

Organised Learning from Small-scale Incidents

Organised Learning from Small-scale Incidents

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.F. Wakker,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dinsdag 26 september 2000 om 10.30 uur
door **Florus KOORNNEEF**
elektrotechnisch ingenieur
geboren te Rotterdam.

Dit proefschrift is goedgekeurd door de promotoren:
Prof.dr. A.R. Hale
Prof.dr.mr.dr. B.A.J.M. de Mol

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof.dr. A.R. Hale	Technische Universiteit Delft, promotor
Prof.dr.mr.dr. B.A.J.M. de Mol	Technische Universiteit Delft, promotor
Prof.dr.ir. W.A.H. Thissen	Technische Universiteit Delft
Prof.dr.ir. E. Backer	Technische Universiteit Delft
Prof.dr. J.G. Bovill	Rijksuniversiteit Leiden
Prof.dr. A.G. Tijhuis	Technische Universiteit Eindhoven
Prof.dr.dr. B. Wilpert	Technische Universität Berlin

Published and distributed by:
Delft University Press
P.O. Box 98
2600 MG Delft
The Netherlands
Telephone: +31 (0)15-2785121
Telefax: +31 (0)15-2781661
Email: DUP@Library.TUdelft.NL

ISBN 90-407-2092-4
Subject headings: Incident Reporting/Organisational Learning/Viable Systems

Cover design: J.A. Geijp

Copyright © 2000 by F. Koornneef
All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilised in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission from the publisher.

Printed in The Netherlands

Everything should be as simple as possible - but not simpler.

Albert Einstein

Contents

LIST OF ABBREVIATIONS	X
SUMMARY	XI
1 INTRODUCTION ON ORGANISED LEARNING FROM SMALL-SCALE INCIDENTS	1
1.1 SCOPE	2
1.1.1 Focus on Small-scale Incidents	2
1.1.2 Organisation's Need to Learn from Incidents	4
1.2 ACCIDENT ANALYSIS FRAMEWORK	5
1.2.1 Accident Models focussed on 'System Under Control'	6
1.2.1.1 Learning Lessons from daily Accidents using the Deviation Model.....	7
1.2.2 Accident Models focussed on Human Behaviour in the Control of Danger	8
1.2.2.1 Learning Lessons from daily Accidents using the Hale-Glendon Model.....	9
1.2.3 Accident Models focussed on Safety Management Systems	11
1.2.3.1 Learning Lessons from daily Accidents using the MORT Chart	12
1.3 ANALYSING MINOR ACCIDENTS AT WORK	13
1.4 GUIDE TO THE NEXT CHAPTERS	14
2. LEARNING FROM INCIDENT DATA	17
2.1 OVERVIEW	17
2.2 EXPLOITING INCIDENT DATA	18
2.2.1 DOHO: many Companies in one geographical Area of Activity.....	18
2.2.2 Company A: one Establishment with several Plants	20
2.2.3 Company B: two Sites with dissimilar Plants.....	21
2.2.4 Company B: corporate Use of a generic Knowledge-Model	22
2.2.5 Medical Centre A: Reported Significant Observations (RSO)	23
2.2.6 Critical Medical Devices: many Members of one Product globally applied - Heart Valve Prosthesis Case	24
2.3 TOWARDS ORGANISATIONAL LEARNING	26
2.3.1 ARA: nation-wide Use of a K-model by Companies in different Sectors	28
2.3.1.1 Discussion	30
2.3.1.2 The ARA Project Results in short	32
2.3.2 Medical Centre B: Dialysis Trial - probing the K-model in a clinical Setting	32
2.4 LESSONS LEARNED	33
3. MODELLING LEARNING EXPERIENCES	35
3.1 LESSONS FROM PROJECT EXPERIENCES	35
3.1.1 Data Models	36
3.1.2 Detection	42
3.1.3 Selection	43
3.1.4 Description	45
3.1.5 Classification	47
3.1.6 Computation	49
3.1.7 Interpretation	50
3.1.8 Monitoring	51
3.1.9 The NMMS Framework revised.....	51
3.2 PROSPECTS FROM THE LESSONS LEARNED	52

4. ORGANISATIONAL LEARNING.....	55
4.1 LEARNING PROCESS FRAMEWORK.....	55
4.2 ORGANISATIONAL LEARNING THEORY.....	56
4.2.1 Theories of Action.....	58
4.2.2 Single- and Double-Loop Learning.....	60
4.3 MODELS I AND II.....	63
4.3.1 Operational Surprise - Learning Modes.....	63
4.4 TRANSITION OF KEY-CONCEPTS FROM THEORY TO PRACTICE.....	66
4.4.1 Organisational Learning Agency.....	70
4.4.2 Communication Channels.....	70
4.5 REQUIREMENTS FOR ORGANISED LEARNING FROM OPERATIONAL SURPRISES.....	71
4.6 CONCLUSIONS.....	72
5. SYSTEMIC INCIDENT NOTIFICATION SYSTEM: SINS.....	73
5.1 SYNTHESIS OF A SYSTEM FOR LEARNING FROM OPERATIONAL SURPRISES.....	73
5.1.1 Data Collection - Lessons Learned.....	73
5.2 SPECIFICATION OF SINS - PRINCIPLES.....	74
5.2.1 Surprise Experience: Channel-1 Communication.....	75
5.2.2 Surprise Notification: Channel-2 Communication.....	75
5.2.3 Surprise Lessons: Channel-3 Communication.....	76
5.2.4 Notification Data.....	76
5.2.5 Organisational Memory – Case Base.....	78
5.2.6 Organisational Inquiry.....	79
5.2.7 Feedback on Notification.....	80
5.3 CONFIGURATION OF SINS.....	80
5.3.1 Notification and Reporting System: SINS-1.....	81
5.3.2 Case Review System - Learning Agency.....	82
5.3.3 Case Retrieval System: SINS-2.....	84
5.3.4 Organisational Causal Factor Analysis System: SINS-3.....	85
5.3.5 Stop rules for Notification and Analysis.....	86
5.4 SINS IMPLEMENTATION LEVEL - SCALING DOWN / SCALING UP.....	86
5.5 CONCLUSIONS.....	87
6. SINS IMPLEMENTATION: THE PIERCE PROJECT.....	89
6.1 MEDICAL CENTRE C: PIERCE - LEARNING FROM ONE ACCIDENT AND LESS.....	89
6.2 GENERAL DESCRIPTION OF THE CARDIAC SURGERY UNIT (MEDICAL CENTRE C).....	90
6.2.1 Complications and Operational Surprises.....	91
6.3 PIERCE PROJECT.....	93
6.4 CARDIAC SURGERY PROCESS MODELLING.....	94
6.4.1 Time-sequenced Process Phases.....	95
6.5 RSO STUDY.....	96
6.6 SAFETY CULTURE ASSESSMENT.....	98
6.6.1 Outcomes of the Baseline Assessment.....	99
6.6.1.1 Safety Awareness.....	99
6.6.1.2 Incidents.....	99
6.6.1.3 Accountability.....	100
6.6.1.4 Openness.....	101
6.6.1.5 Communication.....	101
6.6.1.6 Conclusions from the Baseline Assessment.....	102
6.6.2 Safety Culture Progress Measurement.....	102
6.7 SINS IMPLEMENTATION.....	103
6.7.1 SINS-1 - Low Threshold Notification.....	104
6.7.1.1 SINS-1 Functional Specification.....	105
6.7.1.2 Realisation.....	106
6.7.1.3 Assessment of, and Conclusions concerning SINS-1.....	108
6.7.2 Review Team.....	109
6.7.2.1 Case Review Process.....	110
6.7.2.2 Learning by Case Review.....	111
6.7.2.3 Review Team Continuity.....	113
6.7.2.4 Review Team reviewed: Discussion.....	114
6.8 DISCUSSION.....	114
6.9 CONCLUSIONS.....	117

7. THE VSM VIEW.....	119
7.1 INTRODUCTION	119
7.2 WHAT HAS VSM TO OFFER?	120
7.3 VSM DIAGNOSIS OF OPERATIONAL SURPRISES	122
7.3.1 Communication Channels in the VSM.....	124
7.3.2 Anti-oscillatory Behaviour.....	125
7.4 CONCEPT OF ORGANISATIONAL LEARNING AND VSM	127
7.5 SINS IN VSM	130
7.5.1 Discussion VSM – SINS.....	133
7.6 VSM DIAGNOSIS OF LESSONS LEARNED	134
7.6.1 Incident Data Processing in VSM Perspective.....	135
7.6.1.1 Example 1: scrutinising Medical Device Reporting (2.2.6).....	140
7.6.1.2 Example 2: diagnosing the ARA Knowledge Network	141
7.6.1.3 Example 3: review of the (modified) NMMS Model.....	142
7.6.2 Specification of a Surprise Monitoring System for Organisational Learning	144
7.6.2.1 Example 1 revisited: specifying Requirements for a viable medical Device Reporting System.....	144
7.7 DISCUSSION	146
7.8 CONCLUSION	151
8. CONCLUSIONS AND FUTURE WORK ISSUES.....	153
8.1 VSM DIAGNOSIS	157
8.2 EPILOGUE	159
9. REFERENCES.....	161
APPENDIX A – MORT CHART: MAIN STRUCTURE AND USE IN ACCIDENT ANALYSIS	169
APPENDIX B – ISA PROGRAM FOR INCIDENT DATA COLLECTION AND K-MODEL BASED ANALYSIS	175
APPENDIX C – RSO DATA MODEL	177
APPENDIX D – RAPID & DIRTY PROTOTYPING: K-MODEL + RSO MODEL.....	179
APPENDIX E – ARA I PROJECT RESULTS.....	181
APPENDIX F – EXAMPLE OF AN ACCIDENT ANALYSIS USING ISA	183
APPENDIX G – EXAMPLE OF AN ANALYSIS OF A CLINICAL ACCIDENT USING ISA	187
APPENDIX H – ON CASE-BASED REASONING	191
APPENDIX I – PATIENT THROUGHPUT PHASES INVENTORY	195
APPENDIX J – REPORTED SIGNIFICANT OBSERVATION (RSO) STUDIE CAPU	199
APPENDIX K – INTERVIEW PROTOCOL (SAFETY CULTURE).....	207
APPENDIX L – PROTOCOL REVIEW VAN ABNORMALE SITUATIES BIJ CAPU.....	211
APPENDIX M – CASE OF THE WEEK REVIEW	217
APPENDIX N – CAUSE CODES	221
APPENDIX O – GENERIC CASES OUT OF OPERATIONAL SURPRISES.....	223
APPENDIX P – CYBERNETIC LEARNING - VSM VIEW	225
APPENDIX Q – UNFOLDING THE SIF IN VSM	229
SAMENVATTING.....	241
ACKNOWLEDGEMENTS.....	245
CURRICULUM VITAE	247

Abbreviations

ARA	Accident Registration and Analysis
BRF	Basic Risk Factor
DA	Decision Analysis
DOE	Department Of Energy
DOHO	Diepgaand Onderzoek van HavenOngevallen
E&CF Charting	Events and Causal Factor Charting
FAFR	Fatal Accident Frequency Rate
FDA	Food and Drug Agency
FOBO	Fouten, Ongevallen en Bijna-Ongevallen
GFT	General Failure Type
IDA	Intelligent DAtabase
ISA	Intelligent Safety Assistant
LAS	Leiden Anaesthesia Simulator
LTA	Less Than Adequate
LTI	Lost-Time Incident
MORT	Management Oversight and Risk Tree
MAUDE	Manufacturers And User facility Device Experience
NMMS	Near-Miss Management System
OL	Organisational Learning
PA	Problem Analysis
PPA	Potential Problem Analysis
PIERCE	Patient Incident Evaluation and Registration for Cardiopulmonary Environments
QUABS	QUality Assurance BreakdownS
RIMAZ	RIsticoManagement Academische Ziekenhuizen
RSO	Reported Significant Observation
SIDT	Systemic Incident Description Technique
SiF	System-in-Focus
SINS	Systemic Incident Notification System
SMS	Safety Management System
SOLOS	System for Organisational Learning from Operational Surprises
TRIPOD	a 'three-footed' approach to hazard management aimed at underlying causes of accidents
VSM	Viable System Model

Summary

This thesis describes the search for and findings about methods and mechanisms that help organisations to learn cost-effectively from (small-scale) accidents and other undesired operational disturbances. As it shows, organised learning from such events is less simple than one might think. Key issues like who should learn and what can be learned are as pertinent as topics like how to organise effective learning and what data to collect.

Residual (safety) risks exist in any intentional activity no matter how well this activity has been designed and realised. In all cases, accidents are undesired and represent a deviation in an otherwise wanted, goal-oriented process. As long as such accidents fall within the limits of the residual risk, their occurrence may be accepted, because, due to utility considerations, resources needed to prevent such accidents have not been made available. Incident monitoring is one way to find out whether or not residual risks turn out to be larger than assumed. Accidents and other operational surprises are seen as unwanted operational conditions that have the potential for, or have actually led to harm or other losses. The operations change under influence of such conditions and, therefore, the processes may deviate from the intentional ones. The earlier such a *deviating process* can be identified, the more effectively the intentional operations can be restored. Operational conditions are set up and managed by an organisational *system*. This management system embeds *systemic* organisational *factors* that precondition the occurrence of operational surprises. These factors underlying the surprise are the ones that are worthwhile to identify in small-scale incident investigation, as such factors provide the best opportunities for solving critical process deviation problems.

In this thesis, I discuss the evolution of thinking about learning system design from an epidemiological model for handling accident data to one embedded in an explicit policy for risk control in the organisation. It is an evolution from accident data collection for post-event data crunching guided by generic harm scenarios in the data definition, to data collection based on a knowledge model concerning organisational causal factors so that organisational learning principles becomes pertinent. This thesis traces the process of learning through a series of projects in different organisations. These all had the objective of introducing systems to learn from accidents and incidents. In the course of time the way in which this was done and indeed the whole design of the systems and the philosophy behind them evolved and changed. The resulting SINS concept (SINS = Systemic Incident Notification System) was implemented recently in a cardiac surgery department in its most developed form. The incremental insights into the requirements for effective learning from such small-scale incidents were gained the hard way. Principles of learning by organisations, rather than by individuals within these organisations, were needed to advance towards effective incident monitoring.

The focus of the early projects was on extensive data collection from many incidents in the search for patterns in accident conditions and harm process scenarios. This approach evolved over time and over the projects towards single event review. The introduction of structured knowledge models about systemic factors into the incident data recording process led to possibilities for learning from single event analysis. In other words, cost-effective organisational learning (OL) from small-scale incidents became feasible. Several projects demonstrated the potential of this approach in practice. Only at the end of the last project was the primary question of how to organise learning answered in a more theoretical way. In searching for more fundamental models and theories to explain and anchor the findings, I linked them to Argyris' principles of organisational learning [Argyris 92, 96] and to Beer's Viable System Model (VSM) [Beer 79, 81, 85; Kingston 96]. This made it possible to diagnose the correct organisational position of an incident data feedback-control system.

The theories provided great insight into the success and failure of the projects and made a number of conclusions very clear. As descriptive incident messages are just simplified models of the real world, the context of an incident notification message must be known for adequate understanding and interpretation in order to work out an effective solution to a problem contained in the message. Comprehensive incident notification messages are not a priori more useful than a short one like "something unanticipated happened...!" if the context is already known. Key notions coming from the Organisational Learning principles include those of the "learning agency" that is assigned to learn for the organisation, the "organisational memory" to store and retrieve lessons learned, the Theory of Action concept as well as distinct learning loops.

The Organisational Learning principles were used to help to synthesise and review the concept of a Systemic Incident Notification System (SINS). The learning agency became an explicit part of the incident monitoring and feedback configuration, and formed a key in solving the principal problem of loss of context variety in incident notifications. The other advance was the positioning of an organisational memory that contains lessons from previous experiences. The interaction between the learning agency and the organisational memory is crucial for the realisation of a cost-effective system: the better the memory function is, the less resources are needed for the learning agency function. SINS was partially implemented in a surgical unit of a teaching hospital. The hospital project provided evidence-based lessons about implementation opportunities and pitfalls. It also pointed out flaws in the SINS-concept that needed to be scrutinised. Although not completely implemented, this PIERCE project provides a critical test for the SINS-concept.

The Viable System Model was found too late to influence the vast majority of the projects described, but offers a generic framework for understanding the successes and failures of those projects and generalising the results. It provides a solid tool to scrutinise the SINS concept as well as lessons taken from earlier projects. The positioning of the incident monitoring and feedback function becomes much clearer in a VSM perspective. As it is worked out, the VSM provides not only the means to diagnose problems in any existing SINS, but can also be used to specify new systems, including the type, routing and variety attenuation of data in messages. It provides a powerful meta-tool for future work.

Future work should focus on the design of an organisational memory for storing operational surprises (or incidents) and control measures applied and the interaction with surprise-reporting users. Such a memory might be built using Case-base Reasoning technology. The proposed VSM-based methods for specifying and diagnosing organisational learning systems need to be validated. With the help of VSM, the adequate specification, design and development of an organisational learning system is feasible and calls for a completely worked-through implementation.

Introduction on Organised Learning from Small-scale Incidents

This thesis is about learning from accidents, near misses and other unplanned as well as unwanted events that occur during activities that are organised by people to reach a worthwhile goal. Throughout the text, I refer to such unwanted situations as operational 'surprises'. Such surprises happen daily, most of them as near-miss situations that did not result in any harm or other loss, but many others as accidents and some as disastrous events.

Accidents and near misses do occur once in a while and mostly when they are least expected. The adverse consequences of accidents, such as injured persons, fatalities, material and immaterial losses, are always undesired. This is especially the case in systems in which people work to reach pre-set system goals. If an accident happens, these system goals are jeopardised and rehabilitation of the system up to the pre-accident level of functioning will incur costs additional to the initial damage. Accidents result from unplanned deviations in system operations. These deviations initiate an undesired process which, if not stopped, will lead to the accident [Hirschfeld 63, Kjellén 83a, Hale 87a, Hendrick 87]. A near miss reveals such a process without actual damage being done. By early detection and identification of the deviation process, corrective action can restore operations to the desired mode with minimal costs. Analysis of near misses is particularly relevant to a better understanding of human interaction with risks [Schaaf 91b]. People are often more willing to talk about their actions and errors in relation to a near miss or minor accident than to a major one; the feelings of guilt are less. In addition, near misses can give clues to the human recovery behaviour that stopped them from becoming serious accidents.

In safety-critical systems, analysis of serious incidents will normally be focussed on finding causal factors of system failures in order to improve the system and, thus, to prevent such deviations in the future. The effectiveness of the analysis depends very much on the objectives of the investigation and the methods used [Benner 85]. Accident analysis in systems that are not (officially) classified as 'safety critical' is too seldom aimed at learning lessons to improve the overall system performance. In such systems accident data often serve hardly more than simple statistical purposes. There is a strong belief that for a single company very many accidents will be needed to draw sensible, statistically founded conclusions regarding system weaknesses from the analysis of the collected data. This

seems to present too much of a demand on the organisation, since resources for in-depth analysis of minor incidents are limited. This thesis shows that such a negative view is not justified. If the learning system is well organised, any organisation should be able to draw worthwhile continuous lessons from its small-scale incidents, however few they are.

1.1 Scope

Many of the accident analysis techniques which have been developed for highly technological and safety-critical systems can also be used for the analysis of near misses and minor accidents. Usually and due to limited resources, such analysis techniques will not be applied if the incident-related risk is classified as minor. Also, powerful tools can be too time-consuming to apply to resolve minor problems. On the other hand, one single, serious accident may be foreshadowed by more than ten times as many minor accidents and by an even greater number of near misses having similar causes. Iceberg models have been published to illustrate this point both from the perspective of accident ratios and from the perspective of loss control, e.g., see Bird et al [Bird 76]. While these ratio studies can sometimes be criticised for being too simplistic (ratios vary widely between types of incident and many minor incidents have no potential for being major), they do indicate that minor accidents and near misses that have the potential for being serious should be analysed on the same principles as used in the analysis of serious accidents. This increases the opportunities for learning lessons from incidents considerably and, consequently, the control of the operation under review can be better ensured.

The need to learn from single accidents is probably obvious in safety-critical systems. Less evident is which of many kinds of lessons can usefully be learned from a single accident. The potential is considerable. In-depth investigation, often by independent committees, can spell out lessons with respect to failure of equipment or operational procedures, to human interventions as well as to organisational issues and lead to redesign of equipment, rewriting of procedures, retraining of operators, restructuring of the organisation or replacement of management.

1.1.1 Focus on Small-scale Incidents

The focus in this thesis will be on 'daily' small-scale incidents, i.e. accidents as well as near-miss situations, that have no or limited direct impact on the system in which the disturbed activity is embedded. Man-made disasters like the inferno of the Mont Blanc tunnel in 1999, the El-Al crash on the Bijlmermeer, the Herald of Free Enterprise, Seveso, Tschernobyl, or Lake Aral are not discussed in this thesis, although the approach presented also applies to this category of incidents. Many accidents, near-miss incidents and other abnormal occurrences do not fall under such major hazard legislation as the European Post-Seveso directives. Yet, learning lessons from such events can still be justified. The Abnormal Situation Management consortium has assessed that the sum of economic costs of 'small-scale' occurrences is at least one order of magnitude larger than those of major accidents in the petrochemical sector alone [Nimmo 95]. A task force of the US Food and Drug Agency (FDA) provides another argument when it observes that

"substantial numbers of injuries and deaths occur annually from medication or device errors. In general, these errors are believed by experts to result from systemic problems, rather than from a single individual's mistake. Such errors are not totally preventable, but can be minimised through interventions to the system." [FDA 99, p. 26]

In health care, such incidents are seldom investigated, not to mention that they are scrutinised for underlying factors that can be influenced systemically. Moreover, the investigation of large-scale accidents does not necessarily yield remedies for the causes of small-scale events. Perrow makes the point that major disasters are not simply the larger versions of 'small-loss' accidents [Perrow 84]. However, small-scale events do sometimes hold valid lessons for improved management of high-hazard risks.

Powerful investigation and analysis methods have been developed in the last decades for analysing serious accidents in technologically complex, safety-critical systems. The emphasis on post-accident corrective measures has shifted to assurance of safe functioning prior to the commissioning of the safety-critical system under review [Hendrick 87]. For example, it was unacceptable that on the first flight to the moon astronauts would be killed during the mission while the whole world was watching. Early pre-commissioning risk analysis techniques like Sneak Circuit Analysis [Boeing 70, Rankin 73] or HAZOP [Kletz 92] focus on the system hardware. The new standard IEC 61508 on functional safety of programmable electronic systems (PES) [IEC 99] contains a "safety lifecycle" that must be considered in detail when developing and commissioning a safety-critical or otherwise dependable computer-controlled system. With the commissioning of unique, first-time systems, it became apparent that the system's mission as well as the assurance of safe functioning requires a full systems approach. Operational readiness of a complex and safety-critical operation calls for a perfect match of equipment, human operators, operational procedures and managerial functions [Nertney 75, 78]. The scope of risk assessment and risk control, but also of accident analysis, broadened to the safety management system of the operation. Among the main results of this change in focus on assurance of safety in highly-critical systems was the development of the Management Oversight and Risk Tree (MORT) philosophy and its accompanying concepts and toolbox, including the MORT chart [Johnson 80]. The analytical tools help to bridge the gap between the harmful sequences of events leading to the accident and the organisational functions which are needed to ensure and maintain stable conditions for running safety-critical operations safely.

A main consideration for excluding from this thesis an in-depth consideration of disastrous events is that - in most cases - such surprises already raise a high public awareness. They will be investigated in-depth, virtually irrespective of resources required, by investigation boards like the NTSB (National Transportation Safety Board, USA) or by a team that is specially assigned. Methods like STEP and MORT [Hendrick 87, Johnson 80] have been developed for use in in-depth investigation. In-depth investigations aim at a detailed reconstruction of what happened in order to identify systemic causal factors that reside in or can be influenced by the management system of the organisation(s) involved. However, such methodologies and tools are in practice hardly applied or made available for the analysis of the small-scale surprises. Of course, there are exceptions. Case studies have been worked out for - at first sight - small-scale fatal surprises like the Therac-25

[Leveson 93] or a failure mode of a heart valve prosthesis¹ [Kallewaard 97, Cromheecke 00], but usually resources for looking into even these, and certainly the much smaller scale incidents with minor or no injury are negligible. I use the term 'small-scale' incidents for events that may occur daily and have a limited *direct* adverse impact on the organisational system(s) involved and/or on the society. Indirect losses, though, may become large eventually, and organisational systems may even go bankrupt or lose their share of a particular market. However, most of the 'small-scale' incidents will be events such as falling on a slippery floor, mental slips and lapses in more or less critical task execution, a rupture of a surgical glove or knocking over a cup of coffee onto some critical equipment. An operational surprise occurs when somebody or some monitoring system detects operational conditions that were not anticipated, e.g. a car driver spotting a bicyclist on a motorway, an operator logging some exceptional value in process parameters, or the spare lamp of an overhead projector that does not work after the breakdown of the main lamp. These operational surprises on a small scale can produce useful learning if captured and suitably filtered. In the last five years several tools [Schaaf 96, Groeneweg 98, DSI 00, Paradies 91] have been marketed that promise to help organisations to identify - on the basis of (small-scale) critical incident review - systemic causal factors, or 'root causes', or weaknesses in their Safety Management Systems in a cost-effective way. Such tools usually generate performance profiles concerning distinguishable SMS areas, like training, maintenance, and supervision. By combining such data from several accidents, such a profile becomes evidence-based, which is a strong feature. However, it requires multiple accidents and the 'validated' profile only tells the company where to look first in the SMS, without telling what to look for or whether the highest scoring category should be the one to resolve first. This thesis attempts to go much further by linking incidents and organisational learning more closely.

1.1.2 Organisation's Need to Learn from Incidents

The scope of this thesis is limited to accidents that occur during operations by an *organisation*, i.e. a man-made system that produces some product or service for a market. Such a system, which acts as host for processes, is designed and operated by people. The targets and constraints of the intentional processes are managed through the organisation's management system. The occurrence of an accident at least means that an intentional process has been disturbed and that the host management system did not prevent this deviated system state. From an engineer's point of view, the obvious thing to do is to acknowledge the unwanted system state and to use this process information ('unwanted system state') to recover from the disturbance in the first place, back into an allowed state. Prevention of a recurrence can also be assured by a well-designed feedback control system, so that the system is regulated within its allowed system limits. This line of reasoning leads to the conclusion that accidents and other unwanted process outcomes of organised, intentional operations are valuable operational data when deployed for feedback control. This seems obvious, so why is it not done? One reason is that people have not realised the importance of small incidents, despite the arguments of Heinrich since the 1930s to study them [Heinrich 31]. Another reason might be that people just did

¹ These are small in relation to major disasters killing many at once. Both cases involve single deaths, and were recurring because lessons were not learned soon enough.

not know how to use this kind of operational data to improve operational control. Maybe, this knowledge area was still too young in the late eighties when the DOHO project in the Rotterdam Harbour area² started. The notion that accident data could be used in process feedback control was not apparent to us in this early project. The Safety Science group has progressed a long way, since its early attempt in Rotterdam, along the road towards designing viable learning systems around incidents and accidents. I have been the primary learning agent³ in that process. This thesis describes this process of learning how to learn from these small incidents.

1.2 Accident Analysis Framework

Analysis of accidents and near misses is an essential feedback mechanism for learning from system failures and improving the performance of the current and future systems. This section reviews a number of accident analysis models and methods and their merits with respect to learning from incidents. They are described here in short, because these models provided the framework for the contents of the projects described in this thesis. There are other frameworks described and discussed elsewhere, e.g. based on Reason [Reason 90] like TRIPOD [Groeneweg 98] or PRISMA [Schaaf 96], which will not be presented here. They were not available when the work described here began. They also do not add any new conceptual insights not already present in the MORT approach described here. Their contribution should be sought more in refinement of classifications for error and other causal factors, which are not so relevant for this thesis. The MORT diagram is still the most comprehensive body of structured knowledge on systemic causal factors available. In this thesis the deviation concept of the accident process plays a ubiquitous role, and the deviation model discussed below provides all the strategic handles needed for the information system design. Explicit attention is paid in this section to the phenomenon of human behaviour because persons are systems themselves functioning within the host system, e.g. an organisation.

