenthekanaal (link 57 TWENKAN1) by 9 m^3/s, from -9.25 m^3/s to -18.25 m^3/s (on this link, negative flows are upstream). This capacity is the net of a pumping capacity of 20 m^3/s and a loss of 1.75 m^3/s through the shipping locks at Eefde. Second, we increase the capacity of link 56 OVIJKAN1, which represents part of the Overijsselsch Kanaal, from its present 5 m^3/s to 20 m^3/s. Third, we add 15 m^3/s of upstream pumping capacity on link 54 OMMERKAN. When netted against the lock loss of 0.2 m^3/s, this brings the lower bound on the flow in OMMERKAN to -14.8 m^3/s (again, negative flows are upstream).

The above changes to the infrastructure reduce the amount of water one will expect to find in the IJssel lakes. Expanding the supply capacity to the Northern region allows more extractions directly from the lakes, while expanding the supply capacity to the Northeast Highlands region allows more water to be taken from the IJssel River, leaving less to flow into the lakes. In addition, changing from the present to the maximum sprinkling scenario increases the extractions for sprinkling that are taken directly from the lakes. Accordingly, we add a tactic to increase the amount of lake water available to users, either by increasing the maximum level or decreasing the minimum. We have chosen here to decrease the minimum level by 10 cm. In Vol. XVI, this is tactic 10.3.2.

To expand the supply capacity to the Southeast Highlands and Delta region, capacities of three links are changed. The downstream capacity of link 72 ZUIDWLM2 is increased from 6 to 10 m^3/s (tactic 9.5 in Vol. XVI). The upstream pumping capacity on link 73 WESNVERT is increased by 5 m^3/s, changing the capacity on that link from -2.4 m^3/s to -7.4 m^3/s (tactic 9.11 in Vol. XVI). These two changes increase the capacity to extract water from the upstream part of the Maas, at nodes 27 MAASLOO and 28 BORN+PAN. The third tactic adds 7 m^3/s of upstream pumping capacity on link 76 WILHEKAN (tactic 9.12 in Vol. XVI). When netted against a minimum lock loss of 0.3 m^3/s, this changes the capacity of the link to -6.7 m^3/s. This tactic permits water to be brought from the lower part of the Maas at node 33 GERTRUID, up the Donge (link 77 DONGE) to node 36 OSTRHOUT, and thence via the Wilhelminakanaal (link 76 WILHEKAN) to node 35 HELMOND.

Finally, we include two tactics to reduce the damage caused by the Rotterdam salt wedge. We build a pipeline from the Maas (actually from the reservoirs in the Biesbosh, which are filled with Maas water) to Delfland (tactic 7.7 in Vol. XVI). Maas water is much less saline than the Rijn water that now supplies Delfland, so this tactic essentially eliminates salt damage to the glasshouse crops grown in Delfland. The value of these crops is two-thirds of the value of all glasshouse crops in the midwest, so this tactic can cause a two-thirds reduction in the potential damage from the salt wedge. In MSDM we represent this tactic by modifying the cost function associated with the salt concentration at node 21 MIDWEST. See Chap. 4 for details.

Finally, we construct a groin in the Nieuwe Waterweg (tactic 7.9 in Vol. XVI). According to the discussion in Vol. XVI, this tactic has the same effect on the salt wedge increasing the Nieuwe Maas flow by 10 m^3/s, and the Oude Maas flow by 20 m^3/s. As described in Chap. 4, the salt wedge is represented in MSDM by two functions of the Nieuwe Maas flow. To implement this tactic in MSDM, we simply translate these functions by 10 m^3/s.

These tactics were chosen because screening (Vol. II) showed that they all might have average annual benefits in excess of their costs. There are additional tactics that meet this criteria, which we did not include here either because they cannot be represented in MSDM (due to the degree of aggregation) or because they compete directly with tactics that were included. The final list given here is called "MAXTACS."

It is unfortunate that lists of tactics called MAXTACS are also given in Vol. II and Vol. XVI, and that the lists given there are not the same as the list presented here. Part of the reason for the differences is that the other MAXTACS lists include the tactics that we exclude here because they cannot be represented in the MSDM network. But there are other differences in the lists, in the capacities of links and in which tactics appear. These differences arise because the MSDM "MAXTACS" list is a preliminary one, which has been somewhat revised since its use in MSDM. To remind the reader that we have not used the latest version of MAXTACS in the cases discussed in this chapter, we will always refer to MSDM's "MAXTACS" in quotation marks. We should note that "MAXTACS" and MAXTACS are not very different, and hence our results should be indicative of the effect of implementing either list.

10.2. STRATEGY DESCRIPTION

In this section we compare the optimal strategies found by MSDM for two cases. Both cases have a Rijn flow of 1000 m^3/s, and a Maas flow of 31 m^3/s, of which 22.4 m^3/s is reserved for Belgium. In both cases, we consider a decade with no rain and 66 mm of open water evaporation. The IJssel lakes and the weir ponds of the Maas start the decade at their maximum levels. But one of the cases, called the Present (1976) Case in the tables that follow, has the 1976 shipping cost functions, the 1976 sprinkling scenario, and the present infrastructure (with only a fresh Zoommeer added). In fact, this is the same case that served as the nominal case in Tables 9.9 through 9.20. The other case, called the Possible Future Case in the following tables, has the 1985 shipping cost functions, the maximum sprinkling scenario, and an infrastructure to which has been added not only a fresh Zoommeer, but all the tactics in the "MAXTACS" list as well. We shall call this infrastructure the "MAXTACS" infrastructure.

10.2.1. Overall Strategy Costs

Table 10.1 shows the overall costs of the two strategies broken down into major components. Note that under present conditions, the shipping cost and drought damage plus sprinkling cost components are much smaller, and the damage from salt is much larger, than under future conditions. The difference in the shipping cost is largely due to the change from the 1976 low water shipping cost functions to the 1985 versions, which are much larger. This reflects an evolution of the shipping fleet toward larger vessels, which are affected by low water more than smaller vessels, and a reduction in excess shipping capacity. The change in the damage from salt is, of course, due almost entirely to the replacement of Rijn water by Maas water for the supply of glasshouse crops in Delfland.

At first sight, it seems wrong that the future case should have a higher drought damage plus sprinkling cost than the present case.

Table 10.1

COMPARISON OF MSDM STRATEGIES FOR PRESENT AND FUTURE CONDITIONS: OVERALL STRATEGY COSTS

Case Description

Pollutant Penalties: None
Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum
Power Plant Inventory: 1985

Cost Components (Dflm/dec)	Present (1976) Condition	Possible Future Condition
Shipping	11.57	34.19
Agriculture		
Damage from Salt	181.06	76.24
Drought Damage Plus Sprinkling Cost	36.88	70.52
Thermal	71.58	71.58
Other Costs	1.00	1.53
Total Costs	302.09	254.06

Sprinkling is supposed to reduce drought damage more than the sprinkling itself will cost. In MSDM, however, the only drought damage considered is damage to crops sprinkled from surface water. The area which is sprinkled in the future case but not in the present case certainly would suffer drought damage in both cases. In fact more damage would occur in the present case. Unfortunately, we are unable to estimate the extent of this damage in the present case. To do so would require additional runs of MSDM which cannot be done now that the PAWN project has terminated.

The thermal cost component (i.e., fuel cost for electric power generation) is the same, since we used the 1985 power plant inventory and ignored the three-degree standard in both cases. The remaining costs are not significant.

10.2.2. Strategies for the National Waterways

In this section we discuss both strategies in the national waterways. The reader will find it helpful to refer frequently to Fig. 9.1.

10.2.2.1. The Waal, Neder-Rijn, and Delta. Table 10.2 shows the two strategies in the national part of the network. In both strategies the weir at Driel is nearly closed, allowing a Neder-Rijn flow little more than the minimum of 25 m^3/s in the present case, and only the minimum in the future case. Because there are larger

Table 10.2

COMPARISON OF MSDM STRATEGIES FOR PRESENT AND FUTURE CONDITIONS: STRATEGIES FOR THE NATIONAL WATERWAYS

Case Description

Pollutant Penalties: None
Rijn Flow: 1000.00 m³/s
Maas Flow: 31.00 m³/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum
Power Plant Inventory: 1985

Strategy Description	Present (1976) Case	Possible Future Case
Upper Rivers Flows (m³/s)		
Waal	767.91	771.10
Neder-Rijn	37.65	25.00
IJssel	185.36	187.69
Lower Rivers Flows (m³/s)		
Nieuwe Maas	224.19	172.50
Oude Maas	474.85	448.19
Nieuwe Waterweg	699.04	620.69
Volkerak + Haringvlietsluizen	40.11	59.53
Amsterdam-Rijnkanaal Flows (m³/s)		
Withdrawal at Tiel	2.00	45.71
North from A-R Mond	9.49	9.58
Lek	5.00	32.56
Noordzeekanaal Flows (m³/s)		
Input from IJssel Lakes	6.36	6.36
Discharge to North Sea	20.00	20.00
IJssel Lakes Level Change (cm)	-4.74	-9.24
Maas Weir Pond Level Changes (cm)		
MAASLOO	-47.51	-60.00
BORN+PAN	0.0	-60.00
LINNE	0.0	0.0
ROER+BEL	0.0	0.0
SAMGRALI	0.0	0.0
Maas Flows (m³/s)		
JULCANL1	6.65	6.50
MAAS3	14.19	10.01
AMER1 (Discharge into Lower Rivers)	18.06	-11.77

extractions from node 6 UPRIVER, the lesser Neder-Rijn flow in the future case does not imply larger flows on the Waal or the IJssel. Here, the two cases are almost identical.

Farther west, the present case withdraws the minimum possible amount of 2 m^3/s from the Waal at Tiel, while the future case withdraws more than 45 m^3/s. Both send essentially the same amount north on the Amsterdam-Rijnkanaal. So, in spite of larger extractions for sprinkling at node 12 A-R.MOND, the future case has a higher flow on the Lek.

But the higher Lek flow and equal Waal flow does not imply a higher flow in the Nieuwe Maas. As we shall see, the future case extracts more water from the Hollandsche IJssel at Gouda, so the flow on link 61 HOLIJSEL is larger. Also, increased demands for sprinkling have increased the flow through the Volkerak and reduced the flow from the Maas into the Lower Rivers. All three factors act to reduce the flow in the Nieuwe Maas to 172.5 m^3/s.

10.2.2.2. The Amsterdam-Rijnkanaal and the Nordzeekanaal. The two strategies are essentially identical in their treatment of the Amsterdam-Rijnkanaal and the Noordzeekanaal. Since we have discussed this aspect of the present strategy earlier, we will omit the discussion here.

10.2.2.3. The IJssel lakes. Because of increased extractions from the IJssel River and the IJssel lakes, the lake levels are lower, by 45 mm, in the future case than in the present case. It is interesting to note how large a flow of water must have been diverted from the lakes to account for the difference. The amount is 105 m^3/s, all due to increased demands for sprinkling.

10.2.2.4. The Maas. The present strategy lowers the level of only the first weir pond (node 27 MAASLOO) to the point where its depth is the same as the Waal depth. The future strategy, by contrast, lowers the levels of the first two weir ponds (MAASLOO and node 28 BORN+PAN). (Although the level changes are different in the two cases, both are lowering the levels until the Waal and the Maas have the same shipping depth. In the future case, the Waal is shallower because the withdrawal at Tiel is larger.) In the present case, lowering the level of one weir pond reduces the cost of conserving water on the Julianakanaal, and provides ample water to operate the locks on link 13 MAAS3 in the least costly possible way. But because the future case has increased extractions for sprinkling from the Maas, lowering the level of one weir pond does not supply enough water to minimize the cost of conserving water on the link MAAS3. Indeed, even lowering both weir ponds still leaves the flow in MAAS3 low enough that some conservation measures are required (see Chap. 7, and App. G of Vol. VA).