In the past, many theories or models about accidents and their causes have been developed and used: see for an overview [Hendrick 87] or [Johnson 80]. A number of them remain valuable to meet the objective of learning lessons to improve system performance. These fall into three categories, namely models which:

- a) start from system processes containing intrinsic hazards to be kept under control and incorporate notions of how technology or hardware moves from a safe state to a dangerous state;
- b) describe conditions of human error which may contribute to system failure;
- c) describe organisational conditions which are prerequisites for system failure in a safety-critical operation.

See [Benner 83, 85] for a more detailed overview and assessment of approaches to accident analysis. In the next sections only those which are most relevant to learning lessons from accidents will be briefly described.

² The DOHO project was the first accident analysis project of the Safety Science Group, see Chapter 2.2.1.

³ The concept of learning agency is discussed in detail in Chapter 4.

1.2.1 Accident Models focussed on 'System Under Control'

Accident models whose focus is the system to be kept under control presume that intended activities and operations entail hazards. If energy, or other harmful attribute, is present in an amount above the limit that the recipient element can absorb and if it is not controlled, a harm process may be initiated in which energy within the system will escape, or be transferred to exposed objects and converted into their harmful deformation or degradation.

The idea of energy transfer underlies several accident models [Haddon 66, McFarland 67]. Its simplest form is the 'energy trace and barrier' model as drawn in Figure 1-1.




Figure 1-1: *Energy Trace and Barrier Accident Model*

Any accident can be described as one or more sequences of energy transfers, influenced by more or less successful barriers. If there is sufficient energy, the prerequisites for an accident are present. If a vulnerable object is exposed to an energy flow without sufficient barriers then the accident will become a fact. There is a near-miss incident if an unwanted energy flow occurs without hitting a vulnerable target. The distinction between an accident and a near-miss incident defined this way is, therefore, marginal. Moreover, a barrier analysis will reveal what energy flows - as precursors of the accident - should have been interrupted, where and when [Trost 85]. This model also contains the safety-precedence sequence (SPS) as a strategic concept about energy management [Haddon 73, Johnson 80]. The SPS states that elimination of a hazard source (reduction or elimination of an energy source) is more effective than providing any vulnerable object with its own protective barrier, as long as this elimination or reduction does not threaten the system goals.

MacDonald's accident sequence model [MacDonald 72] and Kjellén's similar model [Kjellén 83b] start from the idea that it is possible to define a stable, normal operating state of a system in which processes run according to plan and hazards are kept under control. Harm or damage only occur if the (designed) control is lost and not recovered before the energy comes into contact with the vulnerable system element. Thus, an accident is a process of deviation from the desired, normal mode of operation. These models describe time-sequenced process phases, although not in terms of detailed event chains. They provide a basis for a classification system for the types of prevention measures that are possible, by identifying the stages in the accident process where intervention is possible.

Figure 1-2 combines the deviation model with the energy trace and barrier model and makes the link to prevention and system design explicit. The types of interventions can be grouped under the following headings:

1. Elimination (e.g. substitution of toxic chemicals by non-toxic)
2. Elimination (e.g. substitution of toxic chemicals by non-toxic)
3. Built-in control functions (e.g. automatic process parameter regulator)
4. Detection and recovery measures (e.g. preventive maintenance, alarm systems)
5. Escape and exposure limitation (e.g. emergency exits, automation of dangerous processes)
6. Secondary safety barriers (e.g. protective clothing)
7. Damage limitation (e.g. fire fighting, rescue and treatment measures)

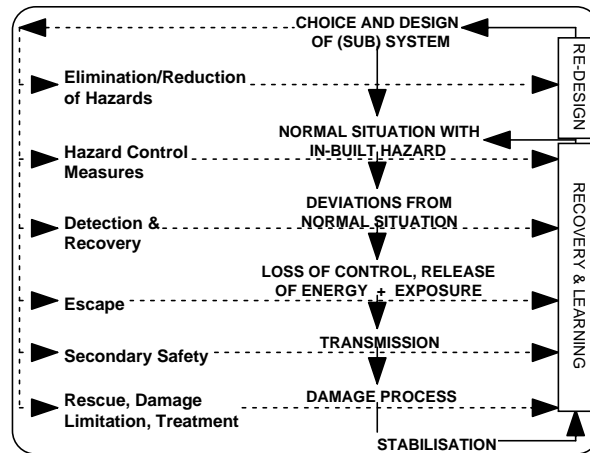


Figure 1-2: Accident Deviation Sequence Model after [Hale 87a]

The model also points out that redesign of the system is necessary if the desired, stable mode cannot be regained and maintained by detection and recovery measures and/or if the system goals can be better achieved with less costly deviations in another way.

1.2.1.1 Learning Lessons from daily Accidents using the Deviation Model

The deviation model invites the user to learn lessons from accidental processes with respect to the choice and effectiveness of risk control measures in relation to the state and phase of process deviation. For instance, in the phase of uncontrolled release of and/or expose to energy, see Figure 1-2, successful escape is imperative to prevent injury. This model indicates the need for system redesign if measures fail to detect the process deviation and restore a stable operational system state before harm or other damage is done. It also indicates that early detection, i.e. when the eventual accidental process deviation (from intended operations) is just beginning, pays off, because recovery measures will be relatively easy to execute. Learning from small-scale incidents and near misses aims at the early detection of emerging operational deviations to learn the lessons to improve the risk control measures within limited resources. By assessing the stage of process deviation once the deviation has been detected, the lessons to learn and resulting recommendations can improve the effectiveness of deviation control measures earlier in

the sequence. They can also point to the possibility of increasing the sensitivity for detecting the evolving process deviation in preceding stages of process deviation. Once matching risk control measures have been identified, their effectiveness and readiness need to be assured.

1.2.2 Accident Models focussed on Human Behaviour in the Control of Danger

More systematic assessment of human behaviour in the control of danger was made possible by the work of Rasmussen [Rasmussen 80], who developed a taxonomy of decision-making errors by control-room operators from which Reason developed his Generic Error Modelling System (GEMS) [Reason 87].

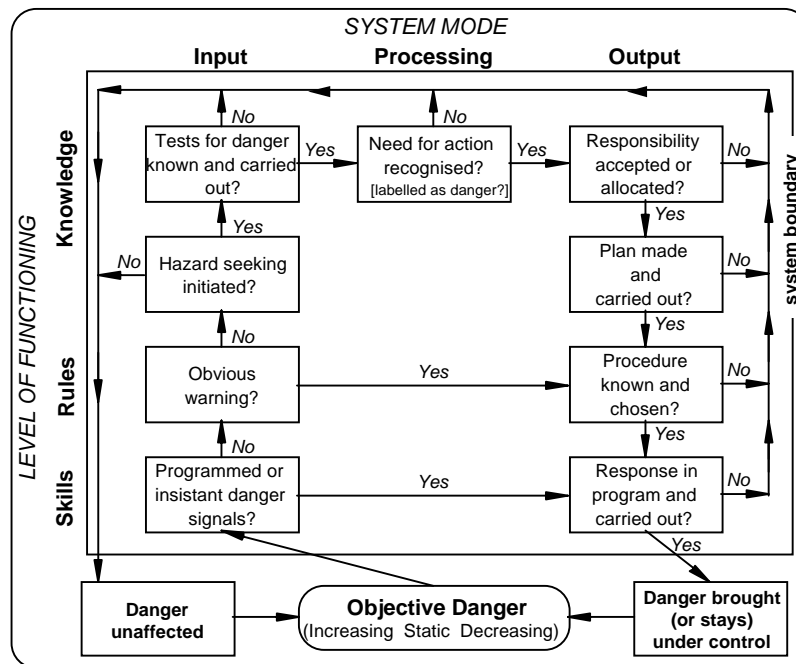


Figure 1-3: *Behaviour in the Face of Danger* (after Hale & Glendon [Hale 87a])

A combination of the GEMS model with other ideas of errors in information processing and problem solving [Surry 69, Hale 70] was developed by Hale and Glendon [Hale 87a], see Figure 1-3. The model starts from the presumption that some danger is always present. In order to work safely in a dangerous environment, persons at risk must recognise the danger and have the knowledge and skills, know the procedures and be motivated to cope with hazards which need action to control them. The task of the individual is to detect the (potential) hazards (input mode), assess the related risks, decide to take appropriate action to reduce the risk (processing mode), and execute this decision accordingly (output mode). For common, known hazards in routine activities, the level of individual functioning during hazard identification is typically routine and largely skill-based, as is the execution

of the learned risk-control actions (e.g. cornering in a car). In more complex situations, where several alternative actions are possible, the selection of a correct procedure to control the hazard is normally at a rule-based level (e.g. routine shut-down or start-up of a chemical plant). The recognition of new risks and learning about new responsibilities and methods to act will be at a knowledge-based level.

The Hale-Glendon model can be redrawn as an event tree [Thomas 95], see Figure 1-4, and thus it is possible to map an accident sequence of events into categories of behavioural errors. Solutions to the identified behavioural problems will differ according to the type of erroneous behaviour and to the phase of the accident process, but can, in most cases, be embedded in the system under review. The model enables the user to distinguish between:

- slips and lapses, e.g. where two routines were confused with each other or the person was distracted and lost his/her place in the sequence;
- misdiagnoses, where the wrong routine was chosen on the basis of confusing information (e.g. wrong expectations, misleading warnings);
- mistakes at the knowledge-based level where the mental model or understanding of the task in hand was a basic cause.

It also distinguishes between input problems (e.g. failure to notice or understand warning signals, lack of knowledge that something is hazardous), processing problems (e.g. unclear allocation of responsibility, lack of knowledge of preventive actions) and output problems (e.g. clumsiness in carrying out an action or failure to respond rapidly enough to an emergency).

1.2.2.1 Learning Lessons from daily Accidents using the Hale-Glendon Model

The Hale-Glendon model recognises the individual as data processing system. The model focuses on the decisions that the individual may, or needs to, make in order to control an imminent (safety) risk. Figure 1-3 shows that failure to respond adequately at the skill- or rule-based level of behaviour leads to the need for knowledge-based decision making. An incident occurring at these two lower levels therefore calls for a system that explicitly raises the attention for the process to the knowledge level in order to solve the problem lying behind the incident. Figure 1-4 reveals that knowledge-based activities are critical and relatively complex. In addition, appropriate action must be available, known and executed well to bring the danger under control: failure to do so raises objective questions about availability, knowledge and execution of the action and may indicate a danger perception problem. Thus, the Hale-Glendon model permits the structured analysis of human decision making in the case of emerging danger. The type of lessons concern the reallocation of danger control behaviour, e.g. moving a knowledge-based assessment sequence into a skill-based capacity by means of recurrent simulator training or by changing conditions that reduce the imminent danger risks.

	A	B	C	D	E
	SKILLS Identification	RULES Identification	KNOWLEDGE Identification/assessment	RULES Control	SKILLS Control
1	programmed or insistent danger signals <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled				response correctly executed <input type="checkbox"/> Y <input type="checkbox"/> N danger - unaffected / increased / decreased / controlled danger brought under control
2		obvious warning <input type="checkbox"/> Y <input type="checkbox"/> N		correct procedure known <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled	
3			responsibility for implementing procedure accepted <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled	correct procedure chosen <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled	procedure executed <input type="checkbox"/> Y <input type="checkbox"/> N danger - unaffected / increased / decreased / controlled danger brought under control
4			need to test for danger known / recognised <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled		
5			responsibility for testing accepted <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled		
6		correct test chosen and correctly executed <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled	need for action recognised <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled		
7			responsibility for action accepted / allocated <input type="checkbox"/> Y <input type="checkbox"/> N danger remains uncontrolled	correct plan / procedure chosen <input type="checkbox"/> Y <input type="checkbox"/> N danger - unaffected / increased / decreased / controlled	plan / procedure correctly executed <input type="checkbox"/> Y <input type="checkbox"/> N danger - unaffected / increased / decreased / controlled danger brought under control

Figure 1-4: modified Hale-Glendon Model [Thomas 95]

1.2.3 Accident Models focussed on Safety Management Systems

When analysing the nature and causes of accidents, the intuitive approach has been to concentrate on the physical harm process, see 1.2.1, and the actors directly involved, see 1.2.2. A third viewpoint, which has evolved with the increased emphasis on management responsibility and on self-regulation [Robens 72, DMSA 93], is to focus on the organisational settings and control systems which were meant to operate and maintain the intended system operations, but which allowed the occurrence of an accident or near-miss. Such models start from the idea that the owner of the safety-critical operation can assure full control of the process-related risks. Full control means that all risks have been identified and either have been brought under operational control or have been accepted as (residual) risks. In order to realise this level of risk control, safety functions and responsibilities must and can be defined which are linked to departments or specific people within the organisation. If such a safety function fails with respect to the safety-critical process, in terms of the deviation models in 1.2.1, the accident process has already begun and damage will result sooner or later.

The Management Oversight and Risk Tree (MORT) [Johnson 80, Benner 85, Hendrick 87] is one of the very few examples in this category of accident models. The TRIPOD model [Groeneweg 98], the extensions to the PRISMA classification [Vuuren 98] and the more recent I-Risk model of safety management processes [Bellamy 99] are other attempts to structure organisational factors in relation to accidents. The first two lack explicit models of the organisational process, confining themselves to lists of General Failure Types or organisational factors. The latter has been developed in parallel with this thesis and therefore was not available to inform the development of its principles. MORT contains a model of the safety management system needed to ensure safe operations and encompasses the energy trace and barrier model described in section 1.2.1. Figure 1-5 shows the top tiers of the MORT chart (see Appendix A for a further discussion of this chart and the philosophy behind the entire MORT approach).

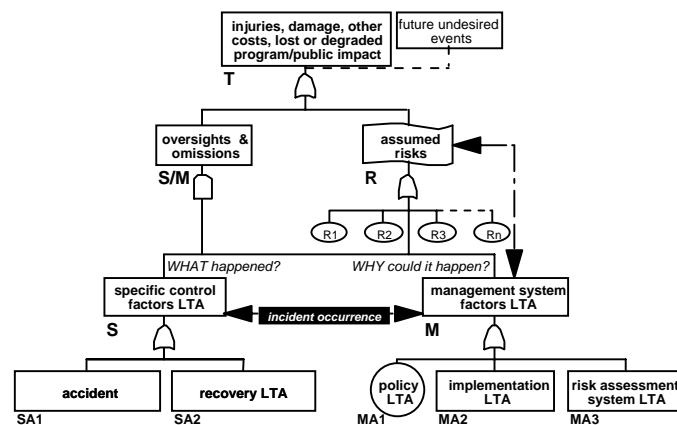


Figure 1-5: the MORT Accident Model

This tree-model has a top-event **T** and three main branches, consecutively, the **S**-, the **M**- and the **R**-branch. The model says that losses due to an accident (**SA1**) either result from

failure of the safety management system or fall within the limits of a predicted residual risk (**R**) that has been accepted by appropriate management. The cornerstones of the idealised safety management system are elaborated in the programme elements 'policy' (**MA1**), 'policy implementation' (**MA2**), and 'risk assessment system' (**MA3**). The specific operational control functions that failed to prevent the accident are worked out in the **S**-branch. One important problem in this accident model is to establish whether the risks, as they have manifested themselves in an accident, are really the known, assessed and accepted residual risks when the safety-critical operations were started. The initial assumption, though, is that accidental losses are due to failures in the safety management system, unless otherwise established by the evaluation of the accident-specific control factors. The model enables the identification of the relevant parts of the safety management system as well as specific control functions that need to be improved in order to realise the desired level of risk control. The expanded MORT tree can be used as an accident investigation technique.

1.2.3.1 Learning Lessons from daily Accidents using the MORT Chart

The structure of the MORT chart enables the user to derive recommendations to management for improved control of the safety-critical process from those tree events which have been evaluated as 'LTA' (Less Than Adequate) or which lack conclusive data for a sound assessment, i.e., which have been classified as 'more information needed'. The problem areas that are directly related to an incident will be found by performing an **S**-branch analysis. Once these have been found, recommendations to management can be derived related to safety management system functions. These can be generated by applying interpretation rules based on tree event evaluation and the structure of each of the relevant branches, keeping the core structure of the **S**-branch in mind, see Figure A-1 in Appendix A:

- **SB1** concerns control of hazard sources (post-design);
- **SC1** also concerns control of hazard sources, but now with respect to design and development, system implementation, maintenance and management of operational use;
- **SC2** considers exposure reduction to prevent harm or damage to vulnerable people or objects;
- **SB3** relates to control of, and evasive action by, a vulnerable person or object.

Recommendations may easily go beyond the span of control of the first line supervisor. For instance, task performance errors might call for improvements in personnel selection, training, or management concern for safety displayed by top executives (vigorous personal action). Residual (safety) risks exist in any intentional activity. In all cases accidents are undesired and represent a deviation in an otherwise wanted, goal-oriented process. As long as such accidents fall within the limits of the residual risk, their occurrence can be, at least temporarily, accepted. This may be for utility considerations, e.g. resources needed to prevent such accidents have not been made available. However, if the residual risks were not known explicitly or the occurrence of 'accepted' accidents exceeds the limits set by the residual risk acceptance criteria, then action is necessary. In

this thesis the MORT methodology forms a general background to the whole approach. It is also explicitly used as knowledge model in the later projects described here.

1.3 Analysing Minor Accidents at Work

The opportunities to learn about system factors that need improvement for increased safety are higher in frequency if small-scale accidents and near-misses at work are used alongside in-depth investigation of major incidents. In particular, humans in the system play a control role in the processes leading both to deviation and recovery. They make errors and miss opportunities for recovery or improvement because of the way in which the system in which they work is designed and organised. The enterprise has an interest in system factors that contribute to a decrease in the incidence of such human error. These errors can often be more easily studied in minor accidents or near misses, because those concerned are less inclined to cover up their error or to have forgotten because of the trauma of a major accident. The assumption is often made, as I make it here, that minor accidents and near misses if analysed will give useful information about the weak points in a safety system comparable to what comes out of the analysis of major accidents. The theory of Heinrich [Heinrich 80] and Bird [Bird 76] is based on the assumption that the underlying causes are identical. I would not ascribe to such a theory unreservedly, since it ignores the fact that people and organisations in general take more care to avoid major accidents than minor ones. This question of the degree of overlap is an open one, which should be resolved by empirical research. Care must, however, be taken that incidents and errors that could never have led to significant harm or losses do not receive too much attention, which would be a waste of resources. This issue will be addressed in chapters 5 through 7.

The old idea that an accident just occurs, has been replaced in the last decades by the acknowledgement of the process nature of accidents. The deviation concept (described in section 1.2.1) not only sets out from the idea of intentional operations, but also provides strategic options for control of risks related to the deviation phase. The safety of these intentional operations is a design and implementation issue for the management system that runs the planned processes. In this thesis, the Management Oversight and Risk Tree (MORT) is deployed to identify systemically the weaknesses in the management system functions that are needed to assure the intentional operations without incidental disruptions.

Although one might expect that the obvious activity, once an accident or other disturbing incident has occurred, is learning, this is hardly the case in practice with regard to small-scale events. Various explanations for this phenomenon may be given. The public or management awareness is too low, only one person is involved and injury may be minor, direct costs involved are relatively low, etc. So why bother? An awareness is growing that the annual societal costs of major disasters like the Erika, Three Mile Island, Piper Alpha, Bophal or (maybe even) Tschernobyl are at least a magnitude less than the aggregated costs of all minor accidents (up to several fatalities) and other operational disruptions [Nimmo 96, OSHA 98, Kohn 99, Thomas 99]. This should also redirect the focus to learning from smaller-scale incidents.

This thesis describes the search for and findings about methods and mechanisms that can help organisations to learn cost-effectively from small accidents and other undesired operational disturbances. As it shows, organised learning from such events is less obvious than one might think. Key issues like who should learn and what can be learned are as pertinent as topics like how to organise effective learning and what data to collect.

1.4 Guide to the next Chapters...

This thesis describes the challenges and findings of ten years of fieldwork that for better or for worse aimed at learning from daily accidents and incidents experienced at work. The projects are described in Chapter 2. They all contributed in some way to the set-up and realisation of the final project described in Chapter 6. The incremental insights into the requirements for effective learning from such small-scale incidents were gained the hard way. Principles of learning by organisations rather than by individuals within these organisations were needed to advance towards effective incident monitoring.

Chapter 3 provides a review of the experiences and findings described in Chapter 2 by means of the Near-Miss Management System framework [Schaaf 91]. Such a framework helps to position a particular approach deployed in a specific project and to criticise approaches within the framework provided. However, lessons drawn from the fieldwork in confrontation with the framework also led to a revision of the latter. Therefore, a modified NMMS model has been developed as a by-product.

In Chapter 4 Argyris' Principles for Organisational Learning are introduced and discussed, in order to pave the ways for converting the lessons from Chapters 2 and 3 into an improved concept for an incident monitoring and feedback system. Key notions coming from these OL principles include those of the "learning agency" that is assigned to learn for the organisation, the "organisational memory" to store and retrieve lessons learned, the Theory of Action concept as well as distinct learning loops.

In Chapter 5, the OL-principles are used to help to synthesise and review the concept of a Systemic Incident Notification System (SINS). The learning agency is now explicitly part of the incident monitoring and feedback configuration and forms a key in solving the principal problem of loss of context variety in incident notifications. The other advance is the positioning of an organisational memory that contains lessons from previous experiences. The interaction between the learning agency and the organisational memory is crucial for the realisation of a cost-effective system: the better the memory function is, the less resources are needed for the learning agency function.

Chapter 6 describes the partial implementation of the SINS-concept in a surgical unit of a teaching hospital. The SINS-concept provided a framework to allocate or assess basic functions needed to elicit signals about operational surprises and process these for feedback control. The hospital project provided evidence-based lessons about implementation opportunities and pitfalls. It also pointed out flaws in the SINS-concept that needed to be scrutinised. Although the implementation was not complete, this PIERCE⁴-project provides a critical test for the SINS-concept.

⁴ PIERCE = Patient Incident Evaluation and Registration for Cardiopulmonary Environments

Chapter 7 looks back at the process of learning about learning from small-scale incidents. The Viable System Model, discussed with respect to diagnosis of Safety Management Systems by Kingston [Kingston 96] is introduced. It provides a solid tool to scrutinise the SINS concept as well as lessons taken from earlier projects. The positioning of the incident monitoring and feedback function becomes much clearer in VSM perspective. As it is worked out, the VSM provides not only the means to diagnose any SINS, but can also be used to specify new systems, including the type, routing and variety attenuation of data in messages. It provides a powerful meta-tool for future work.

Finally, Chapter 8 briefly reviews the entire learning process of this thesis and indicates the areas for further research and validation.

Learning from Incident Data

2.1 Overview

This chapter describes a process of learning to devise a workable incident registration and analysis system for companies. It describes a series of projects on helping organisations to learn from small-scale accidents and other abnormal situations that paved the way for the conception of the SINS, the Systemic Incident Notification System that allows learning from one accident and less. The process of trial and error, which we went through, is illuminating to illustrate and analyse the essential features needed for a successful system. After a brief description of each of these projects, the problems encountered and the lessons learned are discussed in Chapter 3 through a set of themes proposed in the Near Miss Management System framework [Schaaf 91b].

The projects initially aimed at learning by means of statistical analysis⁵ of data on accidents and incidents in participating industrial organisations. Lessons for the improvement of safety management were sought within those organisations. A prime concern was the collection of data that contained information on organisational factors that relate to the incident occurrences. Standard notification systems already in use in the companies were usually too superficial. Therefore, schemes were developed to capture data on hypotheses about harm scenarios and accompanying causal factors. During the first series of projects, data collection focussed on data capturing strategies and statistical validation of the data models. One turning point emerged when it became apparent that the objective of learning statistically valid and significant lessons regarding local safety management system deficiencies cannot be realised through statistical analysis of occupational accident data that has been collected inside a single establishment. Within a reasonable timeframe, the number of cases is simply too small. Other approaches were probed. One line was to apply - existing - structured knowledge on safety management systems factors as a content-laden data model. The data collection would then become an assessment of the actual conditions in terms of the model variables. The values scored within the model and the underlying facts provide a sufficient base for learning lessons within the framework of the model used. Several projects were realised in which the MORT diagram was employed as knowledge model for data collection. Another approach

⁵ Statistical analysis that aims at generalised conclusions on correlations between case-describing factors, requires a sufficient number of cases to be captured. Assuming a normal distribution of cases described by n independent factors, the minimal number of cases to collect is about $5n$.

was explored that starts out from product-related accidents. This lead provided enlightening insights on risk control and risk management opportunities throughout the life-cycle stages of systems or products in critical applications. The relevance of time-sequenced operational conditions in relation to safety and control barriers was also elaborated in several projects in health care institutions.

Some of these projects were designed as a component of the Safety Science Group's research programme on accident and near-miss analysis, namely DOHO, Companies A and B, ARA, Medical Centre B, but other were realised merely by serendipity: Medical Centre A and the BS-cc heart valve projects.

2.2 Exploiting Incident Data

Eight projects provided the grounds for the SINS-concept that has crystallised out and been partly realised in the PIERCE project that is described in Chapter 6. According to application domain, the projects may be divided into 'industrial' and 'health care' sectors, with emphasis respectively on occupational and on iatrogenic accidents. The 'industrial' projects are DOHO, Company A, Company B (two projects) and ARA. The projects Medical Centre A-rso, BS-cc, and Medical Centre B-pilot are rooted in health care. From the research point of view, the distinction between sectors was practical rather than principal: in both sectors, the prime objective has been to learn with regard to safety management systems from accidental deviations in deliberate processes that eventually result in harm, damage or other forms of losses. A significant distinction exists between the earlier projects and the ones that deployed a knowledge model to allow one-off analysis. These latter projects are addressed in section 2.3.

2.2.1 DOHO: many Companies in one geographical Area of Activity

The DOHO project formed the beginning of the sequence in 1988. DOHO is an acronym for In-depth Analysis of Accidents in Ports (in Dutch: 'Diepgaand Onderzoek van HavenOngevallen') [Hale 87b]. The Rotterdam Harbour provides jobs to many people in domains like transshipment of goods, traffic control, ship maintenance and product quality improvement. A regional provider of occupational health and safety services for about 80% of all workers in the harbour participated in the project. The research objective was to develop and validate a data model on factors underlying occupational accidents that would allow learning lessons on safety management system factors of the companies involved. Statistical analysis of partly structured and descriptive data was supposed to lead to identification of crucial factors linked with specific classes of harm mechanisms. Data from many cases were required for model validation. Compared to prevailing accident reporting forms, data for many more variables had to be captured. The use and processing of paper-based forms was impractical. Thus, a convenient tool for data capturing was needed and built as software application for PC platforms: the IDA tool [Hale 91]. IDA stands for Intelligent DATabase and was initially designed as a front-end data entry interface for a powerful database system that runs on a main frame computer. Due to technological developments, the main frame database system became obsolete in favour of applications for new generations of personal computers before it was imple-

mented. The versatile IDA tool provided a meta-language environment in which database definitions and question trees can be built and modified easily, and that also offers access control for different classes of users.

In this project accident data had to be captured mainly by ambulance personnel during the transport of an injured victim to a health centre. Hundreds of cases were collected in a twelve-month period, but without the corresponding data on site-specific safety management system factors, which proved too difficult to retrieve. The primary purpose of the study was to uncover factors in the accidents that could be attacked to improve the level of safety in the system. The most suitable choice of model therefore seemed to be one based on tracing accidents back to factors that represent avoidable deviations from normal system functioning. The energy-barrier concept found in MORT [Johnson 73; 80] was combined with the concept of accidents as deviations from normal or desired circumstances [Kjellén 83b; Hale 87a]. Such a composite model encourages collection of data about the complete accident process from the initial deviation to the damage process and about the barriers and controls that failed or were absent. The latter are susceptible to far more effective and systematic intervention than the events immediately preceding the injury or damage, which had been traditionally recorded on accident forms.

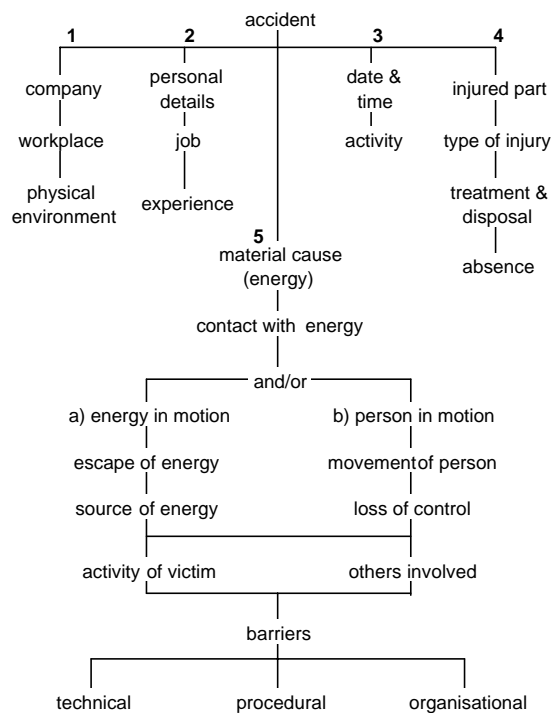


Figure 2-1: *initial question tree structure*

The structure of the data model employed in the DOHO project is shown in Figure 2-1. The question tree (Q-tree) starts with some general questions that identify the accident instance, such as date, time, and physical circumstances, as well as about the victim, his or her activity, injuries and treatment: data item categories 1 through 4 in the question tree.

Starting from 5, the questions then work backwards from the material cause of injury through the events and conditions which led up to it, and finally to the initial deviations from normal or desired procedures, activities and circumstances. The technical and procedural controls and barriers that failed are identified on the way. This backward questioning order is chosen for two reasons. It provides an unequivocal starting point for the questioning and avoids the considerable uncertainty and individual differences that arise if an investigator is asked 'to begin at the beginning'. Also, there is need to structure a tree that asks different questions about different types of accidents. The underlying assumption is that the logic of the choice of questions is much easier to construct based upon the nature of the proximal cause of the injury and the way in which the victim came into contact with the damaging energy. Therefore, the harm mechanism question must come early in the tree. For example, questions about the place from which someone fell and how they came to fall are relevant to falls from heights, while questions about how toxic materials came to escape are relevant to gassing accidents (but not vice versa). The final questions in the tree are aimed at clarifying the deviations from normal as potentially systemic inadequacies of control or safety barriers [Trost 95]. However, the project never managed to reach this step in a consistent way, due to the limitations in data collection from the companies concerned.

The DOHO project demonstrated that consistent capturing of data on occupational accidents cannot be realised from outside the companies involved. Ambulance personnel are not knowledgeable regarding operational conditions inside individual companies. The employers involved were also not lined up to provide necessary data for backtracking issues.

2.2.2 Company A: one Establishment with several Plants

In the subsequent project, the obstructions in data capturing found in the DOHO project were avoided. The management of an establishment of a leading chemical corporation with several plants on one site set itself the objective to reduce the number of accidents of any kind, i.e. occupational as well as plant accidents, further downwards to zero-level by improving capabilities to learn from incidents. One employee was assigned as "learning agent" on behalf of the organisation. Yearly, several occupational accidents and hundreds of operational disturbances were recorded. A hardcopy database on occupational accidents was available as trial data. The idea was to reuse the data model from the DOHO project after some modification of the section below 5 in the Q-tree, see Figure 2-1, this time guaranteeing that specific context information as requested through the Q-tree would be provided by the company. In order to put data captured in this way to the test, the modest number of cases collected annually would be increased through post-processing of hardcopy records of previous years as well as with plant process disturbance notifications. An additional reason for this combination is the argument that the safety management system for plant processes is the same in the case of operational disturbances due to operational control flaws in that for occupational accidents. The information flows initiated by a notification report needed to be defined and differed according to the class and specific aspects of a particular incident. At the time of the project, the prevailing system of distinct notification forms, one for each specific route of information flow, blocked the adequate response to multiple-aspects events, despite a notification contents overlap of about 90%. 'Intelligent forms' in which the routing through a set of questions is

driven by answers to previous questions could become very useful in steering information flows into all appropriate channels.

A limited number of cases were supposed to be selected for a MORT diagram analysis procedure, but this exercise never took place as case-specific data on control and safety barriers needed for MORT analysis were not produced. The existing hardcopy database proved deficient regarding case-specific data on control and safety barriers, so that only a few hundreds of cases could be processed using IDA. A new release of the IDA tool was needed in order to accommodate some additional functional requirements. This project in parallel caused much delay and also allowed room for unplanned modifications of the Q-tree. Eventually, the backtracking scheme including the testing of scenario applicability was basically lost. Consequently, at the end of the day case data were useless, but a new version of the IDA tool was produced and a preliminary notion on smart forms was conceived.

2.2.3 Company B: two Sites with dissimilar Plants

The next project was set-up in co-operation with company B also active in the chemical domain. The IDA tool was available from the beginning. Emphasis was placed on implementation of a fully revised question tree, although still based on the initial one and thus weighted towards occupational accidents. Company B held the view that local recording and review of accident data must primarily be valuable for local management to improve on-site management of safety risks. However, it was also looking for systemic management factors indicators that might indicate safety management system flaws at higher echelons within the corporation and for patterns which might be recognised more quickly by combining data from several sites. The prototype was tested at two locations with production processes in entirely different domains of chemical industries.

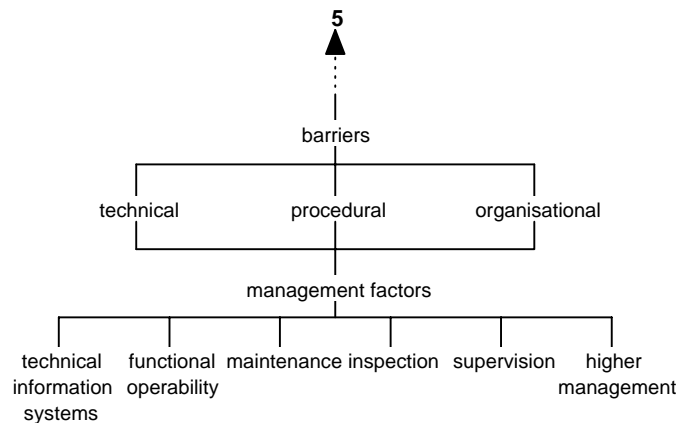


Figure 2-2: extension of initial question tree with management control factors taken from the first tier of the SD-branch in the MORT diagram.

The same question tree was applied after local adaptations regarding job description and naming of geographical areas. Local occupational safety and health practitioners took on the task of learning agents. Demographic data on victims formed a considerable part of

the initial series of questions, see parts 1 through 4 in Figure 2-1. After twelve months, for both locations less than one case per week had been captured: 42 and 41 respectively and far too few for meaningful statistical analysis of the data collected. Figure 2-2 shows the expansion in stage 5 of the backtracking section of the Q-tree with key management control factors from the MORT diagram. This extension proved to be a useful addition to the barrier section, given the distribution in answers provided.

The conclusion was drawn that the number of reported occupational accidents required for sound statistical analysis could not be captured within a reasonable timeframe of five years, see also note 5 on page 14. By the time that a sufficient number of cases would have been collected, the operational system would have been subjected to major modifications, so that conclusions from data analysis would no longer apply to the current local organisation. It was decided to drop the concept of statistical analysis of management factors, and explore the possibility of the employment of an appropriate knowledge model that could be used as reference for single case analysis. As the MORT diagram depicts an ideal and comprehensive set of safety management system functions that are structured in a tree, it was felt that the MORT diagram might serve as such a knowledge model.

2.2.4 Company B: corporate Use of a generic Knowledge-Model

Within the IDA tool, a new data model and navigational question tree were designed to enable the evaluation of individual incidents using the embedded K-model as a reference framework. The top section of the S-branch of the MORT diagram and four tiers downwards, including the top layer of the management control factors SD1...SD6, were taken as the core question tree (see Appendix A). Translation of related MORT questions and explanations in the form of context-sensitive help screen pages represent generic knowledge made available as support for users of the system, now called ISA (Intelligent Safety Assistant). The IDA programming shell allowed not only the rapid design of smart data capturing by means of electronic forms in which non-relevant issues are skipped, but also the programming of a rule-based interpreter of answers to be provided in the context of the reference framework. Thus, it was made possible to generate diagnostic reports regarding the individual incident submitted during data entry. As the ISA system now required that a local user submit evidence-based judgements on issues presented from the embedded knowledge model, special attention had to be paid to maintaining consistency in the process of answering core questions. The design of the ISA system allows multilingual editions in which can be switched between several languages, e.g. between Dutch and English, provided that language-dependent texts are available. Hypotheses concerning for example the role of maintenance in accident causation⁶, or on precursors emerging from operational stages preceding the actual accident phase can and have been built in and can be easily put to the test. The system was built and verified especially regarding the context-sensitive support during data entry. In addition an ISA User's Manual was produced and implementation schemes were described. The implementation of the new programme did not, however, go according to plan. The distribution of the ISA

⁶ The Health and Safety Executive published a series of reports on maintenance accident reports that account for the large majority of fatal accidents at work [HSE 86, HSE 87]

system throughout different branches of the corporation involved did not follow some of the essential steps from the implementation scheme and corporate project advocates moved on to other positions. Consequently, this project ended in silence.

2.2.5 Medical Centre A: Reported Significant Observations (RSO)

At first sight, accidents in health care institutions have an extra dimension in comparison with other domains of organised human activity: the patient is object of diagnosis, treatment and care, and usually much more vulnerable than their care takers. Florence Nightingale's 'doing-the-patient-no-harm' principle means that iatrogenic harm to a patient must be avoided. Health care workers may also become involved in an accident at work as real or potential victims: as area of concern, occupational health and safety also applies to those who work to improve the quality of life of other - sick - people. One undetected infected syringe in the hand of a nurse can trigger an iatrogenic accident and an occupational accident at the same time, if the nurse accidentally pricks her- or himself while administering the drug to the patient. Thus, care providers need to manage operational risks to patients as well as operational risks to themselves. Standard procedures for reporting and handling occupational accidents in health care institutions hardly exist, with a few exceptions, such as hepatitis infection incidents. The feasibility of monitoring operational anomalies and accidents was still only an emerging issue when this project started in 1989. Into this situation the accident reporting research was introduced, to study how it could help to turn hospitals into learning organisations.

New insights were gained from a project that aimed at analysis of operational anomalies of any kind in anaesthesia procedures. This project was embedded in a larger scheme that led to the development of the Leiden Anaesthesia Simulator (LAS) for training of anaesthetists [Chopra 94]. Research aims linked to that project formed the main drive for data collection: conditions embedded in the organisation of a surgical intervention have impact on real-time decision making and action by members of the anaesthetic team during a surgical procedure. A 'neutral' form of critical incident analysis, i.e. the method of Reporting Significant Observations [Eicher 76] was applied [Chopra 92]. For the observer, it may be easier to decide to report a significant observation than a 'critical incident', because in the latter case the observer first must make up his own mind about the criticality of the observed operational condition. Members of the anaesthetic team were asked to report any unexpected operational condition especially if it related to unsafe practices. During a period of 18 months, data were collected by means of RSO-forms tailored to anaesthesia practice. The data model was structured along the time-sequenced phases of an anaesthetic procedure, typically starting with pre-op assessment of the patient the day before the surgery until the patient was dismissed from post-op recovery. The generic phases 'before', 'during' and 'after' the intended process were adopted from the QUality Assurance BreakdownS (QUABS) method [Wells 77]. This is a task-level assessment of iatrogenic risks arising from operational flaws such as task performance errors, system failures and inadequate sequences in protocol steps. A two-month pilot project was used to develop the reporting forms. The versatile ISA tool was used for rapid prototyping of database and electronic data entry forms (see Appendix C for the RSO data model). During the pilot project, 100 cases were recorded and used to develop the classification schemes. Descriptive data on planned procedures and techniques, as well as

on place and procedural phase of the anomaly form the keys to the context of the reported situation. In addition, delayed recognition of the anomaly occurrence as well as a delay in counteractive response were also noted. Conditioning factors associated with the anomaly occurrence were recorded as observed attributes. Examples of such associated factors are failure to check specific critical system states, lack of knowledge for performing a task, long working hours, visual restrictions, time pressure, and poor labelling. A trusted anaesthetist in the accepted role as a 'learning agent' entered the data from the forms into the database after systemic verification of facts. Each report was evaluated individually, checked for accuracy and classified by the learning agent. The reports were classified into three mutually exclusive causal categories: 'human error' or 'procedural fault', 'equipment failure' and health condition-based 'complication'. They were classified further into four groups according to the potential risk of the incident for the patient: no risk = without any potential of adverse consequences, low risk = with potential of reversible harm, medium risk = with potential of irreversible harm, and high risk = potentially fatal. With the data on associated factors the first two causal categories could be preliminarily probed for systemic factors underlying errors or failures observed. After launching the standard form, another 549 significant observations were reported in the following eighteen months, of which 82% were considered preventable. The use of the RSO-forms was voluntary. The motivation of team members to file observations using the RSO-forms was maintained in part by the learning agent who provided positive feedback by regular presentation of preliminary data and selected 'interesting' cases [Chopra 94]. Statistical data analysis clearly indicated 'human error' or 'procedural fault' with 75% as the main causal category. Most frequently reported factors associated with procedural faults were failure to perform a proper check and lack of vigilance or inattention. There were no fatal incidents during the project, but eight cases resulted in lasting serious consequences and twenty-nine in minor consequences. The remaining 512 cases had no clinical consequences for the patients. Nevertheless, 34% of all reported observations were ranked as 'medium risk' and 27% as 'high risk' incidents. 45% of all reported observations originated during the maintenance phase of anaesthesia; of these, 52% incidents occurred during the first 15 minutes after induction of anaesthesia. The identification of causal factor *patterns* relating to specific classes of unwanted incidents to patients was however impossible in practice since so few cases actually 'produced' adverse consequences. Approximately 50% of reportable observations were actually recorded with ratio of about 1 RSO : 5 anaesthetic procedures during the peak period.

2.2.6 Critical Medical Devices: many Members of one Product globally applied - Heart Valve Prosthesis Case

In the previous projects, especially starting with the one described in section 2.2.3, the aim of accident and incident data collection shifted from problem recognition at the level of a population of workers towards improved management and control of operational risks⁷. This was the case even more with the next project. The step was made to expand the scope of reportable instances from operational anomalies that fall within the organisation's span

⁷ In this context, monitoring of incidents or 'operational mishaps' is an essential part in the MORT of 'upstream processes' needed to maintain functional operability [Nertney 87].

of process control towards abnormalities in the functioning of a device or technological system in a critical application, falling partly or wholly outside the organisation's control.

The prime key for reporting significant observations remains the detection of an abnormal situation that appears as an operational condition or process disturbance that needs to be corrected in order to prevent or limit adverse consequences. The breakdown or malfunctioning of a technological system may be such an anomalous occurrence. The purpose of detection is to restore normal operations by timely and effective corrective measures, such as replacement of the failing device. The purpose of notification is that lessons can be learned and disseminated. The detection and notification of device-related operational anomalies might gain added value if incident data from different users could be combined. This allows earlier identification of product-related risks, effective risk management options and possibly underlying design flaws.

A special learning need emerges when a product-failure mode cannot be handled adequately by users, whilst adverse consequences might be severe. Care providers such as physicians, nurses and paramedics usually apply medical therapeutic and diagnostic devices within an organisational setting. Operational problems in the application of such a system can more often than not be handled and corrected by the same staff, and reported afterwards. Learning then takes place wholly within the organisation. Even then root causes⁸ underlying technological flaws in design and manufacturing of a device may not be found if care providers - being device operators - do not combine their problematic experiences. This has been clearly demonstrated in the Therac-25 case study [Leveson 95], where users' meetings formed the key to identification and resolution of the software design flaw problem where the manufacturer's investigators and regulators could not reproduce the fault condition. However, some devices like implants leave the health care institution together with the real user, i.e. the patient. And medical devices are increasingly applied in home care where the patient must take care of the appropriate device control. In both cases, the handling of operational device-related risks may thus become largely, if not completely, beyond control of the care providers after delivery of the system to the real user.

The case of the Björk-Shiley convexo-concave mechanical heart valve prosthesis illustrates the issue of time-delayed system failure that requires active management of product-related risks before the product actually is brought into its operational phase. The system that failed was the cardiac centre that selected and implanted the valve without having the means to detect a beginning fracture, although this device failure mode was known. The system failure comes to light when strut fracture occurs and the disk escapes. At that moment there is no time to counteract, so the user passes away. Effective risk management options are not available for the end user and, therefore, need to be activated before the patient leaves the hospital after the insertion of the valve prosthesis. The project tapped into the topic of remote management of operational risks. The BS-cc mechanical heart valve [Kallewaard 97] killed patients by strut fracture resulting in acute cardiac system failure usually a couple of years after the implantation and outside the direct reach

⁸ The term 'root cause' is widely abused as a buzzword with an implied opportunistic meaning. In this text, the term is taken from the MORT terminology and defined as "a lack of adequate management control contributing to substandard practices or conditions (immediate causes) which result in accidents or incidents. Management control failures can result from the lack of adequate policy, inadequate policy implementation or from the lack of or insufficient risk evaluation" [Horman 84].

of the care providers. Opportunities for mitigation of direct adverse consequences were practically nil.

The mechanical heart valve studies [Koorneef 90, 93] focussed on the breakdown of assurance of functional quality of innovative prosthetic heart valves. The first study accompanied an epidemiological survey [Graaf 92] of a population of patients with the BS-cc valve which was initiated after repetitive reports of mortality due to outlet strut fracture. The Delft study covered the whole life cycle of the critical device, see Figure 2-3, with emphasis on options for passive and active control of risks related to the application of the valve by surgical implantation. The pro-active control as well as reactive management of actual risks of strut fracture turned out to be non-existent in the operational phase of this device. The time-to-correct after a single leg fracture of the outlet strut is at best two hours available for detection, transport to and admission by a cardiac centre. More often, the fracture remains undetected even after fatal cardiac arrest. The life cycle analysis model depicts the main stakeholders and decision-makers with respect to the decision-making about the specific valve through all the lifecycle phases of the valve. In aggregated format, the main stages are [Koorneef 96b]:

- Phase A: design and manufacturing
- Phase B: regulation and marketing
- Phase C: application and clinical risk control
- Phase D: operational use after discharge from hospital

The care providers in phase C have no adequate non-invasive means for detecting pre-fracture fatigue states of the outlet strut [Koorneef 90, Lepelaars 97]. The disk escape due to strut fracture is fatal unless immediately detected and treated by cardiac surgery. This is the kind of risk that actually jeopardises the prime function of the mechanical valve system and needs to be avoided systemically. In the valve's operational phase D the discharged patient as valve bearer is obviously at risk, but left helpless against this hardware failure after discharge from the hospital. Adequate risk control measures regarding this class of scenarios can only be taken in earlier life cycle stages of the product. Relevant control and safety barriers have been identified for each phase of the mechanical heart valve lifecycle more recently [Cromheecke 00]. The BS-cc studies demonstrated the relevance of manufacturer-independent systems for reporting device-related functionality problems and - when appropriate - propagation of 'early warnings' to all users, in addition to the significance of searching for effective risk control opportunities earlier in time, such as consideration of preventive replacement of the Björk-Shiley valve once the extent of the failures becomes known. A global reporting system is therefore necessary to reach the critical number of fractures as soon as possible in order to trigger such a decision.

2.3 Towards Organisational Learning

The projects discussed above were focussed on the collection and use of data on the phenomenon of unwanted operational surprises in order to learn from. The K-model approach, see 2.2.4, indicated the way to cost-effective one-off learning without the a priori need for statistical pattern recognition. The medical device failure case, see 2.2.6, demonstrated the change of problem ownership and effectiveness of residual risk control

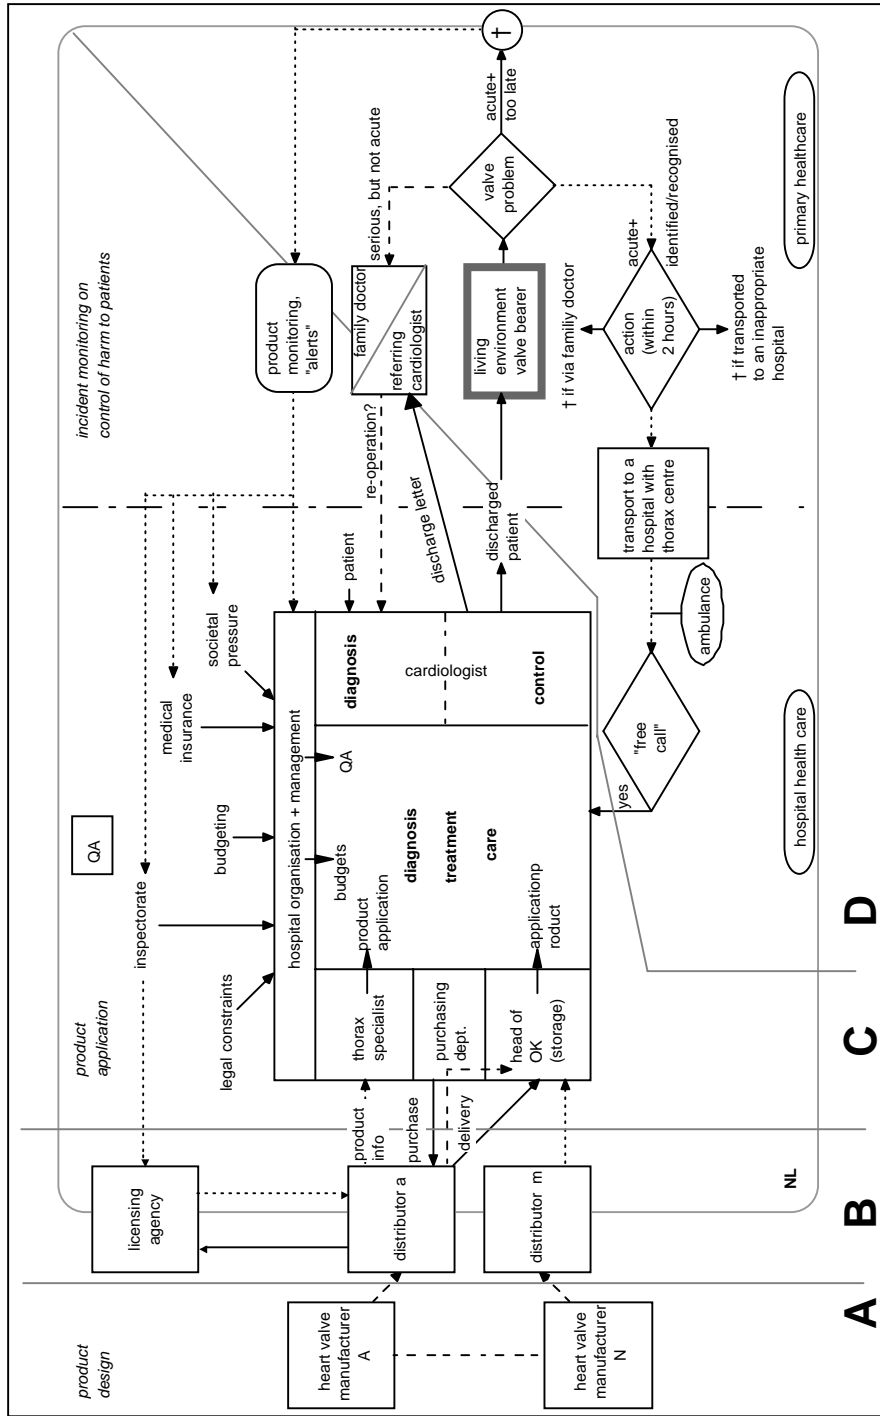


Figure 2-3: Risk Control in a Heart Valve Prosthesis' Life Cycle (ABCD-model)

in subsequent stages of the product life cycle, which raised the question of who needs what kind of incident data to learn what kind of lesson? Thus, the issue of learning by problem-owning organisations needed to be addressed. The projects discussed in this section represented a major step forwards in this respect.

2.3.1 ARA: nation-wide Use of a K-model by Companies in different Sectors

The ISA system described in section 2.2.4 was put to the test in the ARA-project on Accident Registration and Analysis. This Polish-Dutch technology transfer project was initiated from the Polish Ministry of Labour and Social Affairs and aimed at major reduction of work-related accidents in Poland by systemic improvement of safety management at national or macro level as well as at the site or micro level of individual companies. The project deliverables included a pilot-version of an accident data registration system facilitating realisation of accident prevention goals, technology-transfer regarding accident registration and analysis methods, the training of trainers for system users, courseware for training of future users, the education of a first group of users and training centre facilities. The delivery of an operational version of the envisaged registration system was scheduled for Phase II of the ARA-project.

The project started in November 1992 in the early phase of privatisation of state-owned enterprises. It was estimated that in those days the costs of work-related accidents amounted to at least 2% and possibly as high as 6% of the Gross National Product of Poland. Introduction of annual budgeting for state-owned enterprises and a no-claims bonus-system for social insurance employer premiums provided strong economic incentives for companies to invest in accident prevention programmes. Traditionally, in Poland every injury-accident was investigated by at least a local team consisting of a safety practitioner working on behalf of the company and a representative of a trade union on behalf of workers. The main purpose was to establish the rights of victims to receive financial compensation. Such investigations were performed intuitively rather than systematically. The ARA-project focussed on learning useful lessons from occupational accidents by these designated investigation teams. Inspectors from the Labour Inspectorate formed a secondary target group. They investigate fatal accidents at work and other major operational anomalies independently from the local enquiry team.

Over forty years of existing and comprehensive data acquisition at national level on work-related accidents did and still does not contribute to identification of systemic causal factors that can be brought under control by local management. Therefore, several companies had developed extensions to the accident registration systems in order to profit more directly from the administrative burden of the data capturing for the national surveillance systems. These add-ons concerned mainly manageable system factors such as repetitive equipment failures and related risk control measures. The first phase of the project involved a 'Jumbo Task Force' (JTF) with about thirty members from eleven large companies in different industrial sectors⁹, the national Labour Inspectorate, the trade union Solidarnosc, several research institutes and universities, as well as the Ministry of

⁹ Sectors include: mining (coal, copper) , tobacco industry, chemical, pharmaceutical, petrochemical, aircraft and auto industries

Labour and Social Affairs. With the combined set of deliverables in Phase I, including the training of trainers for system users [Karczewski 93], a start could be made with the set-up of an infrastructure for system users to exchange knowledge on lessons learned locally and experiences with the new ARA-registration system.

The first phase finished on schedule after one year of workshops and field trials with the ARA-registration system in the companies of the industrial JTF members. In this year, a series of multiple-days workshops was held to realise the project deliverables, beginning with the evaluation of several methodological frameworks and methods for accident analysis and data capturing. The Polish participants decided to adopt the MORT-philosophy and toolbox, because the MORT-approach is consistently focussed on systemic causal factors within the span of control of company management. The prevailing 'who-is-to-blame' culture has no place in the MORT-approach, a real eye-opener for Polish project members. However, the ARA-registration system had also to be used for compulsory registration of the comprehensive data for the external surveillance systems with data items defined in two forms, the so-called statistical form and the 'protocol'-form.