In the present case, the Maas discharges more than 18 m^3/s into the lower rivers part of the network. In the future case this flow is reversed, and nearly 12 m^3/s is taken out of lower rivers. This water flows via links 19 AMER1, 77 DONGE, and 76 WILHEKAN, satisfying sprinkling demands at nodes 33 GERTRUID, 36 OSTRHOUT, and 35 HELMOND.

10.2.3. Strategies in the Regions

In this section, we describe the present and future strategies in the regional parts of the network. Again, the reader will find it helpful to refer frequently to Fig. 9.1.

10.2.3.1. Southeast Highlands and Delta Region. As shown in Table 10.3, the future case supplies much more water for sprinkling to this region than does the present case. The water comes from increases in extractions from the Upper Maas and the Delta Lake, and by changing a discharge into an extraction from the Lower Maas.

But even the large increases in water supply to this region are not sufficient to meet the entire demand. Nearly one-third of the demand is not met in the future case, whereas the demand in the present case is so modest that even the paltry supply made available by the present infrastructure can readily satisfy it.

10.2.3.2. The Northeast Highlands Region. Table 10.4 shows the two strategies in the Northeast Highlands region. In the present case, some of the modest sprinkling demands are unmet, even though a substantial amount of water is being discharged from the region into the IJssel River. Earlier we cited this as evidence that the infrastructure in this region was unable to distribute water very well. In the future case, the distribution of water has been improved. Now, no water is allowed to escape to the IJssel.

However, the supply to the region is insufficient to meet the full demand in the future case. More than one-third of the demand is not satisfied. Of course, the demands are exceptionally high in this case because of the high open water evaporation. In more normal decades, the full demand could be satisfied.

10.2.3.3. The Northern Region. Table 10.5 shows the two strategies for the Northern region. Again, the major difference is that the demand for sprinkling has increased enormously from the present to the future case. In this region, it is possible to satisfy nearly all of the increased demand by increasing extractions from the IJssel lakes. The small sprinkling cutbacks occur at node 3 GRONETAL because the new capacity of link 49 STABOKAN (25 m^3/s) is not quite sufficient.

10.2.3.4. The North Holland Region. As Table 10.6 shows, the supply capacity to North Holland is barely adequate in the present case, and grossly inadequate in the future case. Only 30 m^3/s can be extracted from the IJssel lakes. Of this amount, a minimum of 7 m^3/s in link 42 ZAAN is required to maintain the water quality near

Table 10.3

COMPARISON OF MSDM STRATEGIES FOR PRESENT AND FUTURE CONDITIONS: SOUTHEAST HIGHLANDS AND DELTA REGION

Case Description

Pollutant Penalties: None
Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum
Power Plant Inventory: 1985

Strategy Description (m^3/s)	Present (1976) Case	Possible Future Case
Water for Sprinkling	3.63	25.57
Other Consumption	0.33	0.33
Extraction from Upper Maas	4.84	10.78
Extraction from Delta Lake	1.49	7.89
Discharge to Lower Maas	2.38	-7.23
Target Sprinkling Water	3.63	37.56
Cutbacks in Sprinkling	0.0	11.99
Regional Agricultural Damage (Dflm)		
Damage from Salt	0.0	0.0
Drought Damage Plus Sprinkling Cost	2.64	13.05

Table 10.4

COMPARISON OF MSDM STRATEGIES FOR PRESENT AND FUTURE CONDITIONS: NORTHEAST HIGHLANDS REGION

Case Description

Pollutant Penalties: None
Rijn Flow: 1000.00 m^3/s
Maas Flow: 31.00 m^3/s
Net Evaporation: 66.00 mm
Initial Lake Levels: Maximum
Power Plant Inventory: 1985

Strategy Description (m^3/s)	Present (1976) Case	Possible Future Case
Water for Sprinkling	4.69	23.10
Other Consumption	-4.56	-4.56
Overijsselsche Vecht	1.80	1.80
Extractions from IJssel	5.89	16.94
Discharge to IJssel	7.36	0.0
Discharge to North	0.20	0.20
Target Sprinkling Water	5.09	36.64
Cutbacks in Sprinkling	0.40	13.54
Regional Agricultural Damage (Dflm)		
Damage from Salt	0.10	0.08
Drought Damage Plus Sprinkling Cost	2.37	12.23

Table 10.5

COMPARISON OF MSDM STRATEGIES FOR PRESENT AND FUTURE CONDITIONS: NORTHERN REGION

Case Description	
Pollutant Penalties:	None
Rijn Flow:	1000.00 m^3/s
Maas Flow:	31.00 m^3/s
Net Evaporation:	66.00 mm
Initial Lake Levels:	Maximum
Power Plant Inventory:	1985

Strategy Description (m^3/s)	Present (1976) Case	Possible Future Case
Water for Sprinkling	15.13	72.01
Other Consumption	41.57	41.57
Extraction from IJssel Lakes	62.49	119.38
From Northeast Highlands	0.20	0.20
Discharges to Waddenzee	6.00	6.00
Target Sprinkling Water	15.13	74.65
Cutbacks in Sprinkling	0.0	2.64
Regional Agricultural Damage (Dflm)		
Damage from Salt	1.03	1.01
Drought Damage Plus Sprinkling Cost	4.53	10.84

Table 10.6

COMPARISON OF MSDM STRATEGIES FOR PRESENT AND FUTURE CONDITIONS: NORTH HOLLAND REGION

Case Description	
Pollutant Penalties:	None
Rijn Flow:	1000.00 m^3/s
Maas Flow:	31.00 m^3/s
Net Evaporation:	66.00 mm
Initial Lake Levels:	Maximum
Power Plant Inventory:	1985

Strategy Description (m^3/s)	Present (1976) Case	Possible Future Case
Water for Sprinkling	6.33	7.20
Other Consumption	13.80	13.80
Extractions from IJssel Lakes	30.00	30.00
Discharges to Waddenzee	2.87	2.00
Discharges to Noordzeekanaal	7.00	7.00
Target Sprinkling Water	6.33	12.86
Cutbacks in Sprinkling	0.0	5.60
Regional Agricultural Damage (Dflm)		
Damage from Salt	1.26	1.28
Drought Damage Plus Sprinkling Cost	2.67	4.20

Zaandam, and 2 m^3/s in link 43 NOHOLKAN to operate the locks at Den Helder. In this very dry decade, almost 14 m^3/s is needed for level control or is lost due to evaporation from the boezems, ditches, and lakes of North Holland. This leaves only 7.2 m^3/s to partially satisfy sprinkling demands almost twice that large in the future case.

This example underscores some questions we raised in Sec. 9.6. We wonder whether it would not be preferable in the future case to reduce the flow in link 42 ZAAN to 5 or even 4 m^3/s, and to allow the boezems and ditches of North Holland to drop by 10 cm. By these means it would be possible to provide an additional 4 or 5 m^3/s for sprinkling during one decade.

10.2.3.5. The Midwest and Utrecht Region. Table 10.7 shows the two strategies in the Midwest and Utrecht region. Again a major difference is the demand for sprinkling. Most of the increase in demand in the future case is met by increasing the extraction from the Hollandsche IJssel to its full capacity of 35 m^3/s. Even this is not enough to fully satisfy the demand; a cutback of 2.23 m^3/s is still necessary.

Table 10.7

COMPARISON OF MSDM STRATEGIES FOR PRESENT AND FUTURE CONDITIONS: MIDWEST AND UTRECHT REGION

Case Description		
Pollutant Penalties: None		
Rijn Flow: 1000.00 m^3/s		
Maas Flow: 31.00 m^3/s		
Net Evaporation: 66.00 mm		
Initial Lake Levels: Maximum		
Power Plant Inventory: 1985		
Strategy Description (m^3/s)	Present (1976) Case	Possible Future Case
Water for Sprinkling	6.29	22.63
Other Consumption	32.72	32.72
Extraction from Hollandsche IJssel	18.73	35.00
Extraction from Amsterdam-Rijnkanaal	24.10	24.10
Extraction from IJssel Lakes	10.00	10.00
Discharges to Amsterdam-Rijnkanaal	7.60	7.52
Discharges to Noordzeekanaal	3.00	3.00
Discharges to North Sea	3.23	3.23
Target Sprinkling Water	7.98	24.86
Cutbacks in Sprinkling	1.69	2.23
Regional Agricultural Damage (Dflm)		
Damage from Salt	157.09	51.25
Drought Damage Plus Sprinkling Cost	3.42	3.74

The other major difference between the cases is the difference in damage from salt. This is due to the fact that the glasshouse crops in Delfland are not supplied by water from node 21 MIDWEST in the future case. Since these crops account for about two-thirds of the value of all glasshouse crops in this region, the salt damage naturally is reduced.

10.3. VARIATIONS IN RIVER FLOWS, RAINFALL, AND LAKE LEVELS

In this section we will consider the effect on the nominal MSDM strategy of varying the Rijn and Maas flows, the rainfall, and the starting amounts of water in the IJssel lakes and the weir ponds of the Maas. Except for these variations, the cases considered will be identical with the original case defined in Sec. 10.1. As in Sec. 9.4, we will consider variations in the Rijn flow at Lobith between 700 m^3/s and 1500 m^3/s, and variations in the Maas flow between 25 m^3/s and 41 m^3/s, in perfect correlation with the Rijn flow. In all cases, 22.4 m^3/s of Maas flow is reserved for Belgium.

10.3.1. Net Evaporation = +66 mm/dec, Initial Lake Levels = Maximum

We will first consider a series of cases which all start with the lake levels at their maxima, and which all have 66 mm of net open water evaporation. Only the Rijn and Maas flows will differ among the cases, as described above. Looking at Fig. 10.1, we see what effect variations in the river flows has on the optimal strategy. For Rijn flows above approximately 810 m^3/s, the weir at Driel is set to provide a Neder-Rijn flow between 25 m^3/s and 40 m^3/s. According to the upper rivers hydrologic equation, this dictates the flows in both the Waal and the IJssel. Farther west, the withdrawal at Tiel is at its minimum level of 2 m^3/s only for Rijn flows larger than 1450 m^3/s. For Rijn flows between 1450 m^3/s and 1150 m^3/s, the withdrawal at Tiel is approximately 18 m^3/s. As the Rijn flow decreases below 1150 m^3/s, the withdrawal at Tiel generally increases, until a Rijn flow of 810 is reached. Then the Neder-Rijn flow increases and the withdrawal at Tiel decreases, until the Rijn flow drops to 700 m^3/s. A constant 9.58 m^3/s is sent north from the Neder-Rijn on the Amsterdam-Rijnkanaal, regardless of the Rijn flow.

The flows described above are sufficient to determine the flows in the lower rivers part of the network, according to the lower rivers hydrological equation. For Rijn flows above 1150 m^3/s, the flows in the Nieuwe Maas, Oude Maas, and Nieuwe Waterweg all decrease steadily and uneventfully with decreasing Rijn flows. At a Rijn flow of 1150 m^3/s, however, the Nieuwe Maas flow reaches 197 m^3/s, the flow at which the Rotterdam salt wedge first begins to affect nodes 20 GOUDA and 21 MIDWEST. (The effect of the groin in the Nieuwe Waterweg is to change this critical flow from 207 m^3/s to 197 m^3/s.) At this point the optimal strategy recommends that the withdrawal at Tiel be increased to help combat the salt wedge.