At the start of the ARA-project, a main hypothesis was that the ISA system developed earlier, see section 2.2.4, would provide a good means for exploitation of existing knowledge about systemic causal factors underlying any sort of injury accident or other loss-producing event within an organisation. It was planned that - no matter their technology - the individual companies could all use the same generic question tree addressing safety management issues directly related to the realisation of harm or losses. Although the answers and rule-based diagnostic reports produced by the ISA system would need to be seen within the operational contexts of the local company, it was proposed that data from one company could be compared with the data from other companies participating in the project as far as the K-model items were concerned. The MORT-based K-model in the ISA system was appreciated as a missing piece in the prevailing accident analysis framework. Data capturing based on the K-model allows single event analysis. The MORT tree analysis is, by its principle, contents and structure, a method for analysis of single loss-potential events (see section 1.2.3.1). Its use requires context judgements on tree branch states compared to a standard of adequacy defined by MORT. In addition, data collected this way are stored in a database and are later available for statistical analysis. A hypothesis was that by aggregating K-model data across companies, common problem areas might be identified, for instance weak job training programmes, so that regional, national or branch-specific solutions could be considered. During workshops, cases from project members were used for validation of the diagnostic module of the ISA-PL recording and analysis system (Appendix B contains a listing of the pilot program). Assessment of ISA-generated model-based topics required supporting facts provided by the informant and then, subsequently, the ISA-system produced a diagnostic report on management system weaknesses underlying the accident presented. In these trials, the ISA-generated diagnostic reports revealed more relevant issues than the usual investigation, according to the informants. An example is given in Appendix F. For the operational version of the ISA-software tool some functional requirements had to be added or improved. One was that the ISA system should become the only data entry interface for users at site level. In other words, external data requirements had to be accommodated by the ISA software. The Polish legal requirements are comprehensive and very detailed. For instance, for a relevant employee one job function code must be

selected out of several hundreds. The report generator had to comply with these regulations. For the sake of learning based on the K-model throughout participating organisations, the K-model data from participating companies had to be combined from all over the country. Especially regarding the coded K-model data, the multi-lingual features of the ISA system were very useful for assessment of the comprehension of the approach by the project members: the answers provided in Polish were directly understandable in the English version mode provided that the incident description itself was translated also.

The pilot system delivered included a mandatory accident notification form generator of which the Ministry of Labour and Social Affairs formally accepted the output at end of the project.

2.3.1.1 Discussion

The ARA-project contained almost all the ingredients for setting up an organisational learning agency at national level across the system boundaries of individual companies. During the project, the members of the Jumbo Task Force grew into the role of learning agents. The infrastructure of a nation-wide organisational learning system began to crystallise out, see Figure 2-4 [Koornneef 97b]. The main idea was that, in the ARA Knowledge Network at national level, users from participating companies share incident data within the network. These data are used to build a validated case base and identify accident scenarios in distinct classes. Users and experts co-operate to define and develop a new version of the ISA system in which the K-model as well as the diagnostic report-generating module are improved. The Change Control Board (CCB) governs the implementation of system improvements and the distribution to the participating companies so that all profit from the ISA system improvements coming from the shared resources.

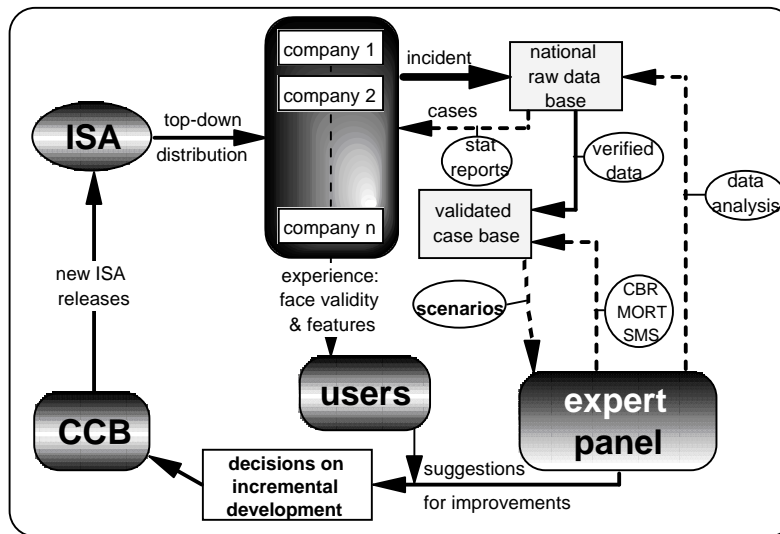


Figure 2-4: the ARA Knowledge Network at national level

At the end of Phase I of the ARA-project at the end of 1994, the pilot version of the ISA-system was available in a Polish version, ISA-PL, whereby both compulsory forms were integrated: the user had now one data entry system for all accident registration demands. Project Phase II was needed for the development and implementation of an operational and maintainable ISA-PL system, and for building the nation-wide knowledge network for ISA-users as depicted in Figure 2-4 and contemplated in the text box below.

How the ISA-PL Knowledge Network was supposed to function...

Lessons learned from aggregated and validated experiences of local, site-based learning agencies inside each of the participating companies would be disseminated through new releases of the ARA-registration system containing an updated K-model with rule-based diagnostic module. On the premise of validity of the K-model applied in the ISA-system and in order to profit more effectively from the fast learning capability of the ISA-system, ISA-users would be requested to submit their cases to a national database. These raw data would be used for periodic review on behalf of the users and for statistical analysis in order to feed a process of incremental improvement of the ISA-system [Koornneef 96a]. By combining data from different companies, a statically significant number of cases can be collected in a much shorter period of time than each company can manage alone. The review of cases and statistical data analysis would be performed by qualified experts from different domains, e.g. data analysis, accident analysis and safety management, and aim at improvement of the knowledge base of the system which would be returned to participating users as a new release of the system. Application of clustering techniques for analysis of the raw data would help to select typical incidents that could be fed to the validated case base after verification and review of the raw data. As short-term return for participating companies, statistical overviews based on the raw data might be produced and distributed. The validated case base was meant as a (public) source of incident data, which would be accessible for scientific research in a way similar to the Aviation Safety Reporting System (ASRS) database [ASRS 98]. The expert panel would use the validated case base for pinpointing accident scenarios and related sets of relevant Safety Management System (SMS) factors. These scenarios would give rise to review and improvement of the diagnostic rules in ISA as well as to advice to companies for improving their management systems. Suggestions for improvements of the knowledge base and the ISA system would be discussed with users in order to take practical experience into account, and to assess the practical feasibility of the proposals and the impact of a new release of the ISA system. A group formed by experts and users would need to approve propositions for a new round of incremental improvement. A change control board (CCB) which would issue new releases of the ISA-application would need to control implementation of approved proposals.

Unfortunately, Phase II was never launched because of political and interdepartmental frictions. A substitute Phase Ia was granted one year after ending Phase I in order to finish the ARA Project gracefully by setting up a network for ISA-users and for controlled release of ISA-software updates. However, the development of maintainable ISA-software was completely dependent on voluntary and in the end inadequate efforts. The time gap of over 12 months between Phase I and Phase Ia and the subsequent prolonged foreseeable delay in the delivery of the operational version of the ISA-PL system had a negative impact on the creation of the learning network for ISA-PL users. These efforts ended when the initial ISA-PL users concluded that it was not feasible at that time to form an ISA-PL Users Club as a legal entity at national level: the heart of the learning network was therefore missing.

The project members felt that the basic K-model and the derived rule-based diagnostic report were useful within the context of their own company. For daily incidents at site level, a 'basic' version of the ISA system would be sufficient. An advanced version of the ISA system would be needed, requiring dedicated extensions of the K-model, for investigation of severe accidents depending on the particular user category, e.g. labour

inspectorate, or on the nature of organisational activities, e.g. regarding plant-related issues in the chemical and petrochemical industries.

2.3.1.2 The ARA Project Results in short

The MORT-approach to SMS and accident analysis opened the eyes of all participants (industry, ministry, inspectorate, institutes and training centres) during the workshops, the trials with ISA-PL and training sessions: it formed a new way of looking systematically at accidents as opportunities for learning. The professional attitude of participating safety practitioners grew and matured through the project, especially with respect to on-site senior management (well-founded advice on SMS issues rather than the search for scapegoats). The validation of the K-model in the ISA system was limited to analysis of cases put forward by project members in workshop sessions, and by user experiences in field tests: the face-validity was thereby improved, but fundamental issues were left open. By review and statistical analysis of a number of incidents, the findings generated, the advice produced, and their relevance for site-based system users, it is likely that an expert panel can validate the diagnostic module derived from the K-model and the knowledge rules for consistency, relevance to end-users and clarity and improve them where appropriate. Co-operation between relevant groups of stakeholders - ministry, trade union Solidarnosc, labour inspectorate, safety practitioners and institutes - improved through the ARA-project. The ARA-project received quite some publicity nation-wide. As a result, by the end of Phase I some 50-80 companies had obtained an uncontrolled copy of the pilot version of ISA-PL and used it without proper training, though the nation-wide distribution was not due until the end of Phase II [Spijkervet 94] (see also Appendix E). Knowledge transfer courses are still being organised in co-operation with the Polish Labour Inspectorate with TU Delft. From Poland the word has spread among neighbouring countries, e.g. Slovakia.

2.3.2 Medical Centre B: Dialysis Trial - probing the K-model in a clinical Setting

The applicability of the generic K-model, as applied in the ARA-project, to review of patient-related incidents in a general hospital department was probed in a small pilot project in 1994. An ISA-application was rapidly prototyped by combining the ISA-PL system from 2.3.1 with the descriptive data model from the RSO application described in section 2.2.5 (see Appendix D). Two completely different incidents taken from a haemodialysis department were used to put the K-model to the test. In one case, an air bubble in the arterial line was a result of incorrect execution of a standard procedure. In the other, a complication due to the health condition of the patient resulted in a fatal shortage of oxygen in the brain. In this pilot, the head of the department answered the questions raised by the ISA system. As the K-model section answers are typically judgements on operational activities and conditions, for each judgement at least one supporting fact had to be provided. The questions, answers and supporting facts gave rise to detailed debates that were logged for use if a dedicated ISA-system were to be

developed in the main (so-called RiMAZ¹⁰-) project. The RSO-part of this prototype was suitable enough for the descriptive needs with respect to the abnormal dialysis situations, despite the fact that it was designed for capturing anaesthesia-related incidents. The translation, from the generic wording in the diagnostic report generated (see Appendix G), back into operational conditions, pinpointed relevant management system factors that had not been disclosed in the incident review that was a part of the standard quality assurance process within the unit. These factors included training, internal and interdepartmental communication and management of human resources as underlying operational failures. Lessons included review of specific protocols, supervision of team members who were still new with respect to dialysis operations, renewed analysis of operational hazards and real-time monitoring of anomalies.

This trial demonstrated that the ISA system could be used in the health care domain. The user had to be trained to answer the K-model questions consistently in order to arrive at evidence-based judgements. The evidence-based analysis guided by the ISA application relies on an objective method of raising questions and rule-based diagnosis whereby the rules are derived from the K-model. See for an example Appendix G. This ISA-system feature of objectivity helps informants to provide evidence for supporting subjective judgements made when answering the K-model questions. Domain knowledge of the operational conditions in which incidents occur is essential for the translation of the diagnosis generated by the system. For instance, a diagnosed lack of supervision may be due to a reduced number of operators working in the evening shift: only someone familiar with the workplace can make such an interpretation, which enriches the bare diagnostic category. The RSO-data recorded provide adequate opportunities for logging facts regarding the incident. The time spent during the trial for case analysis by means of the ISA system was intended to be reduced in the dedicated ISA system operated by a trained user. Validation of the dedicated ISA-dialysis system would be based on controlled field trials in haemodialysis departments of other hospitals participating in the main project. Unfortunately, because of the problem of funding and administrative difficulties the main project was never launched.

2.4 Lessons Learned

From the project experiences described in this chapter, the following conclusions have been drawn:

1. Learning about systemic causal factors just from **large data sets**, i.e. many cases, is a dead end street in accident analysis (DOHO, Company A, Company B, and Medical Centre A). This is due to constraints and limitations regarding the data model, the need for low detection thresholds, and a data collection time frame that is long with respect to business development cycles ('in 5 years' means 'in the long term').
2. Adoption of a **causal model** in the data model is feasible (Company B, ARA) if the causal model relates to the purpose of data collection and analysis.
3. Accidents, near-miss incidents and other **operational anomalies** occur **at operational system level** (all projects).

¹⁰ RIMAZ = Risk Management in University Medical Centres (in Dutch: Academische Ziekenhuizen)

4. The main **purpose** of data collection and analysis should be focused on local learning about improved control of operational risks in prime activities (Company B, ARA, Medical Centre B).
5. The **detection** of an operational anomaly relates to individual sensitivity to notice the **unanticipated surprise**-experience. This sensitivity is influenced by culture, e.g. professional will to learn (Medical Centre B, ARA), by supervisory attitude (Medical Centre A), acknowledgement of the need to manage an emerging operational crisis on hand (RSO), by an undeniable resulting situation (DOHO, ARA), e.g. injured victim, burnt-out building.
6. **People who detect** anomalies can be found in different observer positions: ambulance (DOHO: indirect witness), operators (Company A, Company B), learning agents (ARA, Medical Centre A).
7. The actual **reporting** and **recording** of detected operational anomalies relates to the **notification threshold**. The threshold is higher in cases where it is an administrative burden, e.g. due to a detailed data model (DOHO, Company B, ARA) or a prevailing blame culture. The threshold might be lowered by changing processing constraints, e.g. an active learning agency (Medical Centre A), by feedback of lessons learned, and by explaining the purpose, use and benefits of the recording and analysis system (Medical Centre A) to workers at the operational level.
8. The use of **smart forms** with an **embedded knowledge** body objectifies the elicitation of event and evidence data (Company B, ARA, Medical Centre B).
9. The application of an embedded knowledge model in incident recording and diagnosis implies during assessment a **translation** of model issues into anomaly occurrence and context describing facts, and visa versa after the diagnosis has been made (ARA, Medical Centre B). In other words, the filtering of event describing facts in order to map them onto knowledge model items to diagnose the event in terms of the model, calls for a reciprocal projection of diagnostic findings onto the operational settings in order to make them understandable.
10. The anomaly **description** is a message conveyer between reporter and learning agency about an operational surprise. The less the agency is familiar with the normal operational practices, the higher the demands on completeness of descriptive data concerning the normal process constraints (less: Medical Devices, DOHO, Company A, ARA; more: Company B, Medical Centre A, Medical Centre B).
11. The inference of causal factor diagnostics from the recorded data using a generic knowledge model results in generic recommendations that must be **interpreted within the context** of the reported deviating operations (ARA, Medical Centre B).
12. A smart software tool for recording and analysis of operational anomalies may be a helpful component in an organised learning system, but is no substitute for it (Company A, Company B, ARA, Medical Centre B).

Modelling Learning Experiences

3.1 Lessons from Project Experiences

The projects described in chapter 2 contributed to an understanding of the principles and pitfalls regarding the use of small-scale incident data for learning and can be reviewed following a number of themes. The data to be captured depend on the purpose of and the *models* used for data analysis. The *purpose* of data collection differs between projects and shifts from meaningful statistical identification of organisational causes to learning from single abnormal events experienced as acquisition of organisational knowledge. This brings up issues like *who* is going to learn, *how*, and for whom, or *what* is there to learn at which *level*. Learning agencies could be situated at national level (ARA), at regional level (DOHO), at site level (Companies A and B, ARA), at organisational unit level (Medical Centres A and B), as well as at global level (Products). Sometimes the lessons learned were about what was feasible or practicable in general in designing and implementing an incident-based learning system, in other cases about what approach was appropriate for what purpose or situation.

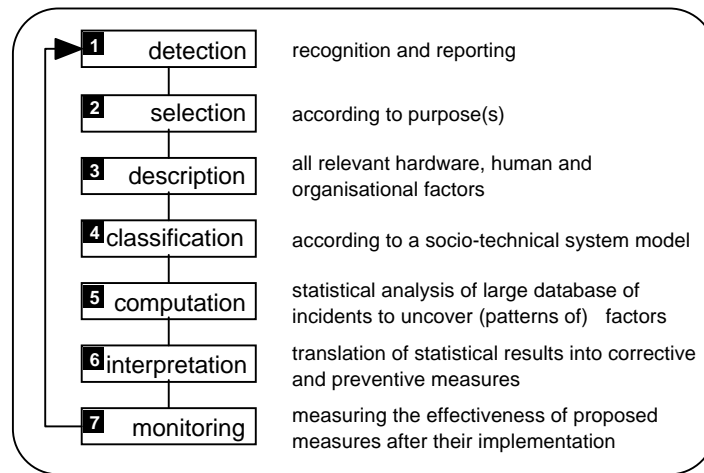


Figure 3-1: building blocks for different types of Near-Miss Management Systems (NMMS) [Schaaf 91b].

A structure for critical review of the lessons learned from the projects described in section 2.2 is provided by the 'basic Near-Miss Management System (NMMS) framework' [Schaaf 91b]. This framework proposes seven modules or building blocks that can be used to assemble different types of NMM Systems, see Figure 3-1. After recording of all near-miss situations detectable by employees in module 1, the interesting reports are to be filtered out in module 2 for further analysis in the subsequent modules. Any selected report must lead to a detailed description of the event sequence (module 3), so that facts and conditions found can be classified in module 4 according to a chosen causal factor model. In module 5, results are fed into a database for further statistical analysis identifying structural factors that must be interpreted in module 6 regarding control measures. Implementation of such measures is monitored through module 7 where effectiveness benchmarks are set, and through module 1: do events keep occurring?

The author mentions as the three main purposes for NMMS design qualitative 'modelling', quantitative 'monitoring' and 'motivation' maintenance. Modelling based on detailed case review allows the identification of 'new' causal factors. Monitoring aims at comparison with previous cases by statistical analysis of event scores on known causal factors or scenarios. The 'motivation' purpose is aimed at activation of the awareness of people about operational risks that have to be dealt with; this may be during training or as a periodic reminder: that danger is (still) present and active prevention is (still) necessary. In the latter case usually, an example case is presented and discussed at work floor level on a periodic basis. The heart of an NMMS is the *data model* applied. It directly defines the 'classification' - 'computation' and 'interpretation' modules that form the information processing section. This section can handle certain types of input data and, therefore, also defines the first three modules 'detection', 'selection' and 'description'. This chapter will therefore first discuss the lessons learned about data models. It will then summarise the more specific lessons per step of the NMMS model, and finally consider what has been learned about the NMMS as model.

3.1.1 Data Models

The data models in the series of projects developed over time. Learning from incidents requires empirical data for hypothesis forming and testing. An early hypothesis was that accident prevalence and characterising factors could be analysed by epidemiological methods in the same way as is done with diseases. The Q-tree data model employed in the DOHO project and in Company A was mainly conceived for epidemiological accident analysis (see the text box on 'Epidemiological and Causal Data Models...' below as well as sections 2.2.1 and 2.2.2). It aimed at proving the links between named management system factors, i.e. different classes of barriers, and the occurrence of accidents. Generic harm scenarios were grouped in two prime categories 'energy in motion' and 'person in motion', see Figure 2-1. Since the data elicitation processes failed for different reasons in the first two projects, the data analysis outcomes were useless in practice and the hypothesis concerning the value of epidemiological analysis could not be tested there. Hence, it was retained for later projects.

The data model was modified for Company B in the third project by insertion of a part of a causal Hazard-Barrier-Target (HBT) model that links accident occurrence with management system factors, including 'control barriers' or managerial 'controls', see

Figure 2-2 on page 21 [Trost 95]. Other factors that were not covered by the model could still be tested by statistical analysis, for example whether or not the accident (deviation) was causally linked to maintenance activity, and the geographical spot itself as direct environment of the deviation. Event evaluation based on worst credible - potential - losses rather than actual losses provided clues to handle near-miss occurrences: the depth of event review related to the credible adverse consequences. In the extreme case that maximum credible losses were negligible, the case under review was treated only as a minor warning possibly flagging emerging problems, but not worthy of action unless repeated frequently.

Epidemiological and Causal Data Models for Analysis of Accidents and Near-miss Incidents

Epidemiological accident surveillance might help to identify 'black spot' zones inside a larger system, e.g. the Rotterdam harbour, with 'outbreaks' of accidents due to repeated failures in particular parts of the prevention system. It might also detect scenarios or causal combinations that are new, i.e. for which no causal model has been defined yet. Without a causal model, hypotheses may be inferred by means of epidemiological accident studies. Such a study can reveal symptoms that are relevant in the occurrence of - specific classes of - accidents in particular in domains of human-organised activity. If one is starting from a total blank with no causal model, an epidemiological study must be based on a comprehensive set of possibly relevant descriptive factors that are all recorded for each reported instance. In epidemiological analysis, a matched control group or population database must also be or become available as a reference in order to allow meaningful analysis. Statistically significant differences in scores on specific factors between data from the accident group and data from the control group form proof that these factors are linked with the occurrence of - particular classes of - accidents.

A large number of cases is needed to identify factors with an incidence that differs significantly from the control group, that in all other aspects must be comparable with the target group. The availability of a case-control group or population data is a methodological 'must' for epidemiological data analysis. A strong feature of epidemiological analysis is its sensitivity for factors that differentiate between individuals, e.g. few sick from many healthy persons. This requires data collection regarding individuals. But with respect to accidents many years of research have led to few useful insights at this level, see for instance [Hale 87a]. It makes more sense to look for conditioning factors in the work place or at other organisational levels. Hence, for epidemiological accident analysis the accident-group, and hence also the control group must be chosen at the workplace or organisation level.

For specific categories of industry, reference data is available beforehand, e.g. in the mandatory risk management plans that EPA-regulated institutions must submit to public bodies [Rosenthal 97]. In other domains such population-wide reference data at company or department level may not be practical, if large numbers of cases must be captured first before epidemiological analysis can start at all.

In causal accident models, the link between event occurrence and causal factors is made explicit beforehand and, therefore, the data model includes scenarios of accident processes or pathways through the causal model.

Typically, data models for near-miss incidents and accident recording and analysis are hybrids containing for one part a causal model and for another part an epidemiological data section¹¹. In the epidemiological approach the part of the model that previous epidemiology has proved can get fixed in a causal model whilst the rest stays open to conjecture and proof. The other approach is to build causal models logically and completely, which is what we tried to do in DOHO and what MORT does.

A further substitution of the initial data model by a causal accident model was achieved by implementing the top section of the S-branch of the MORT chart, including its tree structure. The basic H-B-T accident model in MORT leads into a 'controls and barriers'

¹¹ Data for case identification or other administrative purposes are not discussed here, but may become a large part of the final application data model.

Each relevant specific 'energy flow' - 'exposed target' combination can be assessed in this model. The tree contents, structure and branch context can be represented by a set of rules used for diagnosis of management system factors that contributed to the incident under review, thus enabling single case analysis. If for a particular incident a specific tree item is assessed as 'less than adequate' or 'more information needed', then the model tells what data to look for regarding underlying causal factors. Data input is stored in the database for later use, including statistical analysis over a subset of all cases. This model was applied in the ARA project.

Note that the initial data models in the first few projects required many cases whereas the latest (from Company B - 2.2.4 - and ARA - 2.3.1 - onwards) can be used to analyse a single incident. The projects in companies A and B had shown that the collection of enough cases at that level to conduct epidemiological analysis was likely to take so long as to be impracticable. This drove the search for more causal modelling to benefit from single cases. The hope was still that projects such as ARA at a national level, and the abortive DOHO project at regional level, would provide enough cases at those levels for epidemiological treatment. However, this could not be proved, since neither ran to a successful data collection stage.

Another data model was developed for the RSO-study described in section 2.2.5. In that study, an observed deviation from expected operational conditions was to be reported. No detailed causal model was applied. The staff members were simply told to record what they found to be significant deviations. The data model just needed to be efficient in helping the description of a significant observation. Therefore, anaesthesia procedures were named and the work process was modelled in generic time-sequential phases: 'before' - 'during' - 'after' maintenance of anaesthesia. Any anomaly in the anaesthetic work process falls within one of these phases and can be associated with phase-specific tasks. In the project, these stages were identified in more detail in accordance with distinct phases in anaesthetic procedures.




Figure 3-3: *simplified version of Nertney's Wheel [Bullock 79]*

This was the beginning of the realisation of the importance of the process model to anchor data collection (see text box). Event descriptors were logged in coded form, reflecting the basic ingredients of work described in three categories [Johnson 73]: People - Plant - Procedures (PPP). Thus, the data capturing model used was synthesised from a basic QUABs¹² approach [Wells 77] and 'Nertney's Wheel' [Nertney 75, 87], see Figure 3-3

¹² In a QUality Assurance Breakdowns (QUABs) approach, a critical activity, e.g. a chemotherapy treatment is analysed in subsequent tasks (usually by different persons) in the tree main phases of 'preparation', 'execution',

depicting the basic ingredients of work (PPP) and the three interface domains. Ingredients and their interfaces need to match if safe operations (the bull's eye) are to be guaranteed. The work ingredients attributes were coded using tables that are widely used within the domain.

Causal and Process Models

The process model is used to anchor the instance of an operational anomaly to the intentional process in which the deviation occurred: place of the deviation and process stage are leading attributes. In the early projects (DOHO, Chemical Company A), a rough process model only contained workplace and company site descriptors. The RSO-study in Medical Centre A, see 2.2.5, and the heart valve prosthesis project, see 2.2.6, were more prominently based on process models because the objective was to identify opportunities to manage operational risks rather than how to control them. The main purpose was to look for causal factors earlier in the process that could be influenced so that the unanticipated surprise would be prevented in future. The generic QUABs process model and a generic product lifecycle model were suitable as a basis for data collection. Another rationale for the application of process model based data collection is that only the observed surprise needs to be described, as the intentional processes and normal operating conditions are known beforehand from the model. This reduces the burden of reporting.

Incident analysis based on a causal model implies assessment of the event occurrence within the scope of the causal model. Causal models can be used to identify factors underlying an operational surprise. In incident review, evidence-based event assessment requires information about the actual deviation from the intentional process. The causal model can focus the search for such data, which provides evidence on the deviation process, as was clearly the case in the ARA project and the trial in Medical Centre B. In the ARA project the causal model items indicated the type of data to capture in order to evaluate a model item. In Medical Centre B, the causal model was combined with the descriptive process model developed in the RSO-study in Medical Centre A. The analysis of an incident based on a causal model requires process data describing the operational anomaly, but can help to limit the search to, and focus collection on relevant data only.

In earlier projects, a prime assumption was that reports on all sorts of operational deviations from all stages in the operation process could be considered together in one database in order to identify systemic causal factors, thereby systemically ignoring the issue of linking to the specific operational process in which a particular incident developed. In later projects, one lesson learned was that this relationship with the specific process stage and conditions is crucial for organisational learning, especially when it comes to interpretation of generic - causal - factors and their translation into preventive measures that can fit back into meaningful 'theories of action' (see Chapter 4) that can be applied by the organisational unit(s) concerned.

Thus, in order to learn from individual incidents, the collection of data for review needs to be based on causal as well as on process models. The causal model of systemic factors and the intentional process model are in different (perpendicular) dimensions, and therefore call for different sets of STOP-rules. Essentially, STOP-rules for causal factor analysis are anchored to the applicable span-of-control of the problem owner and the incident severity, which relates to the required mode of (single- or double-loop) learning. STOP-rules for process model analysis are primarily anchored to the initiating and ending events, e.g. as deduced in a STEP analysis procedure [Hendrick 87], usually also within the span of control of the problem owner.

Possible causal factors were adopted from other studies [Gaba 94, De Graaf 98] and evaluated post-hoc for each reported observation. Basic anomaly scenarios and assessment of actual and worst credible risks were designed to allow pinpointing of problem areas in anaesthesia procedures. The origin of the prime causal factor underlying an operational disturbance in the anaesthetic procedure could be assessed as one of the

'aftercare' of that activity. The analysis focuses on assuring the critical steps were problems still can be identified in time.

mutually exclusive categories: 'equipment failure', 'procedural fault' or 'complication' related to the health condition of the patient. This causal model is much less detailed than the MORT model, see Figure 3-2. Although reference data were available from similar studies, epidemiological analysis was not performed [Chopra 92]. This model provided a user-friendly way to describe operational anomalies using context anchor points regarding process conditions and task execution, so that only deviations from normally expected operational situations needed to be described, without requiring an account of the intentional and ongoing process.