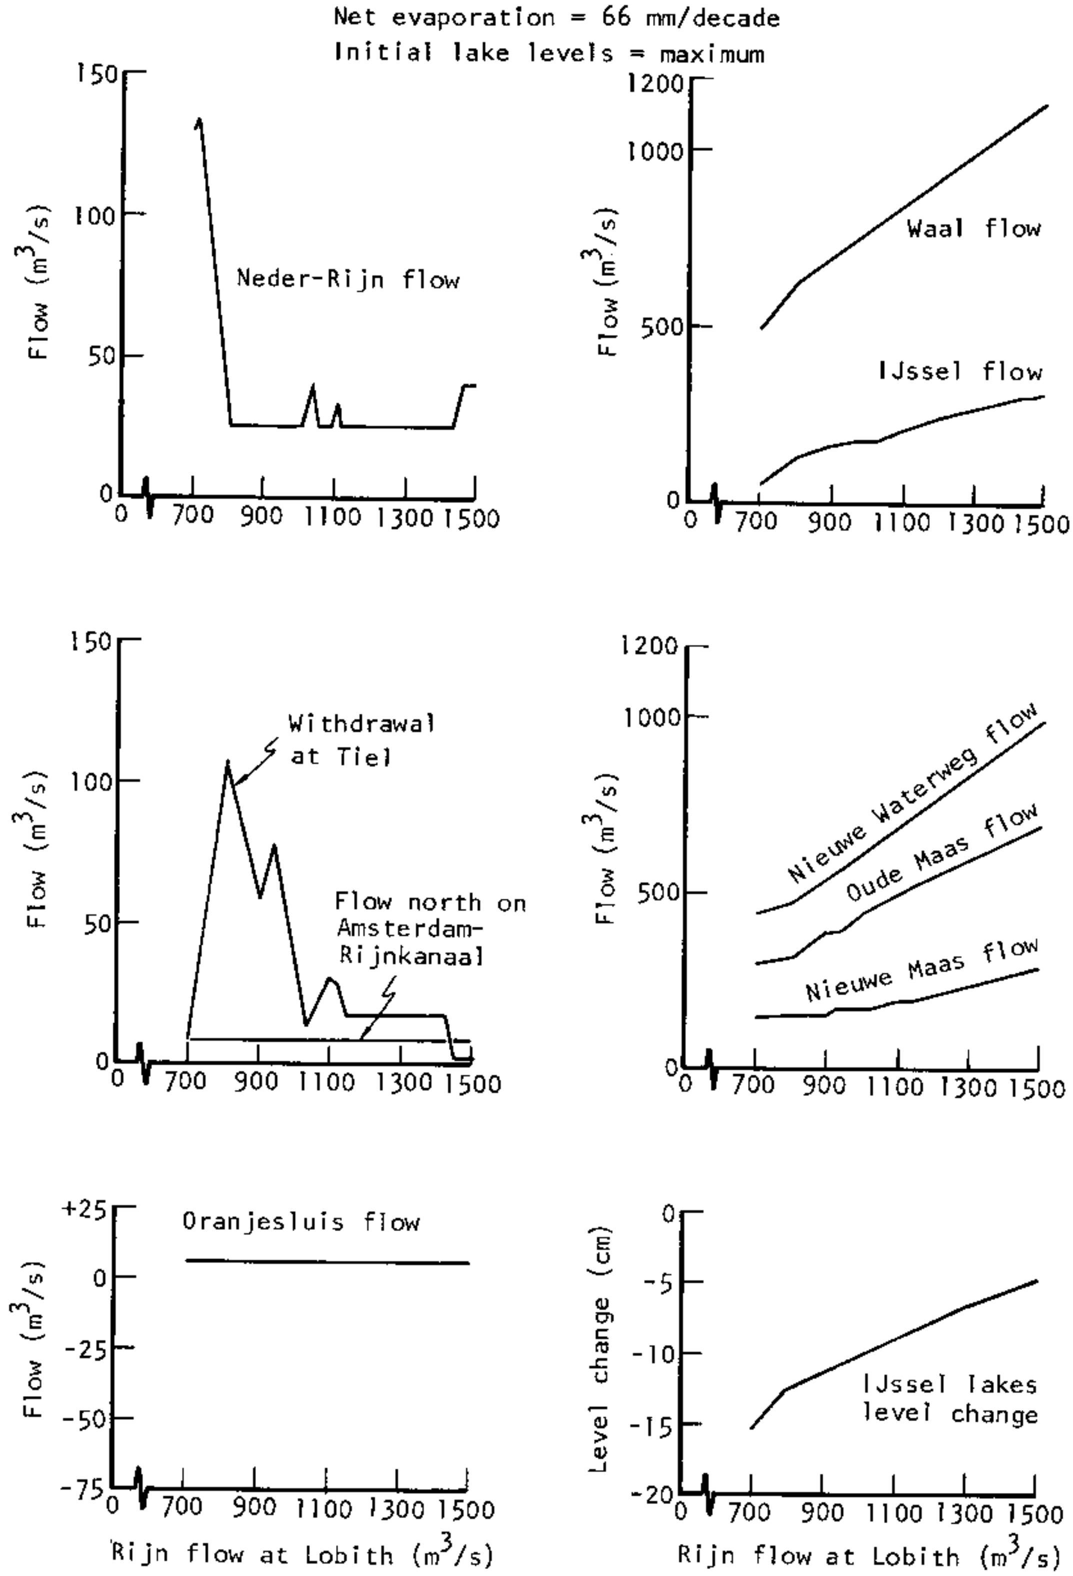

Fig. 10.1--MSDM strategy for "MAXTACS" infrastructure and maximum sprinkling scenario, case 1

Unlike the strategy in the cases considered in Chap. 9, the Nieuwe Maas flow is not maintained at its critical value. For the cases in this chapter, we replaced the 1976 shipping loss functions with the 1985 functions, which yield higher costs. Also, the infrastructure includes the Maas-Delfland pipeline, which reduces by two-thirds the value of glasshouse crops that can be affected by the Rotterdam salt wedge. Thus, for Rijn flows below 1100 m^3/s, the withdrawal at Tiel is reduced in order to protect shipping, and the salt wedge is allowed to affect the salt concentration at GOUDA and MIDWEST.

There are other ranges of Rijn flow in which the Nieuwe Maas flow is kept constant, at levels lower than 197 m^3/s. For Rijn flows between 1020 m^3/s and 940 m^3/s, the Nieuwe Maas flow is a constant 172.5 m^3/s. For Rijn flows between 910 m^3/s and 800 m^3/s, the Nieuwe Maas flow is 152.5 m^3/s. In both ranges, the Nieuwe Maas flow is maintained by increasing the withdrawal at Tiel as the Rijn flow decreases.

The Nieuwe Maas flows at which the strategy seems to "stick" are the flows at which the function describing the transportable component of the salt wedge changes slope. As described in App. D of Vol. VA, this function is represented as a piecewise linear function, with abrupt changes in slope where successive linear pieces are joined. By suitably smoothing this and other functions, this sort of "sticking" behavior could be eliminated.

In Fig. 10.1, we have also presented curves showing the flow of water out of the IJssel lakes via the Oranjesluis (link 38 ORANJESL), and the change in the level of the IJssel lakes. In this series of cases, the Oranjesluis flow is constant, and the IJssel lake levels drop steadily, as the Rijn flow decreases.

Figure 10.1 shows nothing about the strategy for the Maas. The levels of the first two weir ponds (nodes 27 MAASLOO and 28 BORN+PAN) are allowed to drop as far as necessary to provide a total flow of 6.5 m^3/s in the Julianakanaal (link 10 JULCANL1), or to the level of the Waal, whichever is the larger reduction in level. If this drop in the level of MAASLOO is not sufficient to maintain a flow of 6.5 m^3/s in the Julianakanaal, extractions from MAASLOO to the Southeast Highlands and Delta region are reduced. In the 1985 shipping scenario, which we have assumed in all of the cases in this chapter, the minimum flow possible on the Julianakanaal is 6.5 m^3/s. In the cases described in this section, it is never necessary or desirable to reduce the levels of other weir ponds.

Figure 10.2 shows the cost components and the total costs of the strategies described in this section. As we pointed out in Sec. 9.2.1, the magnitudes of the cost components are not important. Rather, one should consider the relative changes in the different cost components. The figure shows that the thermal cost is constant (since we have ignored the three-degree standard in these cases) and the drought damage plus sprinkling cost is very nearly constant. The

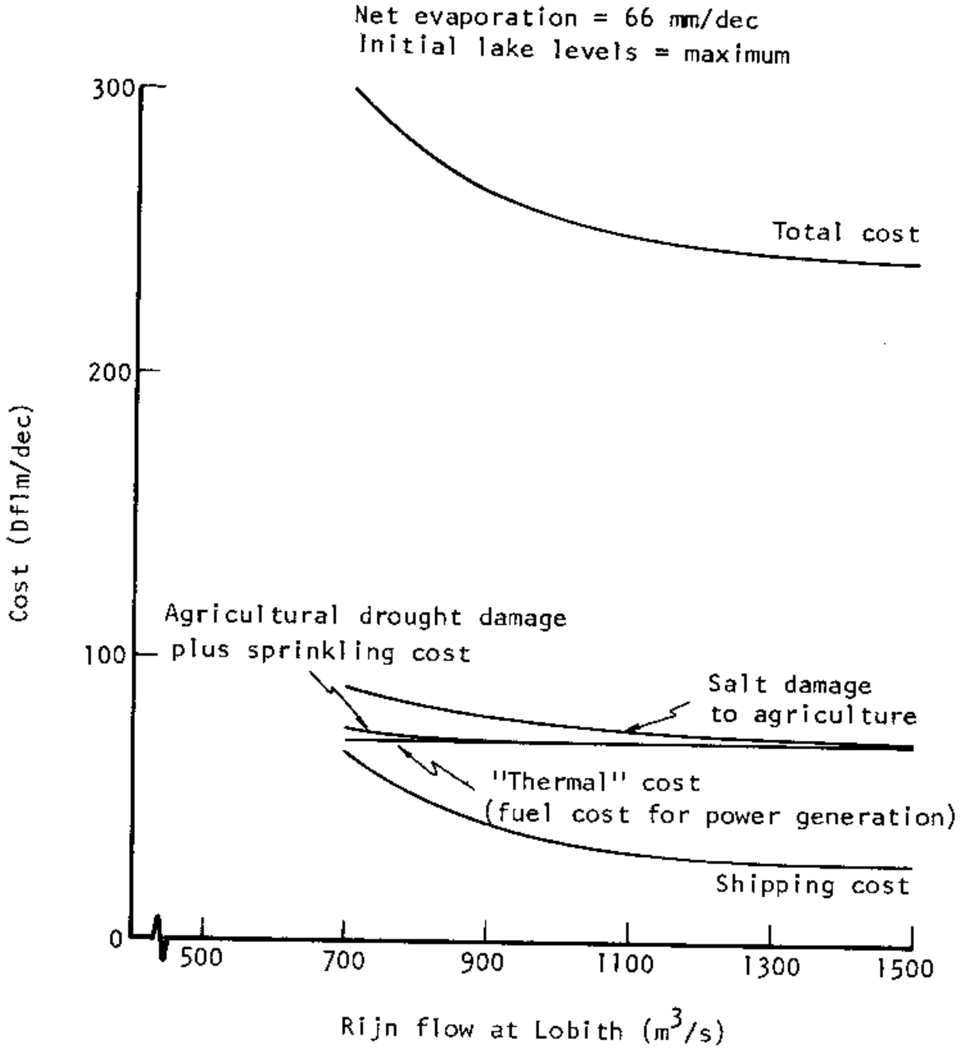

Fig. 10.2--Cost of MSDM strategy for "MAXIACS" infrastructure and maximum sprinkling scenario, case 1

other two components, the shipping cost and the salt damage to agriculture, do show substantial increases as the Rijn flow decreases. The reasons are the same as those given in Sec. 9.4.1 above.

10.3.2. Net Evaporation = +66 mm/dec, Initial Lake Levels = Minimum

Next we will consider a series of cases which start with all lake levels at their minima, and which again have 66 mm of net open water evaporation. Once again, only the Rijn and Maas flows will differ among the cases. Looking at Fig. 10.3, we see what effect the variation in Rijn flow has on the optimal strategy. The strategy sets the Neder-Rijn flow to its minimum allowed rate of 25 m^3/s, withdraws at Tiel at the maximum allowed rate of 120 m^3/s, and sends as much water north on the Amsterdam-Rijnkanaal as is necessary to provide the maximum flow of 73 m^3/s into the IJssel lakes via the Oranjesluis. This is a strategy to maximize the amount of water sent to the IJssel lakes, and it is mandated by the fact that, even with a Rijn flow of 1500 m^3/s, the IJssel lakes level cannot be maintained.

The strategy for the Maas is not shown in Fig. 10.3. When the weir ponds start the decade at their minimum level, there is a minimum Maas flow necessary of approximately 30 m^3/s. This provides 22.4 m^3/s for Belgium, and minimum flows of 6.5, 1, and 0.35 m^3/s on links 10 JULCANL1, 11 MAAS1, and 72 ZUIDWLM2 respectively. Any Maas flow in excess of this minimum is used first to increase the flow in ZUIDWLM2, up to a maximum of 10 m^3/s. Further increases, which are beyond the range of Maas flows considered in these cases, would be used to increase the flow in JULCANL1. But if the Maas flow is smaller than 30 m^3/s, the problem as originally formulated is infeasible. To circumvent this infeasibility, we allowed MSDM to divert some of the 22.4 m^3/s that would ordinarily be sent to Belgium, but imposed a large penalty on each cubic meter diverted.

Figure 10.4 shows the cost components and total costs for the strategies discussed in this section. Once again, the thermal cost component is constant, reflecting the fact that these strategies ignore the three-degree thermal standard. Drought damage plus sprinkling cost is also constant, reflecting the fact that even at a Rijn flow of 1500 m^3/s, it is impossible to maintain the IJssel lakes level. Thus, all sprinkling supplied from the IJssel lakes has already been eliminated at all Rijn flows we consider. Note that this cost component is much larger in Fig. 10.4 than in Fig. 10.2. The other two components increase significantly as the Rijn flow decreases. The increase in shipping cost is not much greater than we observed in Fig. 10.2. And we observe a much larger increase in salt damage to agriculture, reflecting the fact that at low Rijn flows we have no water to spare from maintaining the lake level, and cannot therefore combat the salt wedge.

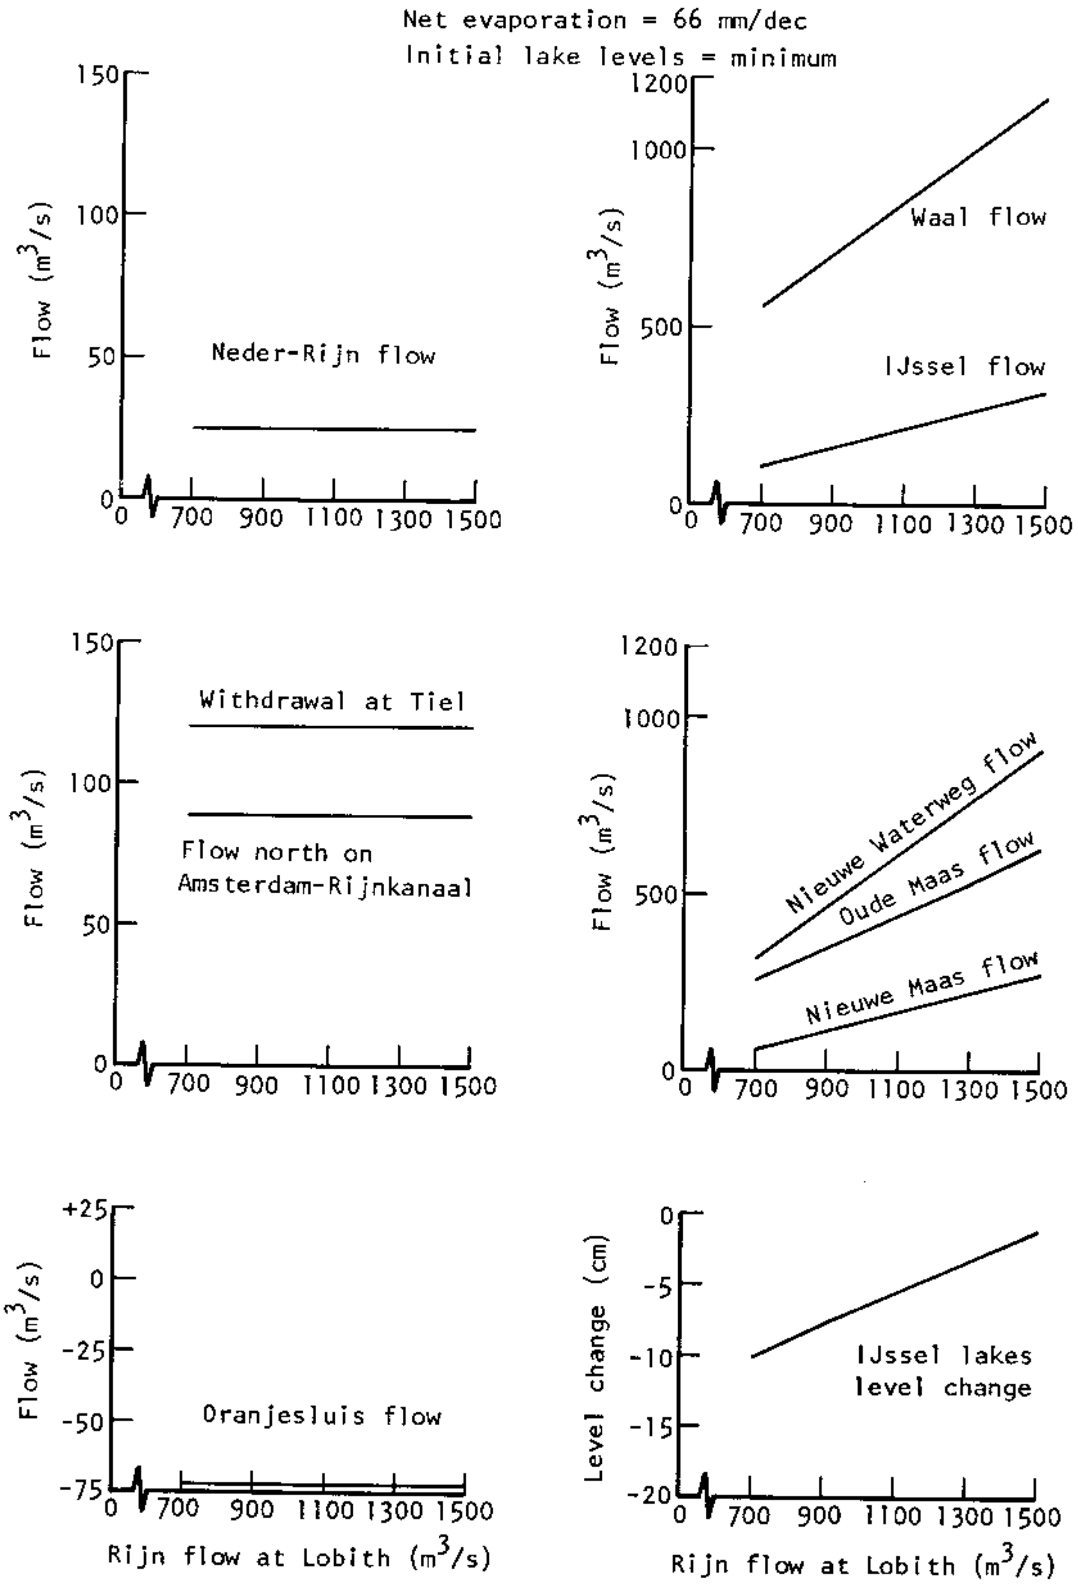
Net evaporation = 66 mm/dec
Initial lake levels = minimum

Fig. 10.3--MSDM strategy for "MAXIACS" infrastructure and maximum sprinkling scenario, case 2

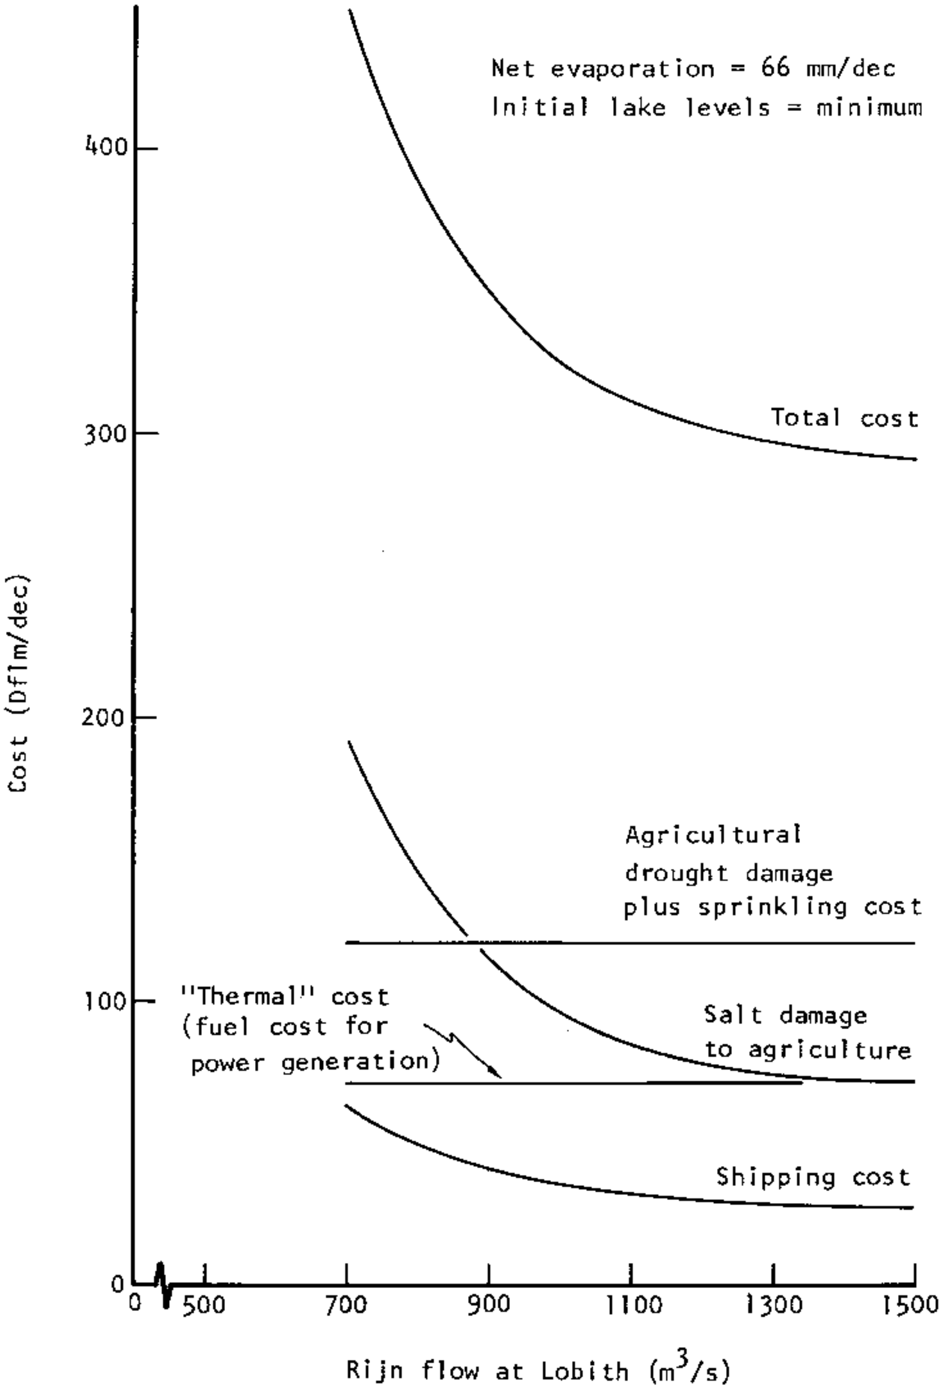

Fig. 10.4--Cost of MSDM strategy for "MAXTACS" infrastructure and maximum sprinkling scenario, case 2

10.3.3. Net Evaporation = -11 mm/dec, Initial Lake Levels = Any

Finally, we will consider a series of cases which have a net open water evaporation of -11 mm--i.e., which have more rain than evaporation. Again, only the Rijn and Maas flows differ among the cases. For these cases the starting lake levels are irrelevant, since as Fig. 10.5 shows, the strategy supplies enough water to the IJssel lakes to raise their level substantially at all Rijn flows. (In practice, if some of this water were not needed in the IJssel lakes, it could be discharged instead through the Noordzeekanaal into the North Sea.)