Initially, no data model was defined in the heart valve project for recording observed product-related operational deviations that occur outside an organisation. Such observations concern anomalies related to the use of the product or system in its operational phase, rather than unpleasant surprises in the work process of an organisational unit. The heart valve case study described in 2.2.6 set out from functional failure scenarios¹³ [Koornneef 93] that may occur or have been observed in the operational phase of the artificial system, i.e. after implantation. However, in a later stage in analysing the results, the need for a process model was felt. Adaptation of the generic QUABs process model to product-related operational anomalies led to a product's life cycle as the generic process model. If such a data model in a product anomaly surveillance system includes a risk assessment feature, critical systems can be distinguished from others more rapidly: a 'high risk' - 'low recovery opportunity' indicates the need for a risk control solution in an earlier phase of the product's lifecycle. A reported example is given in the text box about a pacemaker system-failure-by-design.

System Failure by Design

A pacemaker was found in maintenance read-out mode and hence not operational when supposed to be in active mode. It was designed to switch to read-out mode when exposed to a strong magnetic field. In the real world, the pacemaker-patient may become exposed to a magnetic field anywhere, e.g. when passing a security gate at an airport or nearby electrical devices, and thus step into a deadly trap without knowing even though the system was functioning conform its technical specifications.

Such an observed incident scenario (wrong mode) can be linked to the basic accident scenario of system failure to generate a cardiac trigger pulse when needed. The review of the reported observation then raises the issue of additional analysis of hazard control and risk management options throughout the preceding life cycle stages in order to identify effective, as well as feasible risk management solutions. One may think of product redesign for future users, equipping current users of the specific make with a mode checker and maybe a re-activator device, or even a patient recall for surgical device replacement.

A data model that combined the RSO-model described in 2.2.5 with the K-model described in section 2.3.1 united the strong features of both models. The test described in section 2.3.2 proved that the synthesised data model could be used successfully in the haemodialysis unit, although the RSO-part was tailored for an anaesthesia department and the K-model originated from industrial sectors.

We conclude here that there is a need for both a process and a causal model as basis for collection of incident data. Moreover, we need to move away from epidemiological causal models at least at the individual and workplace levels in order to allow learning from

¹³ In the case of mechanical heart valves system failure scenarios rooted in the bio-mechanical characteristics are mechanical failures e.g. fracture, disc escape or impingement, valve leakage, and clotting-induced thrombosis. The valve characteristics are mainly determined in early stages of the valve's life cycle.

small numbers of incidents. In the following sections the lessons for each step in the NMMS are discussed.

3.1.2 Detection

The NMMS framework states that "the detection module contains the registration mechanism, aiming at a complete, valid reporting of all near-miss situations detectable by employees". Detection is conceived as the double action of recognition of an incident and reporting of the occurrence. [Schaaf 91b, p. 29]

Usually, the recognition of reportable incidents is taken for granted, but it requires thorough reflection. Why would someone in a concrete and familiar situation decide that a specific situation is becoming abnormal? And why would that situation need to be reported, where, how and when? It may well be that the witness of an abnormal, reportable situation is not at all aware of this fact, for instance, because the abnormal situation is perceived as an ingredient of the normal mode of operation. Coping with unusual events may be seen as part of the game, and often one of the most interesting parts.

The detection of occupational injury cases as trigger for accident reporting seemed obvious at the inception of the DOHO project, but was not, because the result was that only victims that needed ambulance transportation were recorded: not lesser injuries which were home treated or where the victim's company did not call the ambulance. This gives a high threshold for reporting, and also depends on the co-operation of the ambulance drivers to complete the report. They had little interest in the learning that might flow from the reporting and the reliability of reporting was therefore not high. The first Company B project demonstrated better results within its planned reporting limits as every lost time incident was reported. However, minor injury cases and near misses were ignored. The industrial participants in the ARA project also traditionally detected only lost time accidents because it was only for those that the social security liability needed to be assessed.

A high level of vigilance regarding operational anomaly detection was realised for a short period in the RSO-study in the anaesthesia department, but required active maintenance by the learning agent. In most cases patients were directly involved, cases which naturally have high priority, but 7.5% of the reports were about unexpected situations of other kinds, like misplacement of stock medicine or an equipment problem during pre-use checking.

The deviation model, see Figure 1-2, shows that the definition of deviation will always be subjective and will depend on how large the participants think a deviation has to be before it is worth reporting. Arbitrary definitions can be agreed on for specific processes and technologies, e.g. air misses for aviation, or Licensee Event Reports in the nuclear industry, but it is hard to arrive at any generic boundaries. The best that can be done is to use a concept of 'unplanned' or 'unusual'. This is close to the definition used in the RSO-project of 'surprises', which experienced employees can define on the basis of their 'gestalt' of the normal process. However, this is subject to very personal, as well as cultural, interpretation and needs to be consciously influenced by training or discussion. The motivation to report or record an incident will be strongly influenced also by the

expectation that it has some value to do so, that something will happen as a result of it. Weegels [Weegels 96] showed that this notion of being modifiable was a strong factor in perceptions of incidents and accidents. Things that the participants in the research saw as inevitable were not reported as causes.

The NMMS requirement of 'complete, valid reporting of all near-miss situations detectable' is unclear. It can not be met as criterion, due to the lack of a base line. The completeness of data might be judged against the data model applied, but any data model is limited in its capability to catch all relevant circumstantial details of the observed reality. The intrinsic validity of a reported near-miss situation may be limited to that of the observer's perception of a surprise experience ("ouch! Something unexpected happened...") even if no operational anomaly occurred.

In the projects described, the drive for completeness in the reports about the incidents was limited to the assessment of items that were included in the data model as potential causal factors, and event descriptors that allowed efficient reporting. Free text descriptions of the course of events, observed circumstances, unintended consequences and suggestions for problem handling will always be incomplete, but may point to new surprising causal factors which can not yet be included in the causal model. Compared to the Q-tree data model, the process-phased data model in the RSO-study provides more anchors to normal operations, thus reducing the need to provide details about operational context. The processing of free text descriptions requires special care in order not to jump to conclusions and not to waste investigative resources. The K-tree data model allows systematic elicitation of relevant facts while assessing the model topics on the basis of evidence. However, to find facts and evidence describing the anomaly, it may be more efficient to employ dedicated investigative methods such as STEP [Hendrick 87] rather than a generic SMS evaluation system¹⁴, e.g. MORT.

The detection phase remains problematic and is now seen as being linked primarily to the motivation to learn. If the incident recording system is seen by all participants as driving learning, the detection phase can be bounded by their own definitions of what is surprising and unwanted, i.e. what they wish to learn to avoid in future.

3.1.3 Selection

"A NMMS that works well will probably generate a lot of "déjà vu" reactions on the part of the safety staff... To maximise the learning effect, some sort of selection procedure is necessary to filter out the interesting reports for further analysis in the subsequent modules." Selection should be done according to purpose. [Schaaf 91b, p. 29]

The purpose of event recording is the key for the applicable data model. Hence, selection of 'interesting' cases out of the whole set of recorded events may serve different goals. Throughout the projects, a variety of case selection goals has been identified: design of prevention plans at sector level, motivation and training, research on causality, resolution

¹⁴ Energy Trace & Barrier Analysis [Troost 95], MORT-based Events and Causal Factor Analysis [Buys 95], or Change Analysis [Bullock 81], i.e. Problem Analysis [Kepner 81], are effective methods for systemic fact finding.

of major safety management system flaws, identification of urgent operational risk management matters, or allocation of investigative and learning resources.

The initial data models were not designed for selection of individual cases. The aim was to collect many cases for epidemiological analysis. Any significant combination of causal factors found could be used later to select individual cases for some in-depth review to understand the causes better and to plan better prevention, provided that a case could be reconstructed with a sufficient level of detail. With the K-model, a useful selection can be made on any combination of relevant K-model item values irrespective of the number of cases available, but the number of cases required may become as low as one.

In the Company B, ARA and medical Centre B projects cases were selected for closer review and assessment for motivational as well as research purposes. Motivational goals were to make a point, e.g. that management system factors are found as causes underlying an accident, or to illustrate the principles of case review and to motivate people to report because of learning lessons. Research or learning goals were to identify scenario - causal factor combinations, or to understand why repeating patterns occur, i.e. why certain problems are recalcitrant and certain so-called preventive measures do not work in practice¹⁵.

A particular example is a public product anomaly recording system, wherein selection can be limited to 'high risk' potential for a specific class of products, for instance, medical implantable devices or toys for young children: see for instance the on-line MDRS database [CDRH 00]. Review of the reported problems and intake findings allows identification of potentially serious system failure scenarios for a class of devices.

The main purpose of initial evaluation of abnormal events lies in managing the allocation of scarce investigative and learning resources. The RSO-method [Eicher 76] included an initial criticality screening of newly reported observations in order to select cases for appropriate action, when needed, without further delay.

It is also relevant to know whether a new incident requires close attention or can be largely ignored. A simple "déjà vu" assessment can become very powerful when a new event can be recognised as one 'known' from the past. A "déjà vu" hit curtails the problem analysis process [Kepner 81], and allows us to postpone a decision to call in more investigative resources until a pre-set incidence criterion is met. For reporters, it is a great advantage if the recording of case descriptive data about a recurring and already 'known' incident type can be shortened drastically during notification compared with a new type of operational anomaly, which is not known to the learning system and for which the database system has no solution. The known incident types can be linked directly to known solutions.

And finally, once a given (new) type of operational anomaly has been recorded and assessed, similar situations will continue to occur unless successful systemic measures were implemented after the previous time to prevent recurrence. These repeating, known events can therefore be used to monitor the success of these measures. If they keep occurring something is wrong. Hence, the incidence rate of a given type of operational anomaly can be used as a selection criterion to trigger other action, such as double loop learning (see Chapter 4).

¹⁵ As operational anomalies occur during intentional activities, one may say that the *main* purpose of case review is to understand why designed preventive measures don't work in practice.

In summary, the selection phase was therefore implemented largely as a process of recognising known as opposed to new cases.

3.1.4 Description

The NMMS framework states that "any report selected for further processing must lead to a detailed, complete, neutral description of the sequences of events leading to the reported near-miss situation." The description should include "all relevant hardware, human and organisational factors". [Schaaf 91b, p. 29]

The description of a reported event is important for a proper understanding of its occurrence. However, the criticism in section 3.1.2 on the 'completeness' requirement regarding near-miss reports also applies here. Furthermore, methods for structuring available facts, identification of missing data and reconstruction of event sequences have their intrinsic limitations with respect to scope and type of analysis outcomes as well as to the resources needed for effective application of such a method. The causal model, see section 3.1.1, provides the key to structuring this step in the process. If the different causal links are in the model, these do not have to be discovered and established each time. The sequence that occurred this time merely has to be indicated.

During the course of the projects the question as to the purpose of the description of the reported abnormal situation became increasingly pertinent. As any set of describing words is just a model and, thus, a simplification of the reported reality, the criteria for assessment of the completeness of the account are far from obvious. If such a description is relevant in the analysis process, then the message conveyed by the description needs to be understood properly by the interpreter of the incident description. Therefore, the actual context of the reported situation must be accounted for one way or another.

In the initial epidemiological studies, free text descriptions were not relevant for analysis as only coded data items on harm mechanisms and possibly systemic causal factors would be processed for epidemiological analysis. Descriptions might just be convenient for a random reader, to provide substance to the report, but were otherwise discouraged.

In the K-model based projects and - in a different way - in RSO-projects, the description of the course of events became more relevant as a "mental movie" was required for the evaluation of the model items by the review team. Context data could more easily be captured when a basic core process throughput-model was applied, but compensation for context information that was lacking was still needed during case review. Additional coded attributes may be helpful here, but not sufficient for context description of task, operational and organisational conditions and state changes.

In all of the projects described, the ultimate aim was to learn lessons at the level of the organisational management system about systemic improvement of the control of operational risks. In the course of the projects, the type of descriptive data shifted, from data collected in conformance with the causal model applied, to description of a process deviation with the purpose to extend the causal model in new cases. Known cases were already included as causal branches, and the description only had to indicate which causal branch was followed, e.g. the heavy line in Figure 3-4. In the causal model tree each path represents a scenario of incident enabling conditions and harm mechanisms in an unplanned accidental process. New branches, i.e. new sets of causal factors may be

identified if data on new incidents is collected in a form independent of the known causal model.




Figure 3-4: causal model tree for incident scenarios

The selection and development of example cases drawn from reported incidents started with the Company B projects. The actual need for a 'complete' description was mainly felt in the need for creating a mental movie of the reported incident occurrence that was shared by the members of the learning agency as a common reference framework. The reasons underlying judgements about causal factors were then tested by them against the course and conditions of the event sequence, and this led to an evidence-based assessment. If this shared movie was inadequate, additional search for facts was triggered to allow a better reconstruction. The more the learning agent or agency was familiar with the actual process that was going on before the anomaly was detected, the less the need was felt for a 'complete' description. In the medical centre projects, the learning agency consisted largely of members from the organisational units that ran the processes in which the operational anomalies occurred. Consequently, case descriptions remained and could remain concise. The procedural method of Events and Causal Factor charting [Buys 95] proved to provide a powerful and objective means to organise facts and hypotheses for descriptive reconstruction of any reported abnormal situation for the purpose of learning lessons about management control factors.

Where case review was done by a learning agency that was less acquainted with the prevailing operations and operational conditions, the need for more descriptive data was felt immediately. The purpose of creating a mental movie that is shared by the members of the learning agency holds throughout the case review process. The more diverse the range of processes and industries concerned, the more this is an issue, as the national ARA project showed. Examples of accident descriptions worked out in the ARA project were presented in workshops by members of the local investigation teams who were very familiar with the daily operational conditions, or by an external investigator who had indirectly acquired such descriptive facts through an investigation report. The descriptions in the ARA project were generated by questioning the informers through evaluation of the K-model items and the search for evidence to support each judgement made. In other words, a dialogue was necessary between the reporters and the learning agency to fill in gaps in the description.

The project on product incidents required a similar treatment. The system failure identified needed to be detailed in a scenario to describe the critical system state and how this state was or could have been produced with respect to failure mechanism and operational constraints (see for example the text box 'System failure by design' on page 41). Identification of underlying organisational factors relevant to the management of product-related operational risks in earlier phases of the product's lifecycle is primarily important if a 'high risk' - 'no-timely-recovery-opportunity' failure scenario has been

identified. To discover these opportunities for earlier correction, the dialogue described above needs to occur. In other cases, a 'wait-and-see' risk management strategy may well be sufficient, and no descriptive data are needed beyond the recording of the failure occurrence.

In the case of critical device failure the micro-environmental context of operational product use is relevant to establish the reported operational failure mode and mechanism, and to assess risk control options if needed. In the case described, the pacemaker was switched to read-out mode unnoticed while the pacemaker bearer was searching for metal objects using a metal detector. In this example, the micro-environment of the product is defined by the activity 'metal search' and the tool applied, i.e. a 'metal detector' with a strong electromagnet¹⁶. It seems necessary that device-related functional anomalies are recorded in a way that the functional problem scenario is described in relation to the ongoing activity, and that essential device state parameters, e.g. the intended mode and expected behaviour, are logged. Actual and credible worst case adverse consequences are other important attributes for early detection of unanticipated product-related system failures and risk evaluation. Thus, a shift from feedback to feedforward analysis emerges.

Regarding this step of the NMMS, the evolution in the projects was to reduce the burden of incident description by transferring it either to the causal model or to a post-hoc interrogation exclusively of new cases, or repeated cases exceeding the trigger threshold for double loop learning. In these cases the learning agent extracts the description through a dialogue with those concerned.

3.1.5 Classification

"Every element in such a - descriptive - tree may be classified according to the chosen socio-technical or human behaviour model, or at least every 'root cause'¹⁷, i.e. the end point of the tree, must be. In this way the fact that any incident usually has multiple causes, is fully recognised." [Schaaf 91b, p. 30]

The purpose of classification in this thesis is to produce and group the set of causal factors that contributed to the reported operational anomaly and, therefore, to abandon the concept of a single or 'main cause'. The socio-technical models applied in the projects described in 2.2 are captured in the K-model depicted in Figure 3-2 and the P-P-P or basic ingredients of work concept shown in Figure 3-3. Some of the key elements in the K-tree were included in the Q-trees of the first projects. The RSO-type of event reporting is mainly based on the P-P-P concept. By definition, the K-tree model states that at least two causal categories are relevant in the case of a near-miss occurrence and at least three in the case of an accident. The P-P-P model states that the basic components of work and their interaction need to match if an intended operation is to run flawlessly: an operational

¹⁶ In contrast, the micro-environment in the case of the strut fracture of the BS-cc heart valve prosthesis is identical with the normal operational environment, i.e. large pressure changes and frequent blood flow conversions... fatiguing the strut welding points.

¹⁷ The term 'root cause' is poorly defined here as: 'the end point of the tree'. In industry, a 'root cause' usually is conceived as a causal factor which if absent would have prevented an accident. In the main text, the term 'root cause' is in accordance with the MORT definition, see also footnote 8 on page 25 above.

upset indicates a flaw in the actual P-P-P configuration due to any combination of mismatch opportunities among these components.

The use of a process model underlying data recording allows classification about operational contexts *wherein* (place and prevailing process stage) events occur. Critical operational stages or places can be identified relatively quickly. Furthermore, the place/process phase event-description attributes allow the reporter to limit notification data to describe the operational surprise without the need to provide a detailed description of the planned and ongoing operations.

The use of a causal model provides an explanation of *why* the events reported could occur, i.e. what kind of systemic causal factors allowed the abnormal situation to occur. Within the scope of the causal model, any single operational surprise can be evaluated for its underlying causal factors. The generic findings then still need to be translated in order to make operational sense and in order to assess the utility or cost-effectiveness and timing-constraints of control measures that resolve specific causal factors¹⁸. By accumulation of data from different reported events, systemic causal factor clusters can be identified that point to specific, particularly weak areas in the organisation's management system or, similarly, in a man-made system's lifecycle phases.

Classification of incidents can be based on the type of harmful process deviation and furthermore on the set of relevant systemic causal factors within the framework of the causal model applied. A basic incident scenario combines a number of characteristic - systemic - causal factors that allowed the occurrence of the abnormal situation in the work process. Once a new type of incident has been reviewed, these characteristics can be used as a blueprint. By comparing a new incident with the available basic scenario's, 'known' incidents can be identified and consequently the yet unknown 'new' scenarios also. Thus by counting the recurrence of basic incident scenarios their incidence rates can be established. Furthermore, as each basic scenario is characterised by a set of relevant systemic factors, most likely, they can be grouped around common mode factors, which indicate the need for higher order solutions. For example, within the MORT model of systemic causal factors a basic scenario will mainly be characterised by factors from the S-branch of the MORT chart, see Figure A-1, whereas underlying common mode factors mainly are found in the M-branch.

Patterns of causal factors encoded in a data model might be identified as being typical for a specific organisational unit, a dedicated operation or, across organisations, for a specific corporate branch or industrial sector. Within one location, this process may be useful as different organisational units share higher management services and facilities and, thus, share common mode problems. Unit specific as well as common mode problems may be made visible. This premise about a causal factors pattern certainly played a central role in the harbour project, but formed also a working hypothesis in the ARA-pilot of organised learning co-ordinated from national level. Proof of the validity of this premise has not been found during these projects because not enough data have been collected to test it. Nevertheless, members of the ARA-team have indicated the need for branch-specific

¹⁸ Post-hoc risk control measures that can be implemented fast apply often to less systemic factors. More systemic ones usually require more effort and time. Furthermore it seems not necessary to resolve all systemic causal factors identified in a particular incident review: an organisation can still perform well if it knows its risk control weaknesses and decides to 'live' with these on the basis of acceptance of assessed residual risks.

extensions for the causal model in the ISA-PL system. They indicated that the lack of context information in records combined from companies in different sectors of industrial activity made it unlikely that patterns found by data analysis can be interpreted in a useful way, i.e. to suggest feasible changes which are appropriate across many technologies and types of company. What is common at that level are national, cultural, and regulatory processes which may only be changeable in the long term, if at all, and only by actors at the political or societal level.

The data models provide the clue to this step with classification being interpreted as matching to a known scenario or grouping of scenarios with common underlying organisational characteristics.

3.1.6 Computation

"Each near-miss tree as such generates a set of classifications of elements, which have to be put into a database for further statistical analysis... A steady build-up of such a database until statistically reliable patterns of results emerge must be allowed in order to identify structural factors in the organisation and plant instead of unique, non-recurring aspects." [Schaaf 91b, p. 30]

The NMMS claim is that large incident databases are needed in order to identify by statistical analysis structural organisational factors. In the projects described above, conclusions were drawn that such a premise may easily become of little practical value for a company at a certain location, if the annual number of reported observations is small. The events then become obsolete due to organisational or technological changes before the statistical reliability criterion can be met.

It is possible that databases that combine incident data from different sites or organisations may be useful to reveal common mode risk management problems at corporate or branch level, as been shown for occupational accidents in Denmark [Jørgensen 98], but this hypothesis could not be tested in the second company B project described in 2.2.4. Cumulation of process data across organisations makes sense only if processes and applied technologies are common. Likewise, causal cumulation is only meaningful if management systems are common or at least comparable. For instance, different companies within a corporation may share incident data to learn about common mode problems in the corporation's high-level management system. Companies may also share incident data when they run similar operations using similar technologies run by similar management systems. Data on product-related incidents can also be shared for a product class or per application sector.

As a section of the predefined MORT chart, the K-model of Figure 3-2 provides a means to systemically assess structural management factors even in single events analysis. By storing such cases in a database, data computation might generate added value by pinpointing areas of serious weaknesses in the organisation's management system as already discussed in section 3.1.5. This relies on the fact that organisational factors map on a many – many basis to incidents and therefore cumulate faster than the specific hardware or process factors.

The RSO-based projects in the medical centres are less suited to statistical analysis of structural factors. The anaesthesia project was aimed to test findings from other research

projects and causal factors were not yet linked to classes of management system factors. The database containing 549 cases did neither reveal significant 'harmful effect - causal factor' relationships nor show strong relationships between the causal factors tested. The value of computation is therefore largely unproven by these projects¹⁹. However, as we will see in Chapters 5 through 7, the computation of incidences of known incident scenarios provides a key to cost-effective double-loop learning.

3.1.7 Interpretation

The NMMS framework states that interpretation is the translation of statistical results into corrective and preventive measures. "Having identified such structural factors, the model must allow interpretation of these, that is: it must suggest ways of influencing these factors, to eliminate or diminish error factors and to promote or introduce recovery opportunities in the human-machine systems and indeed in the organisation as a whole." [Schaaf 91b, p. 30]

In the projects described, statistical analysis of causal factor data was only performed in the RSO-study within the anaesthesia department. Significant relationships between causal factors and adverse consequences for a patient could not be established. The aggregated data did show a substantial contribution of procedural errors to the occurrence of observed operational anomalies. This analysis result supported the proposition to invest in recurrent training using an anaesthesia simulator.

The heart valve prosthesis case study was triggered by the discovery of an unexpected mechanical failure mode of the product, i.e. single leg strut fracture, causing disc escape and resulting in a high magnitude of risk exposure to patients world-wide. This was the outcome of an epidemiological analysis [Graaf 92]. Recommendations on prevention of exposure to the fracture scenario risk were based on life-cycle analysis [Mol 95]. Recommendations regarding prophylactic replacement of suspected valves were to be based on the evaluation of the expected improvement of the health or wellbeing of individual patients weighed against no intervention. The causal analysis in the incident analysis project was limited to highlighting where and when the critical actions should be taken, not what they should be.

In the ARA-approach, the interpretation of data was built into the process of data collection, rather than delayed until a massive data collection had taken place. This heralded the development towards single incident interpretation by building in the causal model. The content-laden tree structure of the knowledge-model applied in the ARA project, see Figure 3-2, and also tested in a hospital unit, gave the opportunity to generate generic diagnostics and recommendations by reading back the appropriate tree branch. Because the analysis reports are stated in generic terms (see the examples in Appendices F and G), the actual interpretations into work process-specific recommendations remained to be done by a domain expert knowledgeable about the daily practice and processes, for instance by the local learning agency. In the ARA project, the domain experts were the

¹⁹ Note that in Tripod Beta, the assumption of accumulation is similar to that in the NMMS framework: accidents need to be stored in databases in order to allow Failure State Profiling of the company. "A Failure State Profile for an individual incident does not necessarily reflect the HSE 'health' of the operation under investigation, but composite Profiles obtained retrospectively by combining the latent failure categories from a number of incident analyses have been seen to correspond closely to those obtained proactively." [EQE 99]

project members from the companies. In the trial with Medical Centre B, the domain expert was the unit manager.

This approach of deployment of generic causation knowledge is much superior to epidemiological approaches. Epidemiological analysis, because it groups many disparate cases, leads to general conclusions, e.g. that 'more training is needed', 'personal protective equipment is not being used', or 'people slip over on pathways outdoors in winter'. What is missing in such findings is the link back to feasible changes in detail to the work process(es) generating the deviations. There is also no opportunity to test and reject a change in a given aspect of the situation on the grounds that it would be impossible within the essential constraints of the work process and its conditions.

The conclusion on this step of the NMMS framework is that interpretation can take place without, and long before, statistical analysis, if the causal and process models are good and complete enough.

3.1.8 Monitoring

The last module in the NMMS (monitoring) is meant to measure the effectiveness of proposed measures after their implementation. "...If the suggestions to management are accepted by management and actually implemented in the organisation, they will have to be monitored for their predicted as opposed to their actual results, that is: for their effectiveness in influencing the structural factors they were aimed at." [Schaaf 91b, p. 30]

In the projects described in this thesis, part of this monitoring function was designed to be implemented as a short-circuit. The NMMS was modified to be able to distinguish different classes of reported operational anomalies and recognise that a new occurrence fitted into a class for which recommendations had already been generated, i.e. it was 'known'. The lack of new notifications falling into this specific class of incidents could prove that a recommended measure had been implemented effectively.

None of the projects described in Chapter 2 considered the fuller implementation of this NMMS monitoring module. However, in the PIERCE project, described in Chapter 6, it became apparent that much can be gained if new notification reports can be compared with previous ones to collect this sort of proof.

3.1.9 The NMMS Framework revised

The review of the projects using the NMMS framework has demonstrated that several modules depicted in Figure 3-1 need revision in order to improve organised learning from small-scale incidents. The presumptions made in the NMMS modules 1 and 3 about completeness of event description need rethinking, as any event description is by its nature a simplified model of reality and therefore limited. Moreover, the detection of an incident that should be reported is negatively influenced by the burden imposed on the observer to provide a 'complete' account of the adverse event. The need to lower thresholds for reporting is pertinent. One way is to transfer the burden of event description and reconstruction to learning agents who are equipped to do a better job in event description once it is clear that something new has happened.