The strategies for these cases are easy to describe. The weir at Driel is kept fully closed, allowing only the minimum flow of 25 m^3/s in the Neder-Rijn. The withdrawal at Tiel and the flow north on the Amsterdam-Rijnkanaal are also at their minima. This strategy gives rise to Waal and IJssel flows, and flows in the lower river branches, as shown in the figure. On the Maas, the strategy is the same as for the cases of Sec. 10.3.1.

Figure 10.6 shows the cost components and total cost for these cases. Again, the thermal cost is constant, since we have ignored the three-degree thermal standard. In these cases, drought damage plus sprinkling cost is constant (and very small) also. Because the decade is so wet (net evaporation = -11 mm), sprinkling demands are zero, and drought damage is suffered only because we assumed that the farmers' soil was quite dry at the start of the decade.

The salt damage to agriculture rises rather little as the Rijn flow decreases. As in the cases considered in Sec. 9.4, there is a break in this cost component at a Rijn flow of approximately 1040 m^3/s. Above this flow, the chloride concentration of the Rijn is less than or equal to 300 mg/l, which is the initial chloride concentration at node 21 MIDWEST. Accordingly, the optimal strategy is to flush MIDWEST with as much Rijn water as possible. For Rijn flows less than 1040 m^3/s, the chloride concentration of the Rijn exceeds 300 mg/l, and all intake of Rijn water ceases.

The shipping cost also increases as the Rijn flow decreases. However, the increase is much smaller than in Figs. 10.2 and 10.4. Because there is no intake of Rijn water into MIDWEST at low Rijn flows, there is no need in the present cases to combat the Rotterdam salt wedge. The flow out of MIDWEST on the Hollandsche IJssel (link 61 HOLIJSEL) prevents the salt wedge from affecting GOUDA and MIDWEST regardless of how far it intrudes in the Nieuwe Maas, or its effect on the chloride concentration at IJSLMOND. Accordingly, the effect of the strategies for these cases is to minimize shipping costs for each Rijn flow.

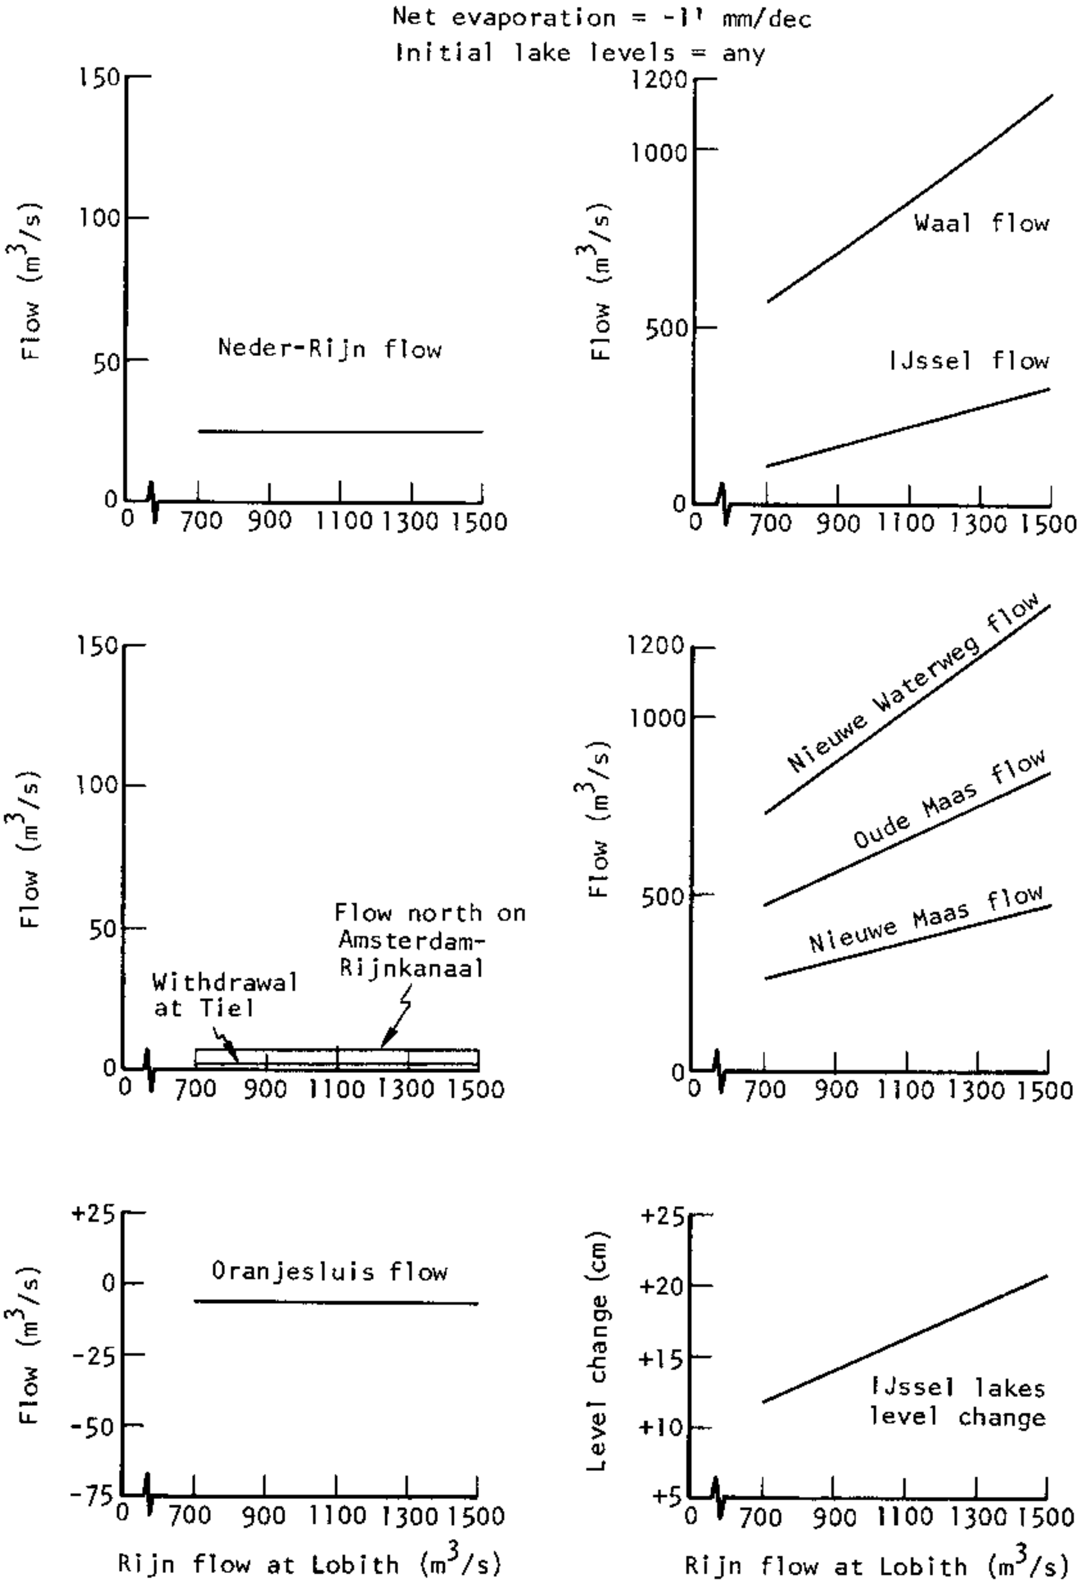

Fig. 10.5--MSDM strategy for "MAXTACS" infrastructure and maximum sprinkling scenario, case 3

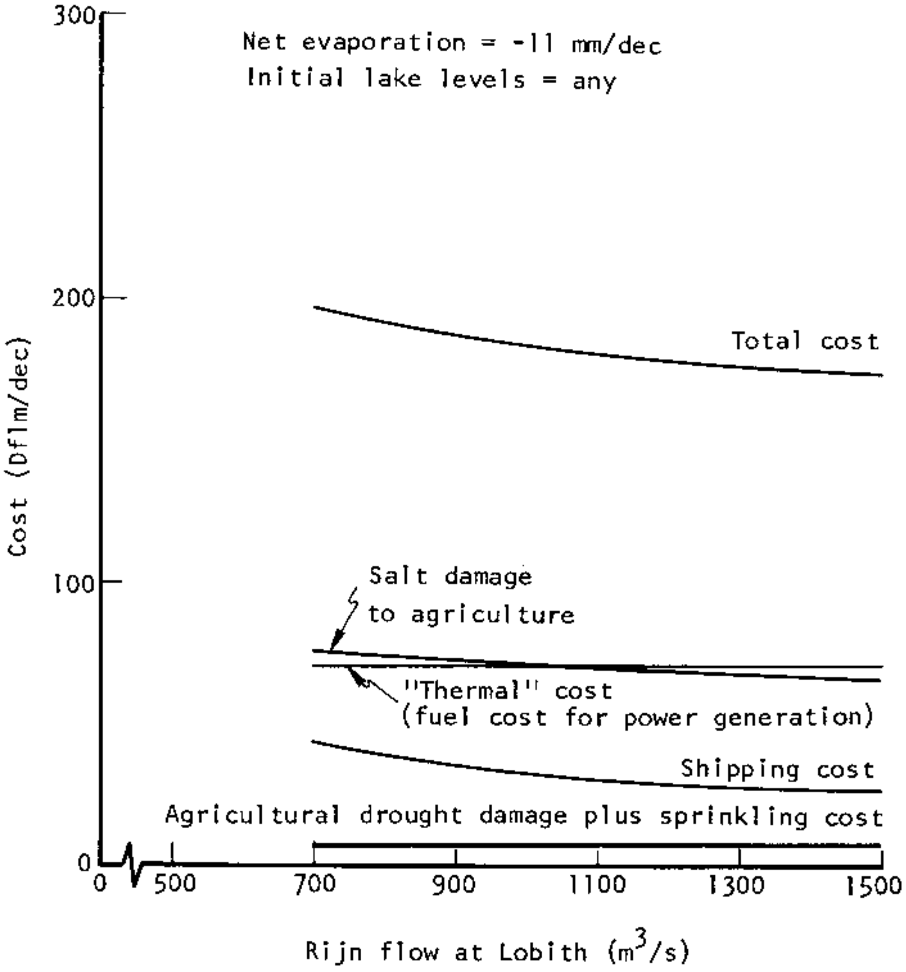

Fig. 10.6--Cost of MSDM strategy for "MAXTACS" infrastructure and maximum sprinkling scenario, case 3

10.4. THE EFFECT OF THE THREE-DEGREE THERMAL STANDARD

In this section, we consider the effect of imposing the three-degree thermal standard on every node of the network. We will compare two managerial strategies which can meet the standard. The first is the nominal strategy, with the proviso added that the power plants must shift the generating load in such a way as to meet the standards. As in Sec. 9.5, we consider only load shifting away from nodes 2 B'GUMMER, 15 AMSTEDAM, and 16 HAL+IJMU. The cost of shifting the load from B'GUMMER is 0.17 Dflm/dec. This is less than the cost increase found in Sec. 9.5 because, with the addition of "MAXTACS" to the infrastructure, and in particular the increase in the capacity of link 49 STABOKAN, the nominal strategy sends more water past B'GUMMER. Thus a smaller part of the generating load must be shifted. The cost of shifting the load away from AMSTEDAM and HAL+IJMU is 0.90 Dflm/dec, the same as before. Thus, the total cost of the shifts is 1.07 Dflm in each decade in which they occur.

The second managerial strategy we consider is the Thermal strategy discussed in Sec. 9.3. As in Sec. 9.5, the Thermal strategy has larger extractions from the IJssel lakes than the nominal strategy. With the increase in the capacity of the Van Starkenborghkanaal (link 49 STABOKAN) from 16 m^3/s to 25 m^3/s, the Thermal strategy needs to send only an additional 15 m^3/s past B'GUMMER to make load shifting unnecessary. Thus, the Thermal strategy differs from the nominal strategy by sending 105 m^3/s more along the Noordzeekanaal past AMSTEDAM amd HAL+IJMU, and 14 m^3/s more along the Margrietkanaal past B'GUMMER.

10.4.1. Average Annual Cost Difference for Power Generation

We first calculate the average annual benefits of changing from the nominal strategy to the Thermal strategy. As discussed in Sec. 9.5, the nominal strategy will fail to cool the power plants whenever the lakes are below capacity, while the Thermal strategy will fail only when they have dropped to their minimum levels. When a strategy must cut back the flow of cooling water from the IJssel lakes by 119 m^3/s, the fuel cost for generating electric power is increased by 1.07 Dflm/dec. To estimate the frequency with which the two strategies suffer this cost, we have carried out a simulation using the Distribution Model applied to 44 years of river flow and rainfall data. This simulation was exactly parallel to the simulation described in Sec. 9.5, except that it assumed the "MAXTACS" infrastructure, the 1985 shipping cost functions, and the maximum sprinkling scenario. In the 44 simulated years, the IJssel lakes dropped below capacity in 31 decades.

To estimate the frequency with which the Thermal strategy fails to cool the power plants, we modified the results of the 44-year simulation in the same fashion as discussed in Sec. 9.5, but using a change in IJssel lake extractions of 119 m^3/s instead of 128

m^3/s. In the "MAXTACS" infrastructure, the minimum lake level is NAP - 50 cm. Using this procedure, we estimate that there would have been five decades during the 44 years of simulation with the Thermal strategy during which the lakes would have been at their minimum levels, and hence when the extractions of cooling water from the IJssel lakes could not be continued. Thus, adopting the Thermal strategy reduces the frequency with which the fuel cost of power generation rises by 1.07 Dflm/dec, from 31 to 5 decades per 44 years. The Thermal strategy thus conveys an average annual benefit from this source of 0.632 Dflm/yr, when compared to the nominal strategy.

10.4.2. Average Annual Cost Difference from Lake Depletion

The other difference in the average annual costs of the strategies occurs because of the different probabilities that the lakes will drop to their minimum levels. A comparison of Figs. 10.2 and 10.4 shows that the strategy cost is between 50 and 150 Dflm/dec higher if there is no water available in storage (and if the Rijn flow is low and the decade has little rain as well).

Figure 10.7 shows the minimum summer levels that we estimate would have been achieved under the nominal and Thermal strategies. The minimum IJssel lakes level permitted in these cases is NAP - 50 cm. To estimate how frequently the lakes might reach this level under the nominal strategy, it is necessary to extrapolate the upper curve of the figure. A rough guess is that the lakes might reach their minimum level with an annual probability of 0.004. From the lower curve of Fig. 10.7, we can estimate the analogous probability for the Thermal strategy to be 0.072.

But this cost is only borne if the decade has little rain and low Rijn flow. We estimate that the cost becomes zero if the Rijn flow exceeds 1600 m^3/s, even when the rainfall is zero and the evaporation is as high as 66 mm/decade. By the same method used in Sec. 9.5.2, we estimate that the probability that a summer decade would have a nonzero incremental cost with the "MAXTACS" infrastructure and maximum sprinkling scenario is 0.054. From our earlier discussion, we know that the Thermal strategy increases the probability of depleting the IJssel lakes from 0.004 to 0.072 per year, or a decrease in the annual probability of 0.068. We have estimated that the incremental cost that will be suffered if the lakes reach their minimum levels, and the following decade is dry enough, will be between 50 and 150 Dflm. Multiplying these three figures together provides an estimate of the average annual difference in the cost due to lake depletion of the two strategies. That estimate is between 0.184 and 0.551 Dflm/yr.

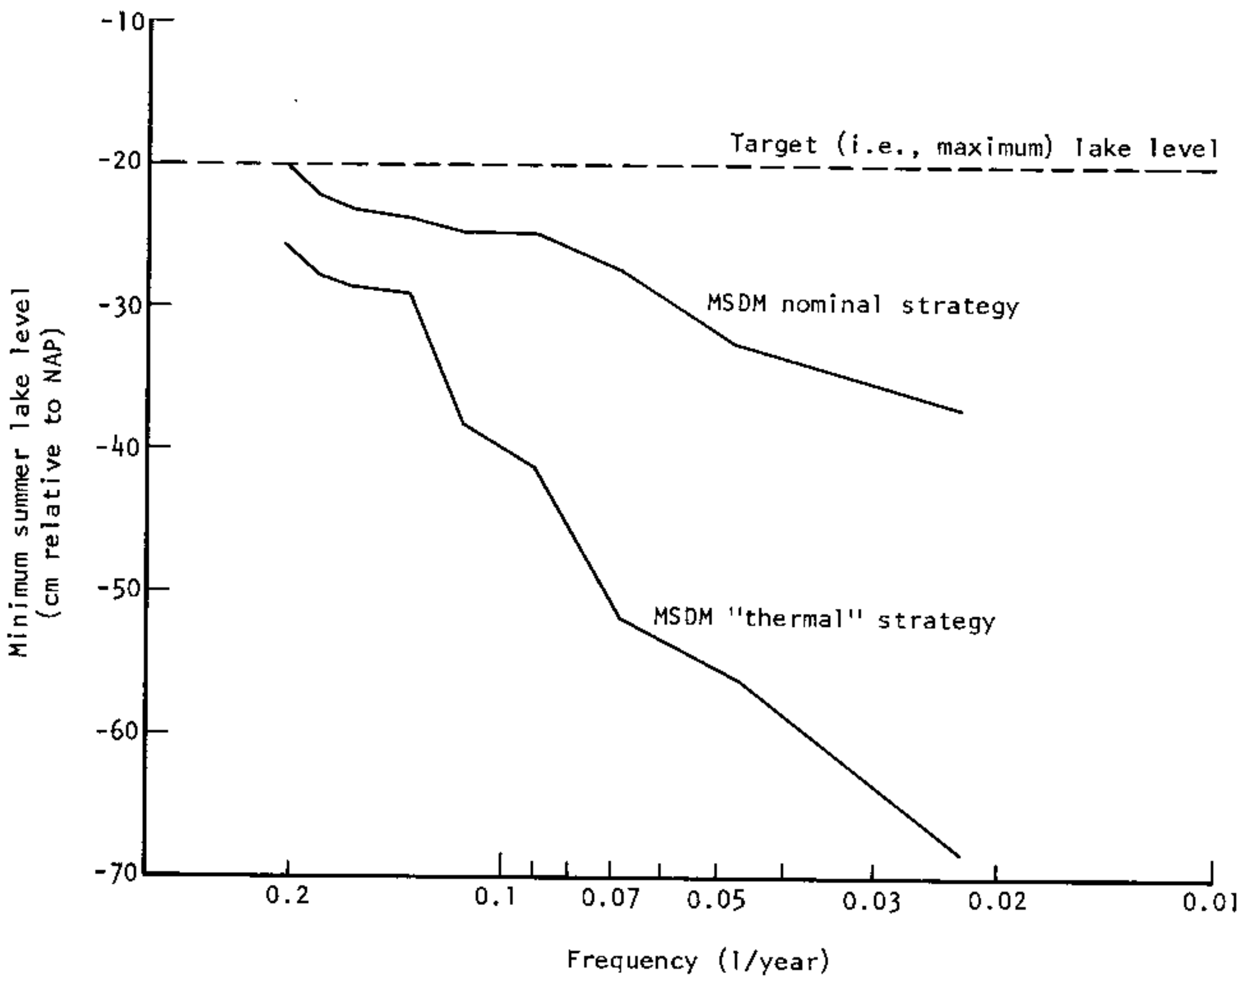

Fig. 10.7--Frequencies of low summer lake levels: "MAXTACS" infrastructure and maximum sprinkling scenario

10.4.3. The Thermal Strategy Is Preferable

In Sec. 10.4.1 we estimated that the Thermal strategy would save an average of 0.632 Dflm/yr in fuel costs for power generation, as compared with the nominal strategy. In Sec. 10.4.2 we estimated that this savings would be partially offset by no more than 0.551 Dflm/yr due to the increased probability of depleting the IJssel lakes. These numbers are too close together for us to conclude that either strategy should be preferred to the other.

Using numbers from Secs. 10.4.1 and 10.4.2, we can form a better estimate of the value of water stored in the IJssel lakes (the insurance functions mentioned in Sec. 9.1). There were 26 decades in 44 years during which the Thermal strategy extracted 119 m^3/s more from the IJssel lakes than did the nominal strategy. This amounts to a difference of 61.60 million cubic meters per year left in storage, at a difference in cost between 0.184 and 0.551 Dflm/yr. The average value of this water is therefore between 0.0030 and 0.0090 Dfl/m^3. This is two or three times larger than the value we estimated for the present infrastructure and 1976 sprinkling scenario. Of course, what we have estimated here is only an average value. Surely, the value of having an additional cubic meter in storage will be much smaller if the lakes are nearly full, and larger if they are nearly empty. This is another way of saying that it is probably worth while to use lake water for cooling power plants when the lake levels are high, but that it will be preferable to conserve the water once the lake levels have dropped sufficiently.

Finally, we note that all results in this section have assumed that power plants at AMSTEDAM, HAL+IJMU, and B'GUMMER will be subject to the three-degree thermal standard. This standard does not presently apply at these locations, and it appears unlikely that it will be soon applied. If the alternative seven-degree standard is applied at these locations, it is approximately true that the amount of water needed for optimum cooling is three-sevenths as large as the amount needed for optimum cooling to a three-degree standard. Thus, a "Seven-Degree Thermal" strategy would have considerably less chance of depleting the IJssel lakes, and would definitely be preferable to the nominal strategy.

10.5. DESCRIBING STRATEGIES IN TERMS OF PRIORITIES

In Sec. 9.6 we developed a description of the optimal strategy in terms of priorities attatched to various uses of water. High priority uses are to be satisfied first. Once no more can be done to satisfy a high priority use, a lower priority use can be considered. The same can be done for the strategy for the "MAXTACS" infrastructure and maximum sprinkling scenario. In spite of the changes in the infrastructure, sprinkling demand, and shipping cost functions, the priority order remains the same. From high priority to low, it is the following:

Priority 1: Supply level control requirements for boezems and lakes, and meet all other constraints on MSDM.

Priority 2: Supply water to farmers for irrigating their crops. Also, establish certain nominal flushing rates for locks at which salt intrusion causes damage to crops grown locally.

Priority 3: Trade off shipping losses due to low water on the Waal and the IJssel, the dredging cost (or other sedimentation cost) due to withdrawals at Tiel, and salt damage to agriculture due to the Rotterdam salt wedge, by simultaneously adjusting the weir at Driel and withdrawals at Tiel.

Priority 4: Use water from the IJssel lakes for cooling the power plants at AMSTEDAM, HAL+IJMU, and B'GUMMER to whatever thermal standard has been set for them.

Priority 5: Raise the IJssel lakes to their maximum levels to meet the possible future needs for water.

Priority 6: Use water for flushing boezems and ditches, and for raising the flushing rates of locks with local salt intrusion above the nominal rates established in Priority 2.