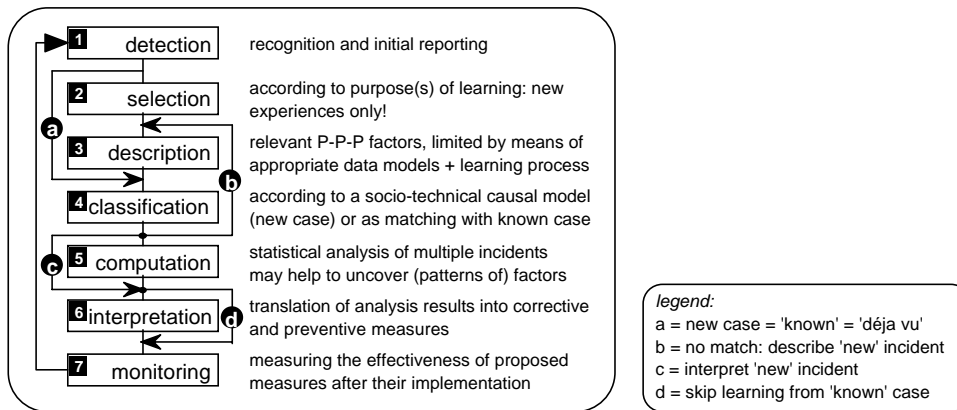


Figure 3-5: restructured modules for Near-Miss Management Systems

The projects have shown that the data model employed for classification is a corner stone in an effective NMMS. A combination of a process model for simplified event occurrence description with a generic model for analysis of systemic causal factors provides the strongest approach when investigative resources are limited. If the learning process is separated from the notification deed, modules can be postponed until special learning triggers are activated, e.g. a predefined number of a specific type of incident, or an apparently new type of abnormal situation. In that event the description and classification of the new incident lead to results that require interpretation rather than cumulation and the new case can be transformed into a 'known' one. Thus, the NMMS framework as depicted in Figure 3-1 can be redrawn as shown in Figure 3-5. Note that the scope of the modules 2, 3, 4, 5 and 6 has been modified compared with Figure 3-1. The principle of reuse of knowledge acquired from earlier incidents allows the short cut **a** when a new event is recognised a 'known' experience. This can be simply recorded and added to the incidence rate for this type of event. If this rate is sufficiently low, then interpretation of computational findings is not opportune: loop **d**. If the new event is not recognised as a 'known' experience then obviously it is a 'new' one, which deserves closer attention, i.e. description in more detail and classification (loop **b**). Subsequently, interpretation of case review findings is necessary to arrive at risk management measures scaled to the reported incident (loop **c**). The causal model applied for the classification allows cumulation of 'known' cases and their recurrence rates so that eventually statistical analysis of systemic causal factors also may become useful. Cumulation of data of multiple events provides opportunities to deepen the search for solutions of systemic causal factors.

3.2 Prospects from the Lessons learned

The basic concepts about purpose, contents, structure and configuration of incident recording systems have evolved over the projects described. The generic objective of all the projects has been to identify systemic opportunities for improvement of control of operational risks by problem owners, usually the management of organisations that own the processes in which the risks are contained. Initially, the incident recording and analysis system in the DOHO project aimed at external monitoring of adverse outcomes

of accidents that occurred in a set of companies in order to identify systemic causal factors underlying harm scenarios. That was the theme also in ARA and to a lesser extent of the heart valve project. However, the projects gave rise to legitimate doubts about the feasibility of this top-down type of approach. Large numbers of recorded cases are needed and, consequently, many harm processes have to occur before data analysis can start to become meaningful. Factors found through such data analysis cannot be interpreted into operational meaning relevant to monitored organisations, unless a common background is shared, for instance, where different companies run similar technological processes, or belong to the same governing corporation. In the same way, technological systems from different brands may share life cycles with similar risk control characteristics.

The shift has, therefore, been made to internal monitoring of unanticipated deviations in normal operations within an organisational unit in order to learn to improve intended processes. The initial objective of learning across individual organisations has moved to one of learning within an organisational sub-unit. The scope has widened from occupational - injury - accidents to any abnormal, unanticipated situation that eventually could result in harm or other losses. The objectives for incident reporting were reversed from a top-down process to individual and group learning in a bottom-up approach.

An organisation that wants to learn from its operational anomalies needs to organise the learning process and allocate learning functions. From the projects described, the first designated learning agent was found in the RSO-project within the anaesthesia department, see section 2.2.5. In the ARA-project described in section 2.3.1 a learning agency was envisaged at national level consisting of ISA-users supported by methodological experts. An equally important notion is that in order to learn from incidents, the organisation needs to become the prime user of its own incident data. In Chapter 4, a framework for organisational learning as described by Argyris is introduced and discussed in order to provide more theoretical basis for this process and the role of the learning agency in it. In Chapter 5, the concept of the Systemic Incident Notification System (SINS) will be developed from that theoretical insight and from the lessons learned from the projects described and discussed in this chapter.

Organisational Learning

4.1 Learning Process Framework

Learning from incidents involves *data* about such event occurrences that must be collected and analysed, i.e. be transformed into *information* about how to *change* the system in order to prevent incident recurrence²⁰. Options for change may be available or need to be created in order to select from them. The selected option(s) must then be implemented and the impact needs to be monitored and evaluated. The collection of data has been dealt with in previous chapters. The goal and embedding of collecting and processing (near-) accident data for 'learning for change' has been discussed in Chapters 1 and 3. The realisation and implementation of such a near-miss management system is the next challenge to face, since for instance the Company A and B as well as the Medical Centre A and ARA projects have made clear that implementation of a system for monitoring and learning from incidents so that it will be used, is far from easy.

This Chapter addresses the dynamics of organisational learning, i.e. how learning processes within an organisation occur and may be modelled, and what sorts of learning processes and blockages can be identified. These notions from Organisational Learning Theory are then translated into generic requirements for an organisational system to learn from incidents, accidents and other unwanted operational surprises.

Argyris et al. have developed and propagated a number of basic concepts about organisational learning [Argyris 96], which are introduced and elaborated here with regard to organisational learning from incidents. Argyris has performed research and consulting work on 'organisational learning' for the last twenty-five years. His work is of particular interest to the issue of learning from incidents for a number of reasons. Firstly, his application of the concepts of single-loop and double-loop learning, and learning how to learn ('deuterolearning') fits the idea in this thesis that incidents, as signals about an intentional process, can be used to produce feedback-control. Secondly, his concept of 'Theory-of-Action', which he conceives of as the simultaneous existence of a proclaimed 'espoused theory' and the really driving, but tacit 'Theory-in-Use', provides a key for better understanding of the problems encountered in the projects of reluctance to report incidents. His findings include the distinction between the prevalent, but defensive *Model*

²⁰ Bateson points out that defining 'learning' as change implicates the notion of repeatable contexts as necessary premise [Bateson 72, p. 292].

I Theories-in-Use, and the *Model II* Theories-in-Use, which enable second-loop learning. Passiveness throughout all levels of an organisation to learn from these unintended operational anomalies can be understood through this approach. Thirdly, Argyris builds his systems approach on basic cybernetics as presented by William Ross Ashby [Ashby 60], on the work of Gregory Bateson [Bateson 72] regarding learning, and on that of John Dewey about 'inquiry' [Dewey 38]. These approaches are compatible with the systems theory framework which is adopted in this thesis.

Central in his work is the '*theory-of-action*' [Argyris 74]. Actors design their actions to achieve intended consequences and monitor the results to learn whether these actions are effective. Actions are realised within an environment that may have been constructed for this purpose. Action strategies are selected such that the intended consequences can be produced within the assumed world that constitutes the environment. If the actual consequences deviate from the intended ones, adaptation of action strategies or the assumptions about the environment may be sufficient to correct the (first-order) process error. *Single-loop learning* does precisely this: detection of the error in producing the intended consequence leads to a change in action strategies or in the assumptions about the world relevant to the action. This results in turn in either another mismatch or the intended consequences. In the latter case, the learning is effective. If not, the governing intentions and related values may need to be changed in order to achieve consequences that are satisfactory after all, thus requiring a *double-loop learning* process. For the distinction between single and double-loop learning Argyris refers to Ashby [Ashby 60]. Ashby distinguished adaptive behaviour of a stable system, in which all essential variables lie within their normal limits, from the change in value in effective parameters which changes the field within which the system seeks stability, e.g. the thermostat of which the temperature set point is changed. The translation of these concepts from individuals to organisations elucidates the limitations in the organisation's capabilities for organisational learning. Organisational learning theory addresses these issues, see sections 4.2 and 4.3. The organisational learning framework provided is interpreted for this thesis in section 4.4 in order to arrive at generic requirements for a system that enables organised learning from small-scale incidents.

4.2 Organisational Learning Theory

Argyris et al. have addressed the question of what it is about an organisation that enables it to learn in an enlightening way.

"The term learning either means a *product* of the learning process - 'something learned - or the *process* that yields such a product... An organisation may be said to learn when it acquires information (knowledge, understanding, know-how, techniques or practices) of any kind and by whatever means. In this overarching sense, all organisations learn, for good or ill, whenever they add to their store of information, and there is no stricture on how the addition may occur. The generic schema of organisational learning includes some informational content, a *learning product*; a *learning process* which consists in acquiring, processing, and storing information; and a *learner* to whom the learning process is attributed." [Argyris 96, pp. 3-4]

A special form of learning consists of 'unlearning', e.g. the removal of an obsolete or incorrect strategy from an organisation's store of knowledge. An important type of organisational learning consists of an organisation's improvement of its task performance over time and is called *instrumental learning*.

"From a normative perspective, instrumental organisational learning should be taken only as a point of departure. Instrumental learning may be good or bad depending on the values used to define 'improvement'. The distinction between single- and double-loop learning differentiates instrumental learning within a constant frame of values from learning to change the values that define 'improvement'. [Argyris 96, p. 4]

The question now emerges *how* organisations learn. Since organisations are collectivities made up of individuals, one may argue that an organisation learns something when its individual members learn it. But this will not hold if individuals leave the organisation, and with them, the lessons learned. In order to determine the conditions under which the thought and action of individuals become distinctively organisational, the concept of *organisational action* is introduced.

"The idea of organisational action is logically prior to that of organisational learning, because learning itself - thinking, knowing, or remembering - is a kind of action, and because the performance of an observable action new to an organisation is the most decisive test of whether a particular instance of organisational learning has occurred...

Organisational action cannot be reduced to individuals that make up the organisation, yet there is no organisational action without individual action." [Argyris 96, p. 8]

The missing link between individuals that make up an organisation and organisational action lies in the rule-governed behaviour of individual members of the organisation's collectivity in the crucial aspects for its operation and survival.

"Before an organisation can be anything else, it must be 'political', because it is as a political entity that the collectivity can take organisational action. Then it is the individuals who decide and act, but they do these things on behalf of the collectivity, as its agents. And in order for individuals to be able to decide and act in the name of the collectivity, there must be rules that determine the boundaries of the collectivity, when a decision has been made and when authority for action has been delegated to individuals. Insofar as members of a collectivity create such rules, which we call 'constitutional', and become a *polis*, they have organised... By establishing rule-governed ways of deciding, delegating, and setting the boundaries of membership, a collectivity becomes an organisation capable of acting." [Argyris 96, p. 9]

In durable co-operative systems²¹, an *agency* is a collection of people that makes decisions, delegates authority for action, and monitors membership, all on a continuing basis. Formal organisations are usually complex task systems that may be tightly or loosely coupled, rigid or variable, but fall within the basic definition of the conditions for organisational action: they are co-operative systems governed by the constitutional principles of the *polis*. Thus, prerequisites for organisational learning are now known.

²¹ Temporary or 'ephemeral' organisations may arise spontaneously, e.g. in response to a crisis, and dissolve as quickly afterwards.

"If a collectivity meets these conditions, so that its members can *act* for it, then it may be said to learn when its members *learn* for it, carrying out on its behalf a process of inquiry that results in a learning product...

Inquiry does not become organisational unless undertaken by individuals who function as agents of an organisation according to its prevailing roles and rules." [Argyris 96, p. 11]

The authors point out that the term 'inquiry' is used here in a more fundamental sense that originates in the work of John Dewey [Dewey 38]:

"the intertwining of thought and action that proceeds from doubt to the resolution of doubt. In Deweyan inquiry, doubt is construed as the experience of a 'problematic situation', triggered by a mismatch between the expected results of action and the results actually achieved." [Argyris 96, p. 11]

4.2.1 Theories of Action

The output of organisational inquiry may take the form of a change in thinking and acting that yields a change in the design of organisational practices. Two distinct but complementary forms in which such knowledge becomes 'organisational' can be recognised. Knowledge can be stored in a way that is accessible, or can be assimilated in processes of organisational action.

"First, organisations function in several ways as a *holding environment* for knowledge, including knowledge gained through organisational inquiry... (in people's minds, in files, maps, as well as in physical objects).

Second, organisations directly represent knowledge in the sense that they embody strategies for performing complex tasks that might have been performed in other ways... Organisational knowledge is *embedded in routine and practices* which may be inspected and decoded even when the individuals who carry them out are unable to put them into words." [Argyris 96, pp. 12-13]

Such organisational task knowledge can be represented through so called 'theories-of-action' which have the advantage of including strategies of action, the values that govern the choice of strategies and the assumptions on which they are based. In a general form a theory of action states:

IF intent = consequence C in situation S THEN do action A

or abbreviated: (**A→C | S**). A theory of action may take two different forms, that of an 'espoused theory' that is explicitly advanced to explain or justify a given pattern of activity, or that of a 'theory-in-use'.

"By a theory-in-use is meant the theory of action²² which is implicit in the performance of that pattern of activity. A theory-in-use is not 'given'. It must be construed from observing the pattern of action in question. From the evidence gained by observing any pattern of action, one might construct alternative theories-

²² Note that theories of action apply to organisations as well as to individuals. Members of an organisation construct their own, inevitably incomplete, representation of the theory-in-use of the whole and strive continually to complete this picture by re-describing themselves in relation to others in the organisation. As conditions change, they remake these descriptions themselves.

in-use which are, in effect, hypotheses to be tested against the data of observation. In the case of organisations, a theory-in-use must be constructed from observation of the patterns of interactive behaviour by individual members of the organisation, insofar as their behaviour is governed by formal or informal rules for collective decision, delegation, and membership...

Like the rules for collective decision and action, organisational theories-in-use may be tacit rather than explicit and tacit theories-in-use may not match the organisation's espoused theory. An organisation's formal documents, such as organisation charts, policy statements, or job descriptions, not infrequently contain espoused theories of action incongruent with the organisation's actual patterns of activity." [Argyris 96, pp. 13-14]

These theories-in-use have particular relevance as they form implicitly the real-time reference for task execution by individual members of the organisation. One can see that the tacit theories-in-use, including norms and values regarding performance are embedded in the organisation's behaviour. On the other hand, an espoused theory-of-action may have nothing to do with the manner in which the organisation operates. It may be no more than the philosophy that is trotted out when someone from outside poses (embarrassing) questions, or what the organisation would like others to believe that they do.

The link between individual and organisational learning may be best summarised as follows:

"Organisational learning occurs when individuals within an organisation experience a problematic situation and inquire into it on the organisation's behalf. They experience a surprising mismatch between expected and actual results of action and respond to that mismatch through a process of thought and further action that leads them to modify their images of organisation or their understandings of organisational phenomena and to restructure their activities so as to bring outcomes and expectations into line, thereby changing organisational theory-in-use. In order to become organisational, the learning that results from organisational inquiry must become embedded in the images of organisation held in its members' minds and/or in the epistemological artefacts (the maps, memories, and programs) embedded in the organisational environment.

The learning products of organisational inquiry may take many forms, all of which, to qualify as learning must include evidence of a change in organisational theory-in-use. Often such changes are mediated by lessons²³ drawn from inquiry. [Argyris 96, pp. 16-17]

In other words, organisational learning not only requires provisions to enable the learning process through its individual members, but also results in verifiable changes in organisational theories-in-use and preservation of learning products of organisational inquiry of operational surprises.

²³ Argyris et al. give the following examples among others: interpretations of past experiences of success or failure; inferences of causal connections between actions and outcomes and their implications for future action; descriptions of conflicting views and interests that arise within the organisation under conditions of complexity and uncertainty; critical reflections on organisational theories-in-use and proposals for their restructuring; and description and analysis of the experiences of other organisations.

4.2.2 Single- and Double-Loop Learning

When inquiry after an operational surprise leads to improvement by changing the assumptions of the organisation or the way in which activities are carried out, without changing targets of task performance, the learning results from a single-loop learning process.

"By single-loop learning we mean instrumental learning that changes strategies of action or assumptions underlying strategies in ways that leave the values of a theory of action unchanged. For example, quality control inspectors who identify a defective product may convey that information to production engineers, who, in turn, may change product specifications and production methods to correct the defect... In such a learning episode, a single feed-back loop, mediated by organisational inquiry, connects detected error - that is, an outcome of action mismatched to expectations and, therefore, surprising - to organisational strategies of action and their underlying assumptions. These strategies or assumptions are modified, in turn, to keep organisational performance within the range set by existing organisational values and norms. The values and norms themselves remain unchanged.

Single-loop learning is sufficient when error correction can proceed by changing organisational strategies and assumptions within a constant framework of values and norms for performance. It is instrumental and, therefore, concerned primarily with effectiveness: how best to achieve existing goals and objectives, keeping organisational performance within the range specified by existing values and norms." [Argyris 96, pp. 20-22]

Single-loop learning leads to improved organisational task performance without changing objectives, norms and values. This contrasts with double-loop learning.

"By double-loop learning²⁴, we mean learning that results in a change in the values of the theory-in-use, as well as in its strategies and assumptions. The double loop refers to the two feedback loops that connect the observed effects of action with strategies and values served by strategies. Strategies and assumptions may change concurrently with, or as a consequence of, change in values. In some cases, the correction of error requires inquiry through which organisational values and norms themselves are modified, which is what we mean by organisational double-loop learning..."

For example, because there are so many maintenance accidents, a company decides to change the process by which the maintenance procedures are made, because the company is clearly missing out on risks or is making procedures which are impossible to implement.

"It is through double-loop learning alone that individuals or organisations can address the desirability of the values and norms that govern their theories-in-use." [Argyris 96, pp. 21-22]

We conclude at this stage that double-loop learning covers any change of values of current theories-in-use.

²⁴ Argyris et al borrow the distinction between single- and double-loop learning from Ashby's *Design for a Brain* [Ashby 60, pp.71-75]

The organisational learning process loops are depicted in Figure 4-1. The learning starts when the actions within the operational unit produce a consequence that is unintended and unexpected: the consequence mismatch or operational *surprise* may be detected by an individual member of the unit or elsewhere. When detected by an individual, he or she may take corrective action aimed at process stabilisation: the lessons learned this way then may modify the individual's version of the applicable theory-in-use. *Individual single-loop learning* may be effective for the individual in his or her task performance, but does not last within the organisational unit of which the individual is a member: learning products produced by the individual are not adopted in the organisation's theories of action.

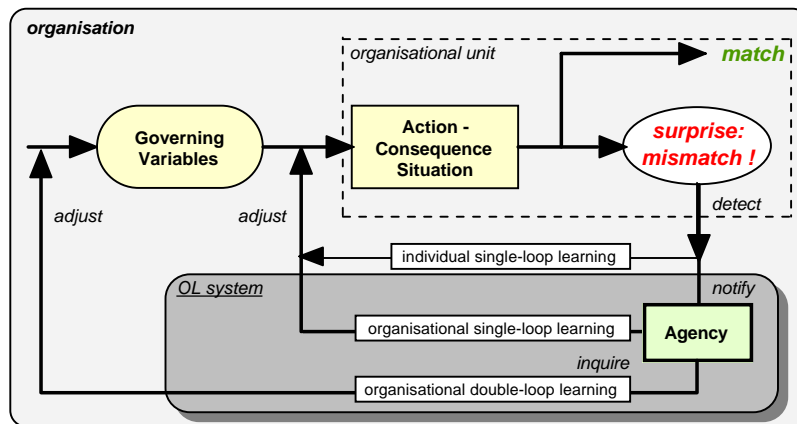


Figure 4-1: Single- and double loop learning after Argyris [Argyris 82]; 'detect', 'notify' and 'adjust' indicate the three channels discussed in section 4.5.

The diagram of Figure 4-1 indicates that learning has not occurred until a match or a mismatch is produced. Therefore, learning may not be said to occur at the point where someone (acting for the organisation) discovers a new problem or invents a solution to a problem. Learning occurs only when the invented solution is actually produced [Argyris 92]. *Learning* includes the effective implementation of solutions to the problem encountered: the learning process 'ends' when the consequences of action match the expectations according to the applied theory-in-use. *Notification* is as a signal a prerequisite for organisational learning from instances of - single - 'surprise' experiences by individual members of the organisation²⁵. *Organisational learning* from individual surprises is based on *organisational inquiry* by an 'agency'. *Single-loop learning* affects strategies of action and underlying assumptions in theory-of-action, i.e. changes in the way operational goals are achieved without changing goals, 'external' values or resources... *Organisational single-loop learning* products are visible in the organisation's theories of action, e.g. as minor modifications in a task protocol. *Double-loop organisational learning* affects norms, values and organisational targets as specified in the governing variable for applicable theories of action. Such changes actually mean changes

²⁵ In this context, Bateson identified the phenomenon of 'Zero learning' which takes place when a message is received by an individual: the option to change the context that existed before receipt of the message has become apparent [Bateson 72]. In the case that the message signals an operational surprise then the response option is to notify this instance to a learning agent. Without 'Zero learning' there is no learning.

in constraints of realisation of organisational functions by the operational process under review. The organisation's function is

"...to decompose double-loop learning issues into single-loop issues, because they are then more easily programmable and manageable. Single-loop learning is appropriate for the routine, the repetitive issue – it helps to get the everyday job done. Double-loop learning is more relevant for the complex, non-programmable issues – it assures that there will be another day in the future of the organisation." [Argyris 92, p. 9].

Organisational single-loop learning can also be triggered from outside operational units, for instance, through HAP triggers²⁶ or analysis of Performance Indicators²⁷. Note that organisational single- and double-loop learning requires a 'designated' learning agency, i.e. someone who learns on behalf of the organisation, or ensures that the learning experience becomes embedded in the organisation.

Example of double-loop learning in incident analysis

An experienced gantry crane driver in a steel works started dropping loads without any obvious reason. The supervisor instructed the crane driver to get focussed again. After a few days, the same crane driver dropped another load, and again another day. The supervisor seriously considered getting the crane driver fired. Then a P&O worker just passing through the workplace noticed another dropping... When the supervisor told him about the recurrent droppings, the P&O worker asked to talk with the crane driver. He found out that the crane driver, who worked in isolation in his cabin the whole day, could not handle this much longer. By assuring daily social contacts during breaks requiring the crane driver to come down for the breaks, the droppings ended as suddenly as they had started. Not used to requesting a talk with his supervisor, the crane driver unconsciously applied the unexpected means of droppings to draw urgent attention to his labour conditions! The instructions by the supervisor following the droppings came from single-loop learning and aimed at a solution without changing goals and values governing the processes in this production unit. The P&O worker initiated a double-loop learning process by questioning motives for the crane driver that could explain his strange behaviour. As a result, values for work in this unit were changed by explicitly acknowledging human social needs, and the prevailing theory-in-use was adjusted accordingly by inserting breaks at regular intervals to enable the crane driver to climb down and chat with his fellow workers daily.

The distinction between double-loop learning outcomes for organisational theory-in-use and double-loop learning in processes of organisational inquiry is correlated with the distinction between first- and second-order errors. *First-order errors* in organisational theory-in-use, i.e. in action strategies and expected results or in underlying values and assumptions, are illustrated by, for instance, excessive costs or too many sign-offs for task approval. First order error conditions like vagueness, ambiguity, untestability, scattered information, information withheld, undiscussability, uncertainty and inconsistency yield organisational situations that individuals find threatening and that are hard to detect by single-loop learning. In defensive organisational learning these First-order errors conditions remain intact. *Second-order errors* arise in processes of organisational inquiry, e.g. in case of a failure to question current practices, allowing such first order-errors to

²⁶ The Hazard Analysis Process (HAP) triggers form one set amongst others that form the organisation's technical information system for management of operational risks. The 'Technical Information Systems' - depicted as SD1 - in the MORT chart is a good example of an organisational double-loop learning trigger outside/above operational levels that may affect prevailing theories of action significantly. [Johnson 80, Knox 92]

²⁷ E.g. as in TRIPOD [Groeneweg 98].

arise and persist. Argyris et al. refer to this condition as the Model I world (O-1), see the next section [Argyris 96]. Inquiry into how the organisation learns, leads to enhanced organisational learning²⁸ being a critically important type of productive double-loop learning.

4.3 Models I and II

Argyris et al. introduce the concepts of Models I and II in relation to theories-in-use to describe the behaviour of individuals in the face of organisational learning and prohibitions. Especially when facing embarrassment or threats, individuals behave defensively according to Model I. Productive organisational learning requires a shift from Model I to Model II. A similar shift is needed to ensure that the organisation's learning system moves from O-I into O-II mode. In O-II mode, current theories of action as well as organisational inquiry practice are questioned in openness and clarity. Productive learning requires a double-loop process to resolve First order errors that blindfold the learning agency in single-loop learning. Transition from an O-I to an O-II learning system requires that Model II prevails for the organisation's members. The distinction between Model I and Model II theories-in-use is summarised in Table 4-1. It can be concluded that a mature system for learning from incidents will require an organisation to operate with Model II. Organisations stuck completely in the mode of Model I will have great difficulty in learning, especially from such threatening situations as incidents which could lead to harm (and hence liability).

4.3.1 Operational Surprise - Learning Modes

The identification of operational surprises that have occurred depends on individual members of the organisation. Incentives may be given to motivate individual members of the organisation to notify operational surprises for organisational learning. Reporting of two categories of operational surprises is eligible for an incentive reward:

- *new* operational surprises that need initial assessment
- '*known*', but recurring operational surprises²⁹ when a pre-set incidence level is reached by the submission of the latest instance and, therefore, a double-loop learning process needs to be triggered (see also the discussion on the idea of cumulation in sections 3.1.6 and 5.2.6).

In 'single-loop learning'-mode, the operational unanticipated surprise condition might be resolved by changing strategies of action and assumptions in relevant theories-in-use without changing norms and target values. This means that the task outcome remains as intended, but the method of realisation may be changed.

²⁸ This type of learning about the organisation's learning system after a operational surprise is a critically important kind of organisational double-loop learning, which is also called organisational 'deuterolearning' or learning how to learn [Bateson 72].

²⁹ Recurring operational surprises are eligible for organisational double-loop learning as these instances can be considered as first-order errors in prevailing theories-in-use that need repair.

Table 4-1: *Model I and Model II theories-in-use [Argyris 96, p. 93, p. 118]*

	<i>Governing Variables</i>	<i>Action Strategies</i>	<i>Consequences on Behavioural World</i>	<i>Consequences on Learning</i>	<i>Effectiveness</i>
Model I	<p>Define goals and try to achieve them</p> <p>Maximise winning and minimise losing</p> <p>Minimise generating or expressing negative feelings</p> <p>Be rational (and not emotional)</p>	<p>design and manage the environment unilaterally (be persuasive, appeal to larger goals – but do not consider other's goals if they conflict)</p> <p>own and control the task (claim ownership of the task, be guardian of the definition and execution of the task)</p> <p>unilaterally protect yourself (speak in inferred categories accompanied by little or no directly observable data), be blind to impact on others and to the congruity between rhetoric and behaviour, reduce incongruity by defensive actions such as blaming, stereotyping, suppressing feelings, intellectualising</p> <p>unilaterally protect others from being hurt (withhold information, create rules to censor information and behaviour, hold private meetings)</p>	<p>actor seen as defensive, inconsistent, incongruent, controlling, fearful of being vulnerable, manipulative, withholding of feelings, overly concerned about self and others or insufficiently concerned about others</p> <p>defensive interpersonal and group relationship (dependence on actor, little help to others)</p> <p>defensive norms (mistrust, lack of risk taking, conformity, external commitment, emphasis on diplomacy, power-centred competition and rivalry)</p>	<p>self-sealing</p>	<p>decreased long-term effectiveness</p>
Model II	<p>Valid information</p> <p>Free and informed choice</p> <p>Internal commitment to the choice and constant monitoring of its implementation</p> <p>Be concerned with others</p>	<p>design situations where participants can be origins of action and experience high personal causation</p> <p>task is jointly controlled</p> <p>protection of self is a joint enterprise and oriented toward growth</p> <p>bilateral protection of others</p>	<p>actor experienced as minimally defensive</p> <p>minimally defensive personal relations and group dynamics</p> <p>learning-oriented norms</p> <p>high freedom of choice, internal commitment, and risk taking</p>	<p>disconfirmable process</p> <p>double-loop learning</p> <p>frequent public testing of theories</p>	<p>increased long-term effectiveness</p>

'Single-loop learning' is effective as long as these changes stay within permitted alteration limits, otherwise 'double-loop learning' is required in order to change norms and values that govern the task performance steered by the theories-in-use. In double-loop learning on behalf of the organisation objectives, settings or constraints for task performance within the relevant organisational unit are changed, and consequently theories of action, including prevailing theories-in-use must be altered to fit the new constraints. At this point we note that by default an operational surprise is evaluated in a single-loop learning process in order to resolve the abnormal situation. It may well be that the operational anomaly is overcome - also in future - by minor changes in a standard operating procedure. The observed situation then could be classified as belonging to the group of 'single-loop learning' instances. Otherwise, the surprise occurrence calls for a double-loop learning process. Thus, a first classification category of a 'operational surprise' could be 'mode of organisational learning' with values 'single-loop' or 'double-loop' learning.