The differences between the infrastructure, sprinkling demands and shipping cost functions make the application of the priorities somewhat different from the earlier cases. The most visible difference is due to the change in the relative values of water to shipping, and for combating the salt wedge. The 1985 shipping cost functions have increased the former, while the Maas-Delfland pipeline has reduced the latter. Thus, the trade-off called for in Priority 3 is now resolved much more in favor of shipping. Another visible difference is the almost fourfold increase in sprinkling demands. Because they are so much larger, they now compete more obviously with other uses.

Chapter 11

OBSERVATIONS, CONCLUSIONS, AND RECOMMENDATIONS

11.1. CRITICAL NOTES ON THE MANAGERIAL STRATEGY DESIGN MODEL

In a study as large as the one reported in this volume, it is inevitable that there will be soft spots in the analysis, sins of omission and commission, and things that the author wishes he could do over. That the study has imperfections, however, does not prevent it from having strengths. (Those who feel otherwise can write their own report.) We have tried to apprise the reader of both strengths and weaknesses in the body of the volume. Here we would like to summarize them, and to discuss their overall impact on the utility of MSDM.

11.1.1. MSDM Soft Spots

First, MSDM relies heavily on parts of the PAWN analysis reported in other volumes. Chief among these are the shipping analysis (Vol. IX), the salt wedge analysis (Vol. XIX), the work done on thermal pollution from power plants (Vol. XV), and the work on agriculture (Vol. XII). Each of these separate analyses has its own assumptions and approximations, which MSDM has necessarily inherited. Furthermore, the assumptions of these separate analyses may not always be mutually consistent. To this, MSDM adds its own approximations and assumptions. It uses a highly aggregated network, especially in its representation of regional and local waterways, and in its representation of the IJssel lakes. The formulation of the network in the lower-rivers area precludes the use of MSDM for wet decades, as do some of the assumptions made concerning agriculture.

Second, the present version of MSDM is expensive and cumbersome to use. A large amount of data is required, and data preparation itself requires the use of other, complex models (notably DISTAG to generate the aggregate plot descriptions). To generate the MSDM strategy for even a single decade requires a certain amount of manual intervention, and it usually requires 4-6 runs before the strategy is satisfactory. Since a MSDM run of a single decade costs about ten U.S. dollars, the cost of generating a satisfactory MSDM strategy for even one decade can be considerable. Moreover, there is no fixed recipe for guiding the intervention; it took quite some time before we could do it effectively, and others would require a similar learning period.

Partly because of the expense of using MSDM, and partly due to lack of time, we did fewer sensitivity analyses than we would have preferred. We would like to have varied the initial pollutant concentrations at nodes with storage, for example, and the initial amount of water and salt in the root zones of our aggregate plots.

We also wanted to investigate the effect of varying some of the cost functions, notably the shipping cost functions and the salt damage function for crops grown under glass in the Midwest. Such sensitivity analyses would have helped us determine how robust are our results and conclusions.

11.1.2. Strengths of the MSDM Approach

On the other hand, we believe our approach has strengths that more than compensate for its shortcomings. We have adopted a methodology that attempts to optimize the management of the water system, rather than relying solely on a methodology that simulates its behavior. (The latter approach is taken by the Distribution Model described in Vol. XI. MSDM and the Distribution Model can be considered as complementary.) MSDM, therefore, made it necessary for us to try to understand the relations between water managerial actions and their eventual consequences to the ultimate water users, even when the eventual consequences were remote in time or space from the original action, and even when the relationships were tenuous, uncertain, or convoluted. Much of this effort at understanding can be avoided by taking a simulation approach; but the insights gained are also avoided.

One insight concerns the value of taking a systems view of Dutch water management. The Dutch have historically operated sections of their water management infrastructure separately. For example, the Upper Rivers District of the RWS is responsible for controlling the weir at Driel and withdrawals at Tiel, but the position of the Rotterdam salt wedge, and any consequent damage to midwest agriculture, are outside its geographical area of responsibility. When water is abundant, the historic Dutch approach results in at most trivial losses, but as this volume shows, when water is scarce, a more coordinated approach would yield significant benefits.

Moreover, we are able to offer constructive suggestions regarding how to achieve better coordination. These suggestions rely on quantitative estimates of the benefits to be obtained through devoting a particular unit of water to any given use. Historically, Dutch water users have expressed their demands for water as requirements, without assessing the harm that would befall them if some portion of their demands were not satisfied, or the benefits that would acrue if they were given extra water. As shown in these pages, such quantitative estimates of harm or benefit can guide water management decisions on a nationwide basis.

Nor is it necessary to restrict this approach to only those uses of water for which a monetary benefit can be estimated. In this volume we have taken into account thermal pollution (which water management decisions can strongly influence), salt pollution (which is influenced by water management decisions at only a few locations), and three other pollutants (which are hardly affected at all by water management). To be sure, quantitative methods are much easier to apply, and more of them are available, when all the impacts of every decision can be monetized and directly compared. But we have shown that quantitative analysis is not helpless in the face of nonmonetizable impacts. And we have shown that water management in the Netherlands has some influence on, and therefore some responsibility regarding, some forms of pollution.

While conducting this investigation, we were forced to the realization that much of the harm or benefit resulting from a water management decision may be deferred. If agriculture is deprived of sprinkling water, the crops will nevertheless thrive until the soil has had time to dry out. Depending on the weather, this may take one, two, or more decades, or it may never happen. Similarly, if the Rotterdam salt wedge is allowed to contaminate the water supply of the Midwest, the damage to agriculture will continue to accumulate for many decades. Water stored in the IJssel lakes may help many users avoid future harm. In this investigation, we have developed estimates for deferred harm or benefits involving a number of water uses, thus illustrating several techniques that might be applied in future extensions of the PAWN study.

Thus, MSDM has both strengths and weaknesses. On the one hand it points a new way to look at water managerial decisions, one that explicitly and formaly considers trade-offs among different users, as well as suggesting a number of concrete, specific conclusions. On the other hand, MSDM is complex, expensive and difficult to use, and necessarily relies on numerous assumptions and approximations. If we could do it all over--not that we would jump at the chance, considering how tired we are--we would build a much simpler model. It would be easier and less expensive to use. More extensive sensitivity analyses would be possible, the better to test the importance of the more questionable assumptions. The priority list presented in Chap. 9 offers an approach for building a "simple MSDM." Indeed, in a later section of this chapter, we describe the start we have made in this direction.

In the sections that follow, we discuss the more important conclusions we have drawn from the analyses in the previous chapters. In our opinion, the net effect of the soft spots listed here should not be to discount these conclusions, but rather to engender caution; the conclusions should not be accepted, nor the recommendations implemented, without further scrutiny. Furthermore, even if some of the specific conclusions or recommendations were to be proven wrong, we feel that the development and use of MSDM has given us insights that could have been gained in no other way.

11.2. "DILUTION IS NO SOLUTION FOR POLLUTION"

We have observed that managerial strategies have little effect on water quality. Most pollutants seem to be distributed rather uniformly throughout the MSDM network, leaving little opportunity to reduce pollutant concentrations by redistributing water. But almost the only effect of managerial tactics is to redistribute water. We expect this observation is accurate for the major lakes and waterways of the Netherlands, since the MSDM representation of these parts of the water management infrastructure is reasonably faithful. But the regional parts of the MSDM network are highly aggregate representations of the Dutch regional waterways, and the observation can only be tentative there.

Thermal pollution in the Noordzeekanaal and at Bergumermeer may be an exception. Managerial tactics are available to provide cooling water to these locations if necessary. Under current practice, the Dutch maintain high enough flows at these locations to keep excess temperatures below seven degrees. Our results suggest that if the three-degree standard were imposed here, it would be worthwhile to increase the flows rather than to cut back power generation (and hence heat discharges) at these locations.

For other water quality problems, especially in regional waters, we recommend further study. We have carried out a preliminary study of problems of algae growth (usually identified as the <u>eutrophication</u> problem) in a number of Dutch lakes (see Vols. III and VI). From these efforts we conclude that regions should be studied individually and individual solutions should be developed. A single solution applied nationwide is unlikely to be optimal.

11.3. PRIORITIES FOR WATER USES

We have described the optimal managerial strategy by assigning a priority to each use of water. Water should be devoted to a use only after everything possible has been done to satisfy higher uses. The optimal priority list is the following:

Priority 1: Supply level control requirements for boezems and lakes, meet requirements for drinking water and other vital uses, and satisfy all other constraints on MSDM.

Priority 2: Supply water to farmers for irrigating their crops. Also, establish certain nominal flushing rates for locks at which salt intrusion causes damage to crops grown locally.

Priority 3: Trade off shipping losses due to low water on the Waal and the IJssel, the dredging cost (or other sedimentation cost) due to withdrawals at Tiel, and salt damage to agriculture due to the Rotterdam salt wedge, by simultaneously adjusting the Neder-Rijn flow (by adjusting the weir at Driel) and withdrawals at Tiel.

Priority 4: Use water from the IJssel lakes for cooling the power plants at AMSTEDAM, HAL+IJMU, and B'GUMMER to whatever thermal standard has been set for them.

Priority 5: Raise the IJssel lakes to their maximum levels to meet the possible future needs for water.

Priority 6: Use water for flushing boezems and ditches, and for raising the flushing rates of locks with local salt intrusion above the nominal rates established in Priority 2.

This ordering of water uses corresponds roughly to the relative economic values that water has in the various uses, which we discussed in earlier chapters. It cannot precisely reflect the values of water, since they change with circumstances. However, with relatively few exceptions, the uses of water belong in the same order regardless of circumstances.

11.4. POTENTIAL SAVINGS FROM AN IMPROVED MANAGERIAL STRATEGY

One of our results is that, in some unusually dry decades, considerable savings could result from implementing a new managerial strategy in place of the present Dutch practice. We think the new strategy ought to resemble the MSDM strategy described in Chap. 9. However, if decades with significant savings occur too rarely, the savings will be trivial on the average, and the change of strategy might not be worthwhile. To estimate how large the average annual savings might be, we tried to implement the MSDM strategy in the Distribution Model, and to simulate its performance over several selected years, instead of over a single decade.

But MSDM as it stands is far too cumbersome, and requires too much manual intervention, to permit incorporating it into the Distribution Model. Accordingly, we wrote a subroutine for the Distribution Model that implemented a simplified version of it. To use the subroutine in any decade, the Distribution Model first calculates the extractions from each node and the flows in each link of the network, using the RWS (or any other desired) managerial strategy. The SIMPLE MSDM strategy then modifies the flows in only some of the links, namely those shown in Fig. 11.1. Flows in all other links and extractions from all nodes are left unchanged.

SIMPLE MSDM imposes several constraints on the changes it makes to flows in the affected links. It requires that the flow in the Lek and the discharge from the Noordzeekanaal into the North Sea must exceed specified minima, which are nominally set to 5 m^3/s and 40 m^3/s, respectively. It also requires that the water level in the IJssel lakes exceed a specified minimum at the end of the decade, which is set to NAP - 40 cm for cases with the present

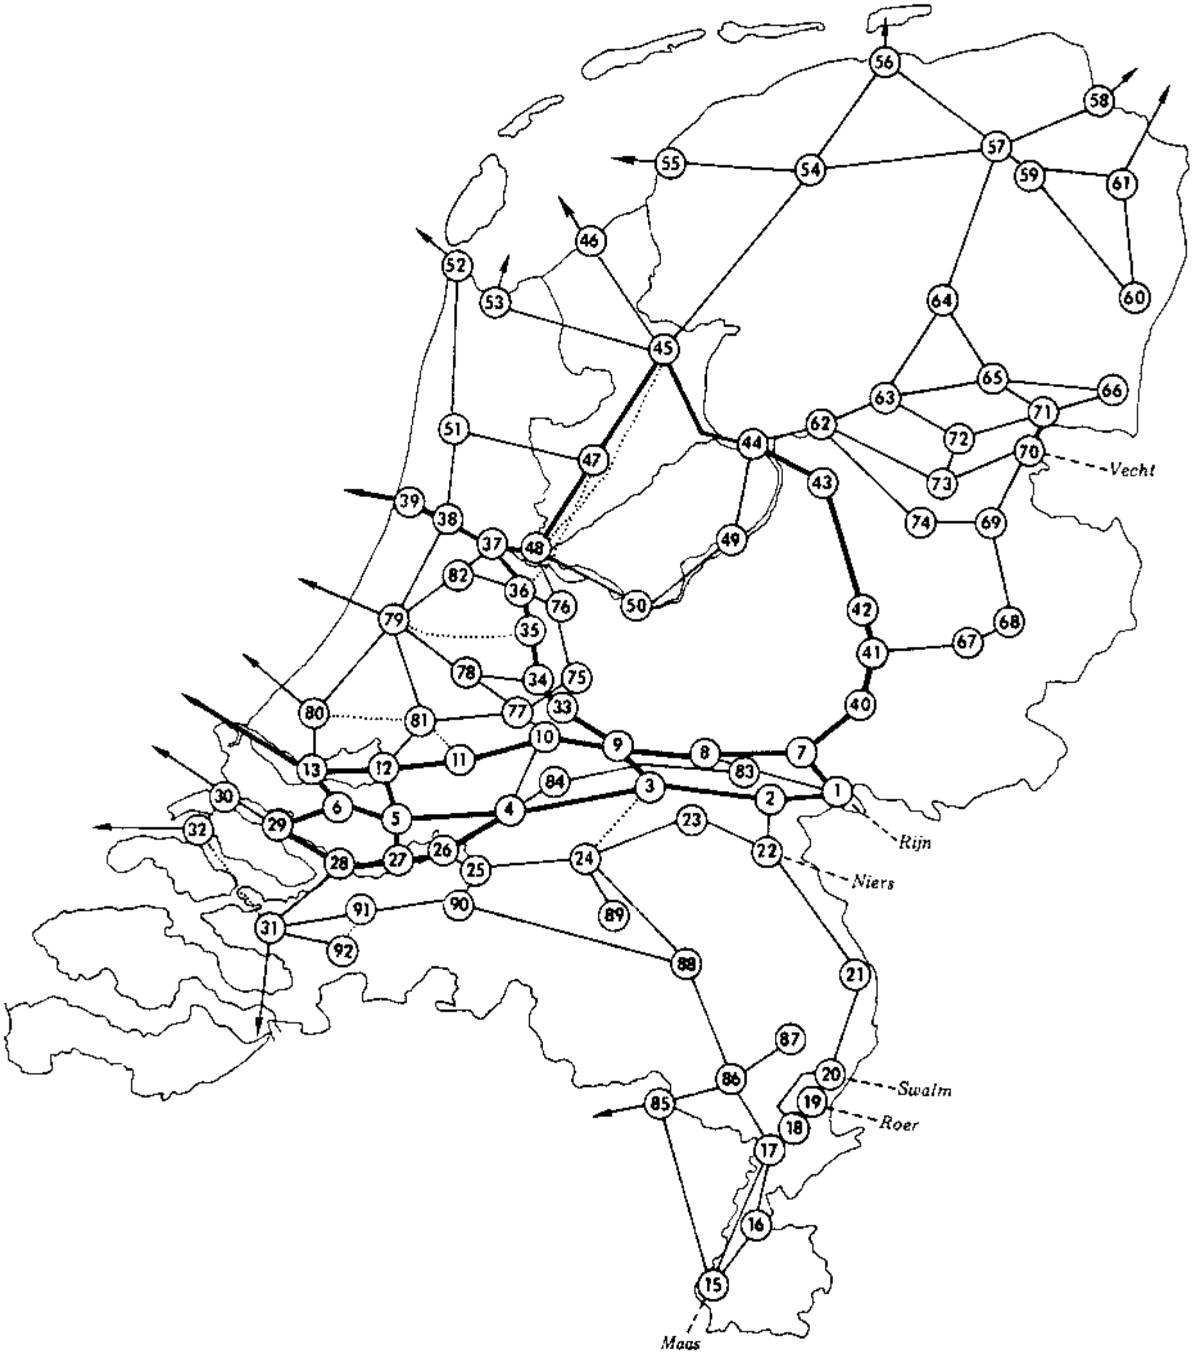

Fig. 11.1--Distribution model network links affected by the simple MSDM strategy

infrastructure, and NAP - 50 cm for cases with the "MAXTACS" infrastructure. Except for achieving the minimum discharge into the North Sea and the minimum lake level, however, the strategy attaches no value to sending water north along either the IJssel River or the Amsterdam-Rijnkanaal, and hence sends the minimum amount consistent with the constraints.

The central part of SIMPLE MSDM is an implementation of the trade-off specified in Priority 3 of the priority list from Sec. 11.2. The flow in the Neder-Rijn is varied between a minimum of 25 m^3/s and a maximum of 18.7 percent of the Rijn flow, while the withdrawal at Tiel is varied between a minimum of 2 m^3/s and 120 m^3/s. (Both upper and lower bounds can be differently specified, if desired.) SIMPLE MSDM tests different combinations of these flows to determine which results in the lowest sum of costs to shipping due to low water on the Waal and the IJssel, cost of sedimentation due to withdrawals at Tiel, and salt damage to midwest agriculture due to the salt wedge. The test involves changing the flows on the links of the Distribution Model network representing the IJssel, the Neder-Rijn and Lek, the Waal, and the Lower Rivers part of the network, and then calculating the costs using the cost functions derived in Chap. 4 (the salt wedge) and Chap. 7 (shipping) of this volume.

Table 11.1 shows the results of a comparison of three managerial strategies tested in Distribution Model simulations. The three strategies are the RWS strategy, the SIMPLE MSDM strategy, and the VEL strategy. We described the RWS strategy in Chap. 9, and SIMPLE MSDM is described above. The VEL (short for Velsen) strategy is essentially the same as the RWS strategy, except that the VEL strategy gets more water for cooling power plants on the Noordzeekanaal (especially the Velsen plant) from the IJssel lakes, and less from the Waal or Neder-Rijn via the Amsterdam-Rijnkanaal.

The table compares the three strategies for both the present infrastructure and sprinkling scenario, and for the "MAXTACS" infrastructure and maximum sprinkling scenario. We obtained the numbers in the table from runs of the Distribution Model, and not from MSDM. Thus, costs were calculated for entire years, not for single decades. The benefits shown in the table are reductions in costs relative to case A, the first column of the table. Unlike MSDM, the Distribution Model calculates costs incurred by agriculture on unsprinkled plots and plots sprinkled from groundwater, as well as costs from sprinkled plots. This is the reason for the large agricultural benefits shown for the "MAXTACS" cases.

Comparing first the cases with the present infrastructure and the sprinkling scenario (first three columns), it appears that the RWS strategy is the best of the three on the average (results for the 1943 external supply scenario are essentially the same as results averaged over all years). This is due to an unintended difference between the strategies, namely that the RWS strategy sent an average of about 5 m^3/s more cooling water past Velsen than did the other

Table 11.1

COMPARISON OF MANAGERIAL STRATEGIES IN AN AVERAGE YEAR AND AN EXTREMELY DRY YEAR

Case Description(a)	A	B	J	H	F	K
Management Strategy	RWS	Simple MSDM	VEL	RWS	Simple MSDM	VEL
"MAXTACS"	no	no	no	yes	yes	yes
Sprinkling Scenario	1976	1976	1976	max	max	max
Shipping Scenario	1976	1976	1976	1985	1985	1985
Dutch Net Benefits (Dflm/yr), 1943 External Supply						
MAXTACS	--	--	--	-50.0	-50.0	-50.0
Agriculture	--	0.3	0.2	360.0	360.0	360.0
Shipping	--	0.2	0.4	--	0.3	0.3
Thermal(b)	--	-2.0	-1.8	0.4	-1.6	-1.4
Total	0.0	-1.5	-1.2	310.4	308.7	308.9
Dutch Net Benefits (Dflm/yr), DEX External Supply						
MAXTACS	--	--	--	-50.0	-50.0	-50.0
Agriculture	--	7.2	4.0	2154.0	2148.0	2150.0
Shipping	--	2.0	5.6	-10.4	-12.7	-12.7
Thermal(b)	--	-4.5	-3.7	-0.7	-4.3	-4.0
Total	0.0	4.7	5.9	2092.9	2081.0	2083.3
Minimum Summer Lake Level (CM Relative to NAP)						
1943 External Supply	-20	-21	-20	-24	-25	-24
DEX External Supply	-31	-35	-32	-50	-49	-49

(a) The letter designations given here are those given to all Distribution Model cases run for the impact assessment phase on the PAWN study. See Vols. I and XI.

(b) In these cases, we have imposed a three-degree standard on excess temperature. If the standard were seven degrees on canals and regional waterways, the thermal benefits would be essentially zero.

two strategies. Since we have imposed a three-degree thermal standard everywhere in these cases, the additional cost of generating electric power under the other strategies more than cancels the benefits to other sectors. If a three-degree standard had not been imposed at IJmuiden, or if the two strategies had been adjusted to cool Velsen with the same amount of water, the SIMPLE MSDM and VEL strategies would have had average annual net benefits of approximately 0.5 Dflm in comparison to the RWS strategy.

We represent a very dry year by a scenario called DEX (Dry, EXtremely). In each decade, this scenario has the same rainfall and evaporation as was observed in 1976, but its river flows are among the smallest observed during that decade in any year. Under the DEX scenario, both SIMPLE MSDM and VEL are preferable to RWS, and become more so if the amounts of water for cooling Velsen are equalized among the three strategies.

The MSDM and VEL strategies give essentially identical results in both the average and extremely dry years. VEL sent more water past Velsen than did SIMPLE MSDM, so the thermal penalty is smaller. SIMPLE MSDM appears to have favored agriculture, and VEL shipping.

We now turn to a comparison of the strategies in the cases with the "MAXTACS" infrastructure and maximum sprinkling scenario. In the average year, except for the differences in the cooling of Velsen, both SIMPLE MSDM and VEL are slightly preferable to RWS, and equal to each other. In the very dry year, both MSDM and VEL appear to be worse for agriculture and shipping than RWS. The reason for this has to do with the fact that under both VEL and MSDM, the lake levels dropped to their minima in some decades, and large cutbacks in sprinkling were necessary. It is dangerous to compare strategies when this happens, because we never gave much thought in PAWN to the question of what to do in such an event.

11.5. FURTHER RECOMMENDED WORK ON MANAGERIAL STRATEGIES

We recommend that investigations into managerial strategies be continued. First, we think that some of the constraints imposed in the present version of MSDM should be questioned. The most important of these are the constraints on the minimum IJssel lakes levels, and the level control requirements in boezems and ditches. Under most circumstances, these constraints have an insignificant effect upon strategy costs. But on rare occasions, they can be extremely costly.

Second, SIMPLE MSDM should be further developed. It should be extended to incorporate priorities other than Priority 3, and to consider managerial tactics for the Maas and the regions. A simplified version of MSDM could have two advantages over the present, more complex version. First, it would be more efficient, so that cases could be run at lower cost. Second, it could be made to find the optimal strategy fully automatically, rather than requiring judicious human intervention (see Chap. 8).

The most straightforward use of a simplified MSDM would be to replace the RWS strategy in the Distribution Model, and possibly eventually to be implemented in practice. The comparison in Sec. 11.3 suggests that with further tuning, significant benefits might result.[1] But a simplified MSDM could also be used to investigate further the value of water stored in the IJssel lakes, and to investigate for the first time the value of water stored in the weir ponds of the Maas (the insurance functions mentioned in Sec. 9.1).

NOTE:

1. The tuning we have in mind involves adjusting the flow of water in the Noordzeekanaal to reduce the thermal penalty. The resulting average annual net benefits would be well under one million Dfl (see Table 11.1), which is insignificant. In a very dry year, however, the benefit could be as high as 10 Dflm. This difference is, we believe, large enough to be real, and not a consequence of inaccuracies in MSDM.

RAND/N-1500/5-NETH