In the case that single-loop learning from a single incident leads to a solution that is considered adequate for prevention of recurrence the incident, it is crucial for verification of the effectiveness to keep track of future operational anomalies. If the resolution conceived is not sufficiently effective under similar operational conditions, the type of incident will repeat sooner and later. By tracking occurrences of abnormal situations, an unexpected repetition of a type of incident can be recognised and eventually trigger a more systemic double-loop learning process.

Criteria for placing a new instance in one category or the other come intrinsically from the solution space that exists for the problem in the relevant theory-in-use underlying the operational surprise:

```

IF <solution(problem(A', S', C'))>  $\supseteq$  <solution space(theory-in-use(A, S, C))>
AND IF <norms and values (theory-in-use(A, S, C))> = <still valid>
      THEN <single-loop learning mode>                                     [1]
      ELSE <double-loop learning mode>

```

This says that if a solution for a problem in the application of a theory-in-use can be found within the current solution space of the theory-in-use, then single-loop learning is sufficient, otherwise double-loop learning is needed for finding an effective solution by changing the governing variables, i.e. values and norms.

Another classification situation is still left open, namely the case where the operational surprise has been experienced before and for which a single-loop solution has already been assimilated in the prevailing theories-in-use:

```

IF <solution(problem(A', S', C'))> = <theory-in-use(A, S, C)>
AND IF <norms and values (theory-in-use(A, S, C))> = <still valid>
AND IF <first-order error flag> = FALSE
      THEN <adjust first-order error setting>                               [2]
      ELSE <double-loop learning mode>

```

This says that if a solution for a problem in the application of a theory-in-use is known within the current solution space of the theory-in-use, reuse of the known solution is indicated until norms or values governing the theory-in-use become invalid or a first-order

error condition has been acknowledged. In all other cases double-loop learning is required for finding an effective solution.

From single-loop to double-loop learning

The latex gloves that rupture once or twice during a surgical procedure form a recurrent annoyance and increase occupational as well as iatrogenic risks. The prevailing known solution to the surgery unit is to have the gloves changed by a nurse of the operating team, which takes about five minutes and one pair of gloves every time. The recurrent rupture experience is a 'first-order error condition'.

Management of the surgery unit decides to set an incidence trigger level at 100 within a specified timeframe. If the incidence level is met, a double-loop learning mode is activated, for instance, by persistent negotiation with the centralised purchasing unit manager to buy surgical gloves of better quality. The purchasing department's governing value of cost-efficiency is with respect to surgical gloves changed into cost-effectiveness: the better gloves become available and the rupture-during-surgery incidence diminishes. Thus, the organisation has learned another - risk management - lesson. Note that the double-loop learning was initiated by the surgery unit, but affected the purchasing department!

The first-order error condition is indicated by repetitive occurrence of a specific operational surprise. As the specific experience is known and some solution has been adopted within the limits of the solution space of the applicable theory-in-use, organisational learning is not an issue, unless the recurrence incidence reaches a pre-set trigger level. If the trigger goes off, the first-order error condition must be resolved via a double-loop learning process, i.e. resulting in changes of norms or values in applicable theories-in-use.

First order errors can also display second-order errors in the double-loop learning process of a previous operational surprise (A', S', C'), resulting in inadequate diagnoses and subsequently insufficient adaptation of the prevailing theory-in-use (A, S, C):

$$\langle \text{solution}(\text{problem}(A', S', C')) \rangle = \langle \text{theory-in-use}(A, S, C) \rangle := \text{problem}(A'', S'', C'') \quad [3]$$

In this case, the new operational surprise (A'', S'', C'') can be tested according to algorithm [1]. Second order errors in organisational learning arise in processes of organisational inquiry, such as a failure to question existing practices that allow first-order errors to arise and persist [Argyris 96].

We have now seen that - in principle - any operational surprise can be assessed for the appropriate mode of organisational learning in single-loop or double-loop using a set of two algorithms.

4.4 Transition of Key-concepts from Theory to Practice

Learning from accidents and near-miss incidents is painful and essentially requires a productive organisational learning process. Unfortunately, theories say little about how to achieve productive learning. Therefore, from the perspective a production unit within an organisation, key issues from the Theory of Organisational Learning introduced in section 4.2 are reviewed in the light of daily work processes.

We assume an *organisational* entity or *unit* whose members perform to produce the unit's intended *output*, i.e. realise the unit's goals within the norms and values set by the organisation's management applying available resources: see Figure 4-1. Typically, a

member of the unit detects the occurrence of an operational *surprise*. However, detection may also come from an outsider, e.g. when visiting the unit or through monitoring the unit's performance. An operational surprise is experienced only when the ongoing process or intermediate states in product transformation deviate from the individual's expectations that are rooted in his or her image³⁰ of the unit's theory-in-use. Thus, the detection of an abnormal situation, i.e. the receipt of the message coming through from the operational reality that something is different, depends on the individual's expectations regarding state transitions in the ongoing processes. If this message gets through³¹ to the observing individual, options for adaptive actions are limited. The individual can decide to leave things as they are or to *initiate an adaptive change process* either unofficially or officially. Individuals often decide for the first, because that seems to be about the only way not to make oneself vulnerable towards colleagues or management: does not the messenger of the bad news easily get the blame or worse? The second option is to initiate a change process oneself, maybe jointly with other members within the unit, but without bothering the organisation's learning system – if identifiable at all. Individual single-loop learning results if the surprise condition is resolved by changes in the theory-in-use of the individual members involved. However, the most effective option for adaptive action is to get the *surprise message* through from the individual member to a designated *learning agency*. Thus we note that

- learning from operational surprises is directly associated with selection of the option to change a prevailing theory-in-use,
- an organisation does not learn from operational surprises when only individual single-loop learning takes place.

A special form of learning is *unlearning*. Unlearning is indicated when implemented lessons learned become obsolete, for instance, outdated. The unlearning process focuses on cleaning up the unit's espoused theory of action, e.g. as spelled out in a quality assurance manual. Omitting this form of unlearning might eventually degrade the unit's performance.

Another, more peculiar form of unlearning is needed when a recurring operational deviation and its correction is seen as normal instead of as an operational surprise, and thus as something which need not (or should not) initiate an organisational learning process. The attitude of acceptance of the not-normal as normal has roots in the prevailing safety culture, as in the macho culture that made woodworkers refuse the use of personal protective equipment when applying a chain saw: getting seriously injured was just a normal part of the job. Such a culture needs to be reversed, so that admitting to the problematic operational experience is seen as a sign of professionalism or skill rather than of weakness. This is 'deutero-unlearning' or learning to learn again.

Let us make an additional assumption. The organisation has decided that it wants to learn from its operational surprises and has designated learning agencies with the mission to learn on behalf of the organisation from detected operational anomalies through single and double loop learning processes. Every member in the organisation now needs to know that every messenger of the operational surprise is appreciated. They need to be urged in

³⁰ Different individuals will hold their individualised and, thus, different images of the prevailing theories-in-use. So, different individuals may judge a specific situation in a distinct way, e.g. as 'normal' instead of 'problematic'.

³¹ If the message gets through, Zero learning has taken place [Bateson 72]

advance to convey the message of an observed operational anomaly to the nearest learning agency. It is crucial that the message can be conveyed in a manner that is practical and invites the observer to pass it on to the learning agency. For instance, complex notification forms not only make the observer reluctant to report, but may also create a paper-based blockage in the administrative processing. A notification requirement for a comprehensive description to capture the operational context may increase the observer's reluctance to report. It may be rewarding to keep the learning within an organisational unit, because of the company-wide 'safety competition'. With reference to Figure 4-1, we therefore note here for future development that:

- in order to lower the notification threshold once an individual has experienced an operational surprise, the surprise-message content and carrying medium need further reflection so that they are minimally invasive and maximally effective;
- of various levels of learning, single-loop learning is closest to the organisational unit's operations and theory-in-use and so will be seen as the first option. This is fine in most cases, but may result in some blindness to the need for more radical change³²;
- an organisational learning agency as receiver of the notifications must be prepared to handle incoming messages about perceived operational surprises and not let them languish unused and without reply (which is demotivating for the reporter);
- message contexts differ for different originating organisational units and hence the reporting context must be accounted for when interpreting the messages.

The learning agency must decide about the *strategy* of learning when a surprise-message is received. Single-loop learning is the prime option and may be adequate for most of the detected operational anomalies. However, to make reporting efficient, the surprises at this level must be recorded in a very compact way, without context. Learning agents being members of the learning agency must compensate for the lack of context information in the surprise-message in order to understand it. This may be realised through *organisational inquiry* by someone from outside the unit, but at the cost of (limited) resources and perhaps of risking activation of a defensive attitude from within the unit. However, a significant increase of effectiveness of organisational learning is gained when the agents are well-informed about the organisational unit's processes, operational conditions and the prevailing theory of action, i.e. the espoused theory as well as the theory-in-use. We therefore note here that

- learning agents may come from the organisational unit itself, as long as they are assigned to learn for the organisation and are able (or assisted) to take a step back from their daily work to view it more dispassionately.

The learning process conducted by the learning agency in single or double loop mode results in a selection of options for change available within the unit's theory-in-use (single-loop) or in a change of values and norms that form the parameters within which the unit must function (double-loop). A special set of options for adaptive changes available in single-loop learning emerges from the opportunity to identify existing but hidden (from the organisation) operational knowledge about surprise-management. Such knowledge may exist in the individual theory-in-use of some members, e.g. as a result of individual single-loop learning. After identification of this knowledge, it may be adopted

³² In a steel works study, it was shown that workers always too rapidly decided to try harder to follow the rules or to add a relatively minor technical fix to the problem identified. The problems remained [Hale 94].

in the unit's espoused theory of action³³, e.g. as a protocol for a specific activity, so that other members also can apply this knowledge under similar conditions. In this way, existing operational surprise-knowledge acquired through individual single-loop learning may be upgraded to organisational knowledge that remains also when individual members leave the organisational unit for whatever reason. We note here that

- the conversion of valuable, but tacit theory-in-use surprise-management knowledge, acquired through individual single-loop learning, into products contained in the unit's espoused theory of action can provide a valuable opportunity for the organisation to capture operational knowledge as organisational asset;
- the learning agency needs to convey its *'lesson to learn'-message* about selected options for change to the organisational unit, so that the unit changes its theory-in-use as proposed.

The effectiveness of the implemented changes must be monitored in order to close the learning loops. The lessons are learned only when the actual operational behaviour matches the intentional behaviour, i.e. the right lesson was correctly implemented, or another mismatch-surprise is acknowledged. This second mismatch contains the message that the selected options for change were either inadequate or incorrectly implemented and an additional cycle of organisational learning is imperative.

So far, three messages have been identified that are essential in organisational learning: a 'surprise' experience from the process to the observing member; the surprise-message from the observer to a designated learning agency, and the 'lesson-to-learn'-message from the learning agency back into the organisational unit. The first two messages have primarily an attention-grabbing function – an unexpected situation has occurred – as the inquiry function is allocated to the learning agency. The third message must be understood and adopted by the organisational unit in order to complete the learning process cycle [Argyris 92, 96]. We note here that

- all of these three messages must exist and get across;
- the learning agency has a key role in the organisational learning process;
- assigned learning agents who are also members of or very close to the organisational unit are better equipped to get the 'lesson-to-learn'-message effectively across to the unit provided that they are accepted in their role by other members of the their unit: this will facilitate the storage, dissemination and assimilation of learning products.

Criteria for becoming a successful learning agent, as a member the unit, include long enough experience to see the consequences of changes in the theory-in-use, but not too long to suppress the awareness of the nature of ones own skills and the willingness to question these when needed. If it is not possible to find this combination in one person, it may be wise to split the function over two people, one inside the unit and one outside, who work in close collaboration.

The role and functioning of the learning agency needs further exploration. This is discussed in the next sections.

³³ The organisational unit's 'espoused' theory of action can be traced back to proclaimed goals and work process documentation like protocols, user manuals, QA handbooks, training contents and other artefacts. Their operational relevance is limited. Such documents will not be read and followed to the letter, unless a specific critical task requires the explicit use of a dedicated checklist, like during start or landing of an aircraft.

4.4.1 Organisational Learning Agency

The learning agency performs the organisational inquiry related to a reported operational surprise. The inquiry comprises the following activities.

Upon receipt of a surprise-message, a test must reveal whether or not a notified anomalous instance is a *new* type of experience. This is an unintended consequence C' in a given situation S following an action A, or a change in S or A resulting potentially in an unwanted consequence C'. If it is indeed *problematic* but known, then the coping strategy has been defined already. If the surprise is new then there is a need to inquire and *assess theories-in-use* for risks and opportunities for resolution of the anomaly, i.e. to define the strategy for situation management. The inquiry consists then of *assessing, learning and retaining* the lesson in organisational knowledge form. Note that 'new' applies to any combination of (A, S, C) that differs from known operational anomalies. For instance, a crisis condition in situation S_1 (with unwanted consequences) that was anticipated is 'known'. The same crisis condition in situation S_2 , which is not anticipated, is 'new' (although the most effective course of crisis management action may well be known elsewhere within the organisation). In short, if (A, S, C) produces the intended consequence, then (A', S, C), (A, S', C), (A, S, C'), (A', S', C), (A', S, C') and (A, S', C') all represent potential operational surprises which may require different adaptations of prevailing theories-in-use.

The main function of the learning agency is to learn the lessons and *retain* both the experiences regarding implemented lessons, and the lessons themselves on behalf of the organisation.

Organisational inquiries become unproductive when inhibitors like scapegoat seeking, doom-scenario-finding, or taboo-maintenance prevail and leave critical issues untouched.

4.4.2 Communication Channels

Three communication channels are vital in the organisational learning cycle, see Figure 4-1:

1. a perception or 'detect' *channel-1* that 'surprises' the individual observer,
2. a 'notify' *channel-2* between the surprised individual and a learning agency, and
3. a lessons-learned 'adjust' *channel-3* between the learning agency and the organisational unit that generated the surprise condition.

Channel-2 and channel-3 communications can be directly designed and systemically influenced by the organisation, but channel-1 communication depends much more on an individual's sensitivity to operational surprise experiences and the prevailing 'safety culture'. The safety culture or climate can be influenced, but climate changes advantageous for productive learning usually take a longer time [Argyris 96]. All three channels are critical in a closed-loop learning process.

The communication channels identified require special consideration. In particular the notification channel-2 is vulnerable, because the individual who experiences an operational surprise, may easily decide not to employ channel-2 because of inability or

negative reinforcements like observed punishment of the messenger, fear based on possible personal involvement, lack of feedback, or administrative workload.

The learning agency was identified as an essential element in an organisational learning system that must be present to capture the operational deviations and control options as organisational knowledge. In order to arrive at a useful lesson learned, the learning agency must inquire into the notification message and understand the relevant operational context. The subsequent feedback message from the learning agency through channel-3 to the organisational unit will need to be understood by the unit's members within the unit's operational context and must fit into the unit's theory-in-use in order to become effective. If the learning agency does not compensate adequately for the lack of contextual variety in the notification message, the feedback message may easily be misunderstood or not fit into the unit's theory-in-use. Hence, for greater effectiveness, the learning agency must embody operational domain knowledge.

4.5 Requirements for organised Learning from operational Surprises

In summary, this leads to the following requirements for an 'operational surprise-learning system':

- a) learn from single surprise occurrences (channel-1 dependency)
- b) simplify notification contents in channel-2
- c) compensate loss of context (variety) in channel-2
- d) organise learning by the organisation (set-up learning agency, channel-2 and channel-3 communication) and acknowledge the need to close a learning loop in order to realise learning
- e) fit recommendations for adjustment of the prevailing theory of action to prevailing theories-in-use and preferably in line with the adjustments to be made in the 'espoused' theory
- f) organise reuse of previous learning experiences (depository, access)
- g) retain a causal knowledge model for use through or by a learning agency.
- h) assign members of the organisational units involved as agents to a learning agency to facilitate compensation for the loss of operational context in channel-2 and to tailor channel-3 messages to the organisational unit receiving them
- i) identify inhibitors for learning, e.g. first-order error conditions and take these into account
- j) provide in organisational deuterolearning
- k) increase sensitivity of Zero-learning

Many of these requirements are at least partly covered in the NMMS-model discussed in Chapter 3. For instance, channel-1 communication is acknowledged in the NMMS model, see section 3.1.2. However, there is a gap in the NMMS model regarding the fundamental problem of loss of context variety in the notification. A learning agency might be assumed to lie behind NMMS-steps 3 – description (inquiry), 4 – classification, and 6 – interpretation, but its reliance on having a description that should include all relevant factors rather

indicates the contrary, see 3.1.4. A concept like that of the 'theory of action' is completely missing in the NMMS.

These considerations have resulted in the development of a Systemic Incident Notification System (SINS). In chapter 5, the SINS concept, its components and data flows are described in detail.

4.6 Conclusions

Organisational learning from operational surprises is a post-event activity, meaning that the event or conditions that formed the operational surprise for individual members of the organisation, has been brought under control before the process of organisational learning is initiated. An organisation learns through its individual members and must organise organisational learning in order to facilitate productive learning. The individual must be challenged to notify the learning agency about any mismatch observed. Lessons are learned when the learning process starts and closes. The lesson may well be another mismatch, conveying the message that the adjustments applied were not a solution to the initial surprise. Lessons learned by organisational inquiry about the operations of an organisational unit need to be adaptable to the unit's theory of action. Failure to do so will usually result in recurrence of the surprise. This is a context issue that is relevant both in the notification phase and in the generating and conveyance to the surprise-generating unit of the adjustment measures. The organisational memory must be accessible for members of the organisation. However, the context-specific features of lessons learned contained in the memory and accessed for retrieval must be taken into account, because reuse of lessons learned elsewhere requires similarity of operational contexts.

In Chapter 5, principles of organisational learning are applied to design a system to facilitate the set-up and implementation of an organisational learning system for cost-effective handling of operational surprises.

Systemic Incident Notification System: SINS

5.1 Synthesis of a System for Learning from operational Surprises

The central thread in this thesis is the issue of how an organisation can learn in a systemic way from its own adverse event experiences, such as accidents. The prime goal of such learning from small-scale³⁴ accidents, incidents and other operational surprises with potential adverse consequences would be to improve control of operational risks.

5.1.1 Data Collection - Lessons Learned

Our view on learning from accidents and other near-miss situations has evolved drastically over the years through the series of projects that we have discussed so far. In the early projects, it was believed that the learning could be realised by means of collection of many records with descriptive accident data in order to identify harm scenarios linked to causal factor patterns. Initially, the unsystematic and superficial acquisition of accident data for the purpose of causal factor identification and modelling was seen as an important flaw in accident databases. The ISA-tool facilitated event recording to increase the number of reported accidents. It was expected that inclusion of harm scenarios would help to focus collection of causal factor data. The idea was that performance indicators could be identified for any system level considered. However, causal factor modelling by means of many accident records did not pay off for many reasons, which have been described in chapter 2.

Learning to improve control of operational risks requires a relatively short duration of learning cycles for the system-in-focus (e.g. for a company learning results must come out within hours to months, for an industrial branch maybe between 3 months to 5 years). Data collection methods requiring a larger number of incident records to be sufficient for statistical validity of data analysis, are therefore inadequate for timely learning by a

³⁴ A distinction between small-scale incidents and major incidents seems to be that in the latter case single-loop learning will not be sufficient for adequate problem resolution. On the other hand, small-scale incidents may very well require a double-loop learning process.

particular organisation at site-level. A more fundamental point is that descriptive incident data as captured through forms in a record provide an incomplete model of the reported real-life system states and constraints, i.e. incomplete "*context*". This context must be known when data is interpreted for improving control of operational risks, i.e. when changes will affect or be affected by that context.

Data collection on the basis of a knowledge-model of causal factors allows learning from a single operational surprise. With an adequate knowledge model it should be possible to diagnose this one incident within the framework of the knowledge model and generate advice about resolution of the problem. Thus, instead of the need to collect data from many incidents before causal factor analysis can start to make sense, the knowledge-based approach allows single-event analysis without loss of opportunities for statistical analysis of multiple surprises. However, the kind of data elicited, i.e. judgements about operational conditions expressed in model items, requires the person doing the reporting to be a domain expert trained to link the causal model items with the relevant reality or context. The ARA project demonstrated the potential of this approach, but the project also revealed two other key issues, namely, that learning from accidents needs to be organised and the contexts of different industrial activities have impact on the requirements for the knowledge model employed.

From the considerations above we conclude that

1. extensive descriptive data on surprises captured in a notification form are not on their own useful for organisational learning from (near-) accidents;
2. in any surprise description, the variety of the relevant context of the surprise is almost completely lost which makes adequate interpretation of the descriptive data impossible if there is no compensation for loss-of-variety; and
3. organisational learning must be explicitly organised.

The compensation for loss of context, incl. context variety, in the notification data of a particular operational surprise can be realised by domain experts. These are individuals who are well acquainted with the intentional, normal operations of the organisational unit that runs the surprise-generating process.

5.2 Specification of SINS - Principles

SINS is designed to facilitate organisational learning from all sorts of operational 'surprises' by individual members. The organisation³⁵ must first decide and communicate to its individual members that it wants to learn³⁶ actively from their experiences in order to improve organisational task performance over time. It must demonstrate that members benefit from such organisational learning in order to motivate them to report operational

³⁵ In this context, the term 'organisation' may initially be limited to a department or production unit within a formal organisation such as a company or institution.

³⁶ In MORT, this is summarised as 'management vigour & example' [MA2-a10] as issue of 'policy' and 'policy implementation'.

surprises. Once individual members are willing³⁷ to share their operational surprises in order to allow organisational learning, the following enablers³⁸ can be identified:

- a *notification system* that enables individuals to report the operational surprises simply and efficiently;
- a single- and double-loop *learning process* - catalysed by organisational inquiry *agents*;
- an *accessible depository* containing the lessons learned, i.e. knowledge of previously experienced and consequently 'known' instances of operational surprises and applicable solution strategies.

5.2.1 Surprise Experience: Channel-1 Communication

Individuals differ in sensitivity and perception of operational surprises. The safety climate within an organisation has impact on prevailing theories-in-use. If the safety climate is poor, individuals will have great difficulty to recognise an emerging operational risk, and the more so when it comes to notifying. Strong organisational defences [Argyris 96] may be triggered where there are conflicts between theory-in-use and espoused theory or within espoused theory. A policy of open safety culture combined with a practice where individuals who have had two accidents in a year are cautioned and threatened with discipline, will result in such a conflict which blocks productive organisational learning. In such a case, individuals who experience an operational surprise (channel-1 communication), will make up their mind to ignore the surprise. Different strategies and provisions can be thought of to improve such a safety climate. Concerning identification and handling of operational risks, management may take a high profile, organisation-wide safety programmes might be launched to raise attentiveness and stimulate openness, or individuals may (have to) receive awareness training on the job.

5.2.2 Surprise Notification: Channel-2 Communication

Zero-learning [Bateson 72] takes place when an individual experiences an operational surprise. This confronts him or her with the decision to act to change the current situation, or to let it be. A decision for change means primarily action aimed at bringing the operational system back into a stable state, if needed. The subsequent triggering of an organisational learning cycle is less obvious. The decision whether or not to report an experienced operational surprise is a balance between the utility of reporting as understood within the individual's theory-in-use and (other) conflicting operational goals. If all reported situations will be processed for assessment and inquiry, then the surprise observer does not have to select 'reportable' surprises from operational conditions, which would represent a notification threshold. Hence, any operational surprise experience may be captured in order to trigger organisational learning whether or not the specific surprise experience is relevant after all.

³⁷ This topic of motivation will be addressed further in Chapter 6 in the section 6.6 on 'safety culture'.

³⁸ In organisational learning theory, enablers are organisational structures which facilitate organisational inquiry [Argyris 96, p. 28.]

For surprise notification, the threshold for signalling to the organisational learning system that an operational surprise experience has occurred, must be made as low as possible so that the act of signalling becomes an anticipated part of the prevailing theories-in-use within the organisational unit. The signal coming - through channel-2 - from the individual, who had the surprise experience, must *always* lead to follow-up action of one sort or another.

5.2.3 Surprise Lessons: Channel-3 Communication

The learning agency produces learning products and recommendations to adjust the theory of action of a specific organisational unit, either directly in single-loop mode or through a double-loop learning process. In either case, adjustment messages are conveyed through channel-3 from the learning agency in due course to the organisational unit(s) involved. Whether a change in norms and values in the essential variables of the unit results from double-loop learning or a single-loop learning process was invoked, the organisational unit must adapt to these interventions and update its theory-in-use and assess the operational outcome: match or surprise! For this purpose, it is crucial that the adjustment message conveyed via channel 3 from the agency to the organisational unit arrives intact to maintain the correctness and variety in the message. Communication channels and their transducers need to be designed for this purpose.

5.2.4 Notification Data

An organisation as entity cannot learn for itself. Only a learning agency, i.e. individuals who are assigned the task to organise the learning from these surprises for the organisation, can take on such an action systemically. In communication channel-2 between the individual and the learning agency the main message is that an operational surprise has occurred during some intentional process. This notification message needs optimisation. In order to keep the notification threshold as low as possible, it must be kept simple, but it must also be sufficiently detailed to indicate what happened where and when. The learning agency as receiver of this message must take the context of the reported surprise occurrence into account in order to understand the notification message (see 4.2).

The notification of an operational surprise regarding organisational functioning is in essence a *signal* from an individual member who - based on his or her perception of theories-in-use - has observed a deviation from expectations. The important part of the notification is the depiction of the identified operational *deviation*. If we were to require that a notification must also contain at least a description of applicable 'violated' theory-in-use, we would have created a serious motivational problem. Not the least problem is that, more often than not, relevant theories-in-use are tacit and sometimes cannot even be deduced by observations (e.g. [Argyris 96], pp. 12-15). Hence, the individual must work to make them explicit and describe them. The need to describe the violated theory-in-use is avoided by coding the report to an explicit throughput process step, see text box below. Experts in the system know from the code what to expect at this step. If the information is later to be used by, for instance, external non-experts, the theory-in-use can be reconstructed by interrogating domain experts in retrospect.

Throughput Modelling

It is sufficient for organisational inquiry to describe only the deviations from the expected state changes if the condition is met that the prevailing set of theories-in-use at the operational level where (**A**→**C** | **S**) are 'known' to inquiry agents, see 4.2.1. This condition can be met by modelling the prime operations of the organisation in meaningful operational stages that are distinctive regarding task execution, resources and, therefore also, are governed by distinctive sets of theories of action. The notification data then must include a reference to the appropriate process stage. See Figure 5-1 below for an example.

In this thesis, this modelling is called throughput modelling. A particular throughput phase is usually assigned to one organisational unit. The transition from one stage to the next one is typically based on a go/no go decision by the initiating (receiving or sending) unit.

The throughput stages of relevant hospital processes can be modelled from the admission to the discharge of individual patients according to a product life cycle. Peculiar in health care is that individual patients form the product. The means of adding value to the product during the process is by delivering care, diagnostic services and treatment. Once a patient has been admitted for treatment and care, he is supposed to come out again in an improved state of health or quality of life: product quality control practice does not allow the rejection of that patient before discharge (unlike in manufacturing industry).

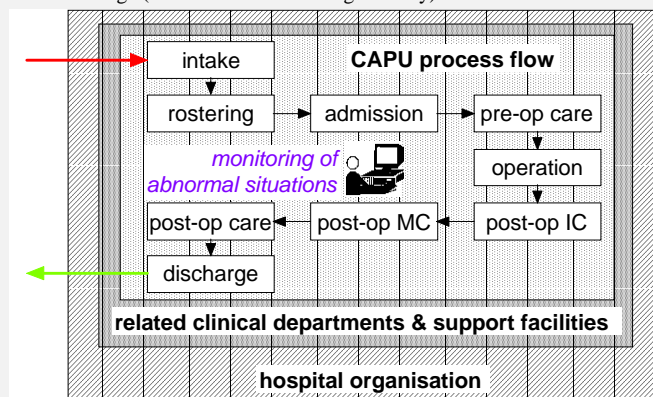


Figure 5-1: throughput model of a *cardiopulmonary* surgery department as 'seen' from the point of view of a CAPU patient. The monitoring of abnormal situations covers all stages

Once an incident has occurred, been identified and taken care of, this success and the learning experience will remain of very limited value if the learning remains at individual level, i.e. without utilising channel-2/channel-3 communications. Channel-2 communication is vulnerable. For instance, if operational personnel are given different objectives for reporting and review of different types of anomaly, the threshold for reporting operational surprises increases [Wickens 92]. Different objectives often go along with different reporting forms, from which the reporting person must choose. The evidence needed to select the appropriate reporting form will increase with each additional form. Each regime will have its own requirements for data. Consequently, the burden on individuals to discriminate appropriately between forms increases so much that it will be more rewarding not to report at all, except for indisputable occurrences (often already visible to the recipients of reports before the official notification). Thus, individuals should not be bothered by different objectives for data capturing when channel-2 communication is essential.

Once barriers and defences inhibiting individual members of an organisational unit from communicating an operational surprises have been overcome, the receiving end of

channel-2 must also be in place: each notification message must be handled with care. We need storage and retrieval facilities along with a learning agency. The notification module SINS-1 (see Figure 5-2, p. 81) is conceived as the channel-2 communication vehicle and includes a database as depository of the raw, i.e. original notification records.

5.2.5 Organisational Memory – Case Base

In order not to have to learn the same thing twice, the reported instance of an operational surprise must be compared as rapidly as possible with known occurrences in the past, preferably during the recording process. This means that new cases need assessment and validation through organisational inquiry resulting in validated, now 'known' cases to be added to the case base of a comparator. When a pre-set incidence level of these known cases is met, a double-loop learning process must be triggered. This means that the case, though known, is not solved and norms, values or resources for the applicable theory-in-use must be changed, possibly requiring changes in action strategies and assumptions as well. The algorithms [1] and [2] in section 4.3.1 apply here.

So as to compare a new operational surprise with previous situations, the previous instances must be stored in a depository – the case-base – that is accessible for instantaneous case comparison. The learning agency should govern the submission of a previously 'unknown' case in order to maintain the integrity of the case base. Requirements for case base data and their accessibility need to be defined. Then, a case base can function as a part of the organisation's memory.

The first action triggered by submission of the channel-2 notification message is that the learning agency tracks down relevant known cases and feeds back through channel-3 to the relevant unit(s) the lessons and recommendations from these to adapt the prevailing theory-in-use. By exploiting the case base for retrieval of similar surprises experienced in the past, and by retaining new, previously unknown surprise experiences to add to the store, the organisation's memory is mobilised for daily use. If a case similar to the new surprise is *retrieved* from the case base then the measures or solutions employed in the previous situation can be considered for *reuse*, usually after some *revision*. The revision, conveyed to the unit through channel-3, will work out more or less adequately. This experience of the success or failure of the recommendation can be fed back to and *retained* in the case base after verification of the impact, i.e. after completion of the learning loop. When a new experience has no match in the case base, the update of the organisational coping response requires active inquiry by the learning agency. The system must generate and validate a new way to adapt the unit's theory-in-use through a channel-3 message.

The case base comparator can be realised through methodologies of 'Case-Based Reasoning' (CBR) for which application development tools are becoming more and more available [Aamodt 94, Kolodner 93, Bareiss 89]. A very important feature of case-based reasoning is its coupling to learning. To a large extent, the driving force behind case-based methods has come from the machine learning community. Case-based reasoning not only implies a particular reasoning method, irrespective of how the cases are acquired, it also denotes a machine learning paradigm that enables sustained learning by updating the case base after a problem has been solved. Learning in CBR occurs as a natural by-product of

problem solving [Bergmann 98]. When a problem is successfully solved, the experience is retained in order to solve similar problems in the future. When an attempt to solve a problem fails, the reason for failure is identified and remembered in order to avoid the same mistake in the future. See Appendix H for a brief introduction to CBR. At this stage, I conclude that the issue of case representation, i.e. what are the essential features for learning from surprise cases needs to be elaborated from the perspectives of input to the case base and case matching during data entry. The different technological principles for CBR systems [Kolodner 93, Bareiss 89] for building a cost-effective system for learning from incidents have not yet been assessed.

Summarising, we conclude that at the receiving end of communication channel-2 the handling of notification data requires an accessible organisational memory in which previous surprise experiences and their solutions are retained, and resources for review of unknown surprises are present. The latter can be fulfilled by a learning agency, whereas case-base reasoning methodology can help to realise and maintain the organisational memory function that is envisaged as the case comparator module SINS-2.

5.2.6 Organisational Inquiry

In the case of a 'new' operational anomaly, the learning agents must validate and assess the case, and set an incidence trigger-level **IncT** below which single-loop lessons apply for specific theories-in-use. Repetitive recurrence of a 'known' operational surprise may become unacceptable even if the consequences are small, if the rate of occurrence is high. In that case norms or values governing the specific theories-in-use need to be revised in a double-loop learning process, resulting in a noticeable change in organisational task execution. Indicators for shifting from single- to double-loop learning are, amongst others, measures that were supposed to work but do not, and cumulated consequences that make the incident serious enough in aggregate to want to prevent recurrence after all.

The incidence trigger level mechanism provides the means for cost-effective inquiry into 'small-scale' incidents. Formula [1] applies to the incidence trigger level **IncT** of a particular case in the case base:

$$1 \leq \mathbf{IncT} \leq \infty \quad [1]$$

For **IncT** = 1 the recurrence of such an operational surprise requires immediate response. This operational surprise was not supposed to happen again after the first round of organisational learning! Organisational inquiry in double-loop learning mode, possibly requiring *in-depth analysis* of organisational causal factors is needed. On the other hand, for **IncT** = ∞, persistent reappearance of the reported anomaly can and should be ignored without any concern. In practice, such anomalies should not be reported, since they merely clog up the system. What is a suitable maximum **IncT** to be actively reported and retained in the system, will be a matter for discussion with the learning system.

Therefore, through the case comparator SINS-2, see Figure 5-2, any operational surprise reported in SINS-1 by individual members of the organisation can be evaluated cost-effectively. New cases are passed through to the learning agency for first-time assessment and validation, and so are the ones for which the incidence trigger-level **IncT** is met as a result of the comparison of new notifications that match with a particular case in the case base.

If a new surprise instance should not have occurred at all because of the extent of its adverse consequences, the learning agency must act immediately and start-up a double-loop learning process as if the new instance had been assessed earlier with a pre-set value for the incidence trigger level **IncT** = 1.

The learning agency may decide, for instance when the incidence trigger-level **IncT** is met, to go back and perform an in-depth analysis of all recorded instances of a particular (type of) incident, as part of a double-loop learning inquiry. Potential resources for in-depth inquiry encompass existing causal model knowledge models, e.g. the management oversight and risk tree (MORT). Such a knowledge model can be made available to a learning agency as an expert system application, similar to the one applied in the ARA project, see section 2.3.1. This is envisaged as the SINS-3 component of the learning system (Figure 5-2).

In summary, we conclude that organisational inquiry of operational surprises is run by a close interaction between the learning agency and organisational memory contained in a case comparator. The learning agency can scale the employment of scarce resources for in-depth inquiry, by setting an incidence trigger-level **IncT** for a particular type of operational surprise depending on the acceptability of its recurrence. Below this level, lessons from previously experienced similar surprises can be reused in a single-loop learning process moderated by a case comparator SINS-2 without direct involvement of the learning agency. Above that level, the learning agency initiates a double-loop learning process and performs in-depth inquiry employing expertise and other resources as needed. Knowledge on organisational causal factors underlying operational surprises can be made available as a SINS-3 expert system module.

5.2.7 Feedback on Notification

Feedback to the individual messengers (notifiers) from the case comparator after submission of a notification of an operational surprise should take place for motivational reasons. Feedback information relates to the classification of the operational surprise and the follow-up action strategy adopted as described above.

The feedback information on follow-up action depends partly on the contents of the case base. If case base data includes action strategy recommendations, e.g. known solutions, this can be offered to the reporting individuals. Usually, this will be too late for implementation on that occasion, but the information may reinforce or even refresh learning to cope. Also, if inquiry agents propose changes in action strategies or assumption of theories-in-action, parts of the implementation are retained visibly in the espoused theories of action, for instance, as changes in a task protocol.

5.3 Configuration of SINS

Now we can configure the building blocks identified above into a systemic incident notification system (SINS).

- A notification submission system with a database: SINS-1
- A case comparator and notification feedback system with a case base: SINS-2

- An agency consisting of individuals who learn through inquiry for the organisation and fill the SINS-2 case base, e.g. Review Team
- An expert system to support double-loop learning in organisational inquiry: SINS-3.

These building blocks are linked together for organisational learning as depicted in Figure 5-2.

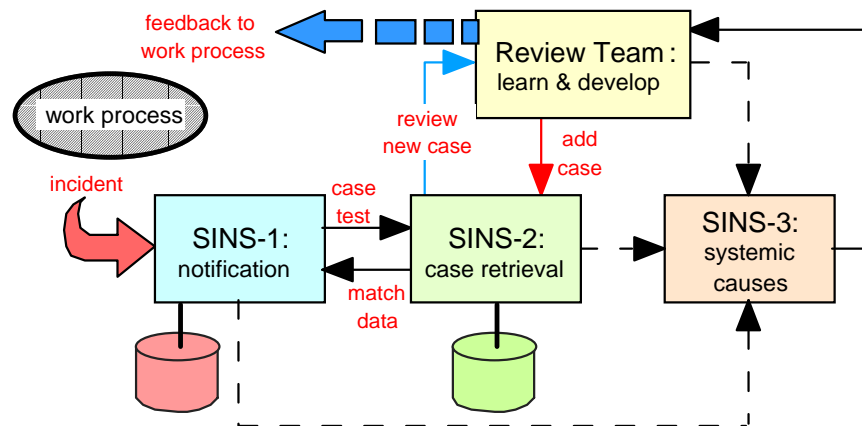


Figure 5-2: *the internal configuration of SINS*

Observed deviations in the work process of the organisational unit-in-view are reported to SINS-1 as an incident notification and stored in the SINS-1 database. Comparison of the new incident data with the case base in SINS-2 results in a 'match' or 'no-hit'. In the latter case, the learning agency, i.e. the Review Team in Figure 5-2, assesses the anomaly notification and produces a validated case with its incidence trigger-level **IncT** for the SINS-2 case base. If there is a match, then validated case data are made available to the SINS-1 user and the **IncT** test is executed: if the **IncT** value is met, a systemic causal factors inquiry (SINS-3) for organisational double-loop learning is triggered. The dashed lines indicate that a SINS-3 inquiry might also be triggered directly by the learning agency or through SINS-1 during the entry of the notification. External links to other units and organisational system levels are not depicted.

In the following sections this 'prototype' system is described as it has been developed and tested in the field in parts in several of the already mentioned projects. The most comprehensive set-up was realised in the PIERCE project (Medical Centre C), which will be described in detail in Chapter 6.

5.3.1 Notification and Reporting System: SINS-1

Ideally, the notification system SINS-1 would be a data recording utility for all members of organisational units where operational surprises may occur. Thus, the channel-2 notification message can be conveyed directly from the individual who experienced the surprise. In addition, the prevailing theories-in-use should include that operational surprises are logged by means of SINS-1. Therefore, the user-friendliness of SINS-1 is of

prime concern. The variety in contents of the notification message determines to a great extent how a SINS-1 subsystem can be realised. Less variety in the notification message through channel-2 requires more 'loss-of-operational-context' compensation upon reception. More variety in the notification message places higher demands on the notifier and does not necessarily solve the loss-of-context problem: the keyhole view remains to be validated when the message is interpreted.

Reported abnormal situations must be processed effectively and efficiently. Also, the reporter must not be discouraged from notifying. The surprise occurrence can be linked to a valid process throughput model, so that the message content is therefore limited to a description of the operational deviation. The description of this deviation must be sufficiently detailed to allow comparison with cases in the SINS-2 case base. Furthermore, a causal incident description model, such as based on the H-B-T accident model (see section 1.2.1), is used to refine case comparisons and speed up inquiry by the learning agency (see sections 2.3.1-2.3.2). The embedding of causal models for case data collection is feasible, see section 2.4. We have seen though in section 4.2.2 that in organisational learning the prime message of a notification is just the signal that an operational surprise has been experienced. Without any further details, the learning agency must consider the reported instance as 'new' and start an inquiry in order to compensate the lack of context information.

A paper-based notification system is doomed to fail, because filling in a notification form interferes, at least at first sight, with the primary process of the organisational unit. It takes valuable operating time. In addition, the prompt processing of such forms is difficult to ensure, because of limited resources. Using a computer-based reporting system can circumvent these problems. 'Intelligent forms' help to report any type of reportable abnormal situation without having to answer irrelevant questions or fill in forms twice for two different purposes, e.g. a notification to authorities and for organisational learning. Notification data is stored in a database for future use. The computer makes it easy to generate tailor-made notification reports and send these (semi-) automatically to defined stakeholders. Thus, the reporter can concentrate on the notification message instead of on report writing.

5.3.2 Case Review System - Learning Agency

A notification channel-2 communication message must be processed. The receiver is the learning agency who can cope with the lack of contextual data in the notification message. The reported surprise experience is then verified with the organisation's memory to establish whether a similar surprise experience has occurred before. This part of the review activity can be allocated to a computerised case comparator SINS-2 system, e.g. as a CBR-application. Lessons about management of a particular type of operational surprise, i.e. adaptation of a prevailing theory-in-use, once stored in the case base, increase the utility of case matching.

When a matched case is identified, case comparison also results in a quick assessment of the criticality or the potential risk of the reported operational surprise. If the criticality is sufficiently high, the review leads to instantaneous risk control measures and/or an in-depth investigation. Examples are operational deviations that could plausibly result in

disastrous consequences for the organisation or the external world. Also, case review leads to the discrimination of unique abnormal situations or new accident scenarios and the validation of known abnormal situation descriptions, including experience about the relative success of the intervention applied when such data are available. The resulting validated cases are stored in the case base for later retrieval. In addition to individual surprise review, the case review system also allows cumulation of descriptive and causality case data in the SINS-2 case base as well as of descriptive data that are collected over time and are stored in the SINS-1 database. This allows statistical use of notification data for priority setting or risk projection. Furthermore, it allows the verification of the operational value of cases in the case base, e.g. by analysis of the case match frequency, i.e. incidence per unit of time. This is vital for unlearning obsolete lessons when needed. Derivation of a root cause profile³⁹ for the organisation-in-view may be based on the causality descriptors added to a case by the learning agency, see section 5.3.3, and on case incidences or match frequencies. The frequency of new surprises that are still 'unknown', i.e. for which the notification message cannot be matched with a case from the case base, may serve as an indicator for the quality of descriptive notification data, i.e. the risk inventory and evaluation system of the organisation. Poor quality results in an unjustifiably high frequency of 'new' surprises. Thus, a learning agency has opportunities to generate added value to surprise notification by utilising SINS-1 and SINS-2 subsystems both in single case review and for multiple case analysis.

Lessons learned in the review process feed back to the notifier, the organisational unit generating the operational surprise and to relevant management levels elsewhere in the organisation. Finally, the learning agents generate criteria or triggers for in-depth analysis of abnormal situations, which are known, but recur for whatever reason.

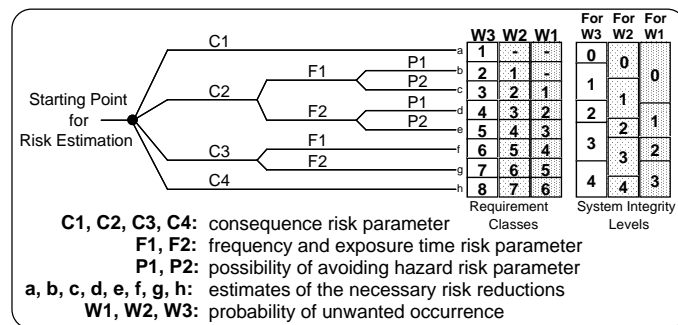


Figure 5-3: risk graph example in recent standards [DIN 94, IEC 98]

These review functions are distributed over organisational learning agents, e.g. designated review teams of domain experts, and computer-based subsystems. A quick assessment of the criticality of an abnormal situation can be made using a risk graph algorithm that is built into the notification system, e.g. as depicted in Figure 5-3. A notification reviewer, e.g. the unit's manager, may also use his domain expertise for assessment of the criticality of the incident. The review function of case validation, though, requires a learning agency consisting of work process domain experts with the input of methodological expertise. Once the case base is available, the initial review functions can be computerised in a

³⁹ like the GFT- or BRF-profile in Tripod...

manner supportive to the operators (SINS-2). The review team functions then move into areas such as supervision of the notification system, dissemination of lessons learned and analysis of underlying systemic causal factors (SINS-3).

5.3.3 Case Retrieval System: SINS-2

Descriptive data on the operational surprise are recorded by means of the SINS-1 subsystem. These data are compared with the cases in the SIN-2 case base. Such a case must have descriptive features that make the case comparable with the notification data. Before the SINS-2 case is added to the case base, the learning agency must link descriptive case features with systemic causal factor categories, such as listed in Appendix N. In this way, a SINS-2 case consists of descriptive features and, additional, 'causality' features filled in by the learning agency. The descriptive features also contain information about the intervention applied to handle the previous incident and, when the old solution is reused, feeds a single-loop learning process. Thus, when a notification message matches a case from the case base, the systemic causal factors identified in the previous surprise can be fed back to the reporting person in addition to the values of the descriptive case features. Furthermore, by aggregation of systemic causal factor feature values from the case base cases, an extra mechanism is created to trigger double-loop learning directed towards the most relevant root cause categories.

In the early phase of the SINS-2 subsystem, the case base is empty and, in practice, each new notification of an operational surprise will be 'unknown' to the case base. The learning agency must therefore review all notified operational surprises and then add a new case to the case base. Later, new surprise experiences will have become known, so that most of the notifications will match with a case, maybe except for some minor feature values. The SINS-2 comparator system then practically merges with the SINS-1 notification system.

When the new surprise situation matches with a case from the case base, the known case recognised can be presented to the reporting messenger who may decide to reject the case presented because the similarity is not close enough. If the case base includes data on interventions to handle specific type of abnormal situation or known lessons learned in the past from similar operational surprise experiences, these may be presented as suggestions for resolution of the surprise conditions reported. Again, the notifier can reject the suggestion as not appropriate for this instance and annotate the case to that effect. If case base matching takes place during notification data entry, early candidate case matches can be presented. Acceptance of a presented case can then be used to shortcut entry of notification data, e.g. when a particular anomaly is known to occur, but has been accepted for the time being, i.e. until the incidence trigger-level **IncT** has been reached. For instance, the occurrence of a rupture in a surgical glove in an operating theatre does not need an elaborate description of the deviation in the ongoing work process. The operational annoyance is 'known' and the most appropriate (but also most expensive) solution of better quality gloves requires a double-loop learning action that involves the purchase department outside the surgical unit-in-view. However, this step is not warranted if the incident is relatively rare. Meanwhile, the 'second best' solution of replacing the gloves can be followed. It could be suggested by the system, though this is hardly necessary in such a simple and frequent case. Meanwhile, the notification only needs to

adjust the appropriate incidence counter until the **IncT** trigger goes off, indicating that the consequences of a high glove-failure rate are great enough to warrant searching for a better type of glove.

New cases for the case base are found when there is no match at all, or when there seems to be a match, but it is judged not to be close enough after review of discriminating attributes. Acceptance of a proposed known case generates additional descriptors of the incident: there is a match (descriptor 1) with a specific case from the case base (descriptor 2). The incidence of specific case matches provides an objective, peer-review-based mechanism for triggering in-depth analysis of reported abnormal situations.

5.3.4 Organisational Causal Factor Analysis System: SINS-3

It may be necessary to perform an in-depth analysis of a specific type of operational surprise. In-depth analysis is indicated when the initial problem remains unsolved, the problem underlying the surprise or incident is not well understood, or the risks involved if there is a recurrence are so unacceptably high that the organisation must ensure it is prevented. In double-loop learning mode, the organisational inquiry focuses on systemic factors that contribute to the observed - recurrent - abnormal situations. Double-loop learning mode can become far-reaching⁴⁰, for instance, may require strategic changes at the highest levels of the organisation. Relevant systemic factors that can be altered by the organisation are prime candidates for lessons to learn about recurrence prevention. Variables governing the organisational unit's work processes will have to be changed and, consequently, prevailing theories of action will need to be adjusted more or less drastically.

The analysis of organisational factors underlying abnormal situations in intentional work processes is performed by means of a Management Oversight and Risk Tree⁴¹ investigation as has been discussed in Chapter 1 and demonstrated in Chapter 2. The structured model of management system functions provided by the MORT chart allows effective identification of relevant 'root causes' that are within the span of control of the organisational system-in-view. Diagnostic rules as applied by MORT investigators can be captured into an expert system to support organisational learning agents with this type of causal factor analysis of small-scale incidents, as has been demonstrated in the ARA project. The systemic review by means of a relevant causal model objectifies the inquiry by the learning agents (ARA, Medical Centres B and C). Thus, computer technology helps to make available systemic causal factor analysis expertise to problem owners. Rule-based inference engines help the user to perform a structured assessment of relevant issues based on a consistent knowledge model, and to generate diagnostic reports which may be based on rules provided by external experts (ARA, Medical Centre B).

⁴⁰ In terms of ambition levels in knowledge management, Argyris' concept of double-loop learning links with the third (out of four) level of ambition, that is where knowledge about work processes is built up and used for improvement of systems rather than of task execution [Spek 97].

⁴¹ Although MORT is taken as an appropriate method, others do exist, e.g. SCAT, Tripod [Groeneweg 98], REASON.

5.3.5 Stop rules for Notification and Analysis

As we do not know in advance whether the operational surprise is a precursor for potentially severe losses, rules are needed to distinguish relevant incidents from other abnormal situations. The development of these stop rules is based on the dynamic, incremental development of the SINS system. We apply a number of different mechanisms to arrive at a decision to perform systemic causal factor analysis in more detail. During entry of notification data, a risk graph algorithm (DIN 94; IEC 98) is used as a trigger. For instance, if the assessment of algorithm parameters results in a high-risk class judgement, a SINS-3 inquiry may be initiated: IncT_g in Figure 5-4 depicts such a high-risk operational surprise. A learning agency may also decide to initiate further - SINS-3 - analysis in the process of assessment of notified incidents. Finally, threshold criteria IncT regarding the number of matches of incidents with specified cases in a - SINS-2 - case base are set to recognise a recurring incident at a given frequency as one deserving an in-depth (SINS-3) investigation of underlying organisational causal factors (IncT_j in Figure 5-4).




Figure 5-4: organisational inquiry resources for a specific type i of operational anomaly: none for $\text{IncT} < \text{IncT}_i$, as much as needed when $\text{IncT} = \text{IncT}_i$

At any time, a line manager may decide independently from SINS-rules to initiate a detailed systemic causal factor analysis of an abnormal situation and for that purpose activate the organisational learning agency. The learning agency of an organisational unit and process owners, i.e. the unit's line management, base decisions about incremental development changes in the SINS system on consensus. Thus, the process owners as the prime users are and remain in full charge of their tailored SINS.

5.4 SINS Implementation Level - scaling down / scaling up

The actual realisation of SINS depends directly on the level of implementation within the system of the organisational structures or 'enablers' of organisational inquiry (Argyris 96). The example of the intelligent anaesthesia monitor, see text box below, demonstrates how the SINS-1 notification system could, in certain circumstances, be reduced to a set of one to three simple push buttons.

Scaling of SINS Components – IAM Example

The Intelligent Anaesthesia Monitor [Graaf 98; Eijkel 99; Vulling 99] employs real-time fuzzy set pattern recognition to identify physiological crisis conditions in an early stage. During the anaesthesia process data are logged automatically. During surgery, the anaesthetist continuously watches the patient and takes action if something is wrong. This action of the expert defines an alarm situation [Eijkel 95]. The goal of the IAM-project was to develop an anaesthesia monitor that supports the anaesthetist during surgery in order to improve the safety of the patient. This support should be based on patient signals measured during regular surgery. An important objective in the monitor design is the recognition of alarm situations. A requirement for this is that alarm situations should be explainable to the anaesthetist. For this reason a knowledge-based approach is followed.

The SINS concept fits in as follows. The Monitor contains the SINS-2 comparator function by means of a neural network inference system. Known sets of parameter values and trends are associated with specific crisis management triggers. If such a set occurs, the monitor alerts the anaesthetist and on demand, may provide support regarding crisis management. When an emerging crisis is detected by the anaesthetist – before the IAM produces an alert – the complete real-time notification message only consists of a few button pushes: alert (yellow), acute alarm (red), and crisis under control (green). If such a notification is done in real-time, than the logic push button pulses are recorded, synchronised with the anaesthesia data and automatically evaluated by IAM. The result of the initial notification (yellow followed by red or green; red followed by green) is either a delayed match with known crisis conditions by the neural network or the full description of a new type of physiological crisis.

After completion of the surgical procedure, a learning agency consisting of domain experts, e.g. members of the operating team, analyse the time-based recorded data and decide on appropriate crisis management measures. It may be concluded that the crisis relates to factors beyond the control of the anaesthetist, e.g. a quality assurance problem elsewhere or limitations of state-of-the-art knowledge, so that a double-loop learning process becomes necessary for a more permanent crisis management solution. At this level of system implementation, the SINS building blocks have the same functions, but their realisation qua contents and methods differs from the implementation at organisational unit level.

The SINS-concept also allows scaling up, as was partly attempted in the ARA project, see section 2.3.1, where a learning agency – the expert panel in Figure 2-4 – was projected at national level, and on-site SINS-1 notification was integrated with a SINS-3 module in the ISA system. The validated case base was mainly needed for incremental improvement of the knowledge model in the localised ISA-PL systems, thus sharing incrementally improved causal model knowledge beyond individual organisations. This was not implemented for reasons connected to project funding. However, with the insight gained since and certainly also through the ARA-project, application of the SINS-concept would have led to a modified conceptualisation of a national ARA network for sharing causal knowledge and lessons learned.

The important notion here is that the functions of the components and their interactions as captured in the SINS concept apply at any system level of implementation.

5.5 Conclusions

The SINS system concept has been unfolded in this chapter from the perspective of the organisational unit-in-view within an organisation. Key functions identified are notification message communication (SINS-1), case matching for management of inquiry resources and for reuse of lessons learned (SINS-2), case review and validation (Review Team/learning agency), and systemic causal factor investigation (SIN-3 to support learning agency).

Of the four SINS building blocks, three were adopted and adapted from predecessors in previous projects (Chapter 2): SINS-1 (intelligent forms), SINS-3 (knowledge model based diagnosis) and learning agents in different configurations (ARA, Medical Centre B, Medical Centre C). The former ISA-PL system combined the functions that now have been reallocated between SINS-1 and SINS-3. The new building block is SINS-2 that became necessary to capture and reuse lessons learned as organisational knowledge assets for management of abnormal situations. An equally important *raison d'être* for the SINS-2 building block is that it generates an incentive for members of an organisation to notify recurring operational surprises, because it mediates in an objectified manner the allocation of resources for organisational inquiry through the learning agency and the **IncT** trigger mechanism.

When the notification includes abnormal situations, such as near-miss incidents, the proposed application of databases opens opportunities to learn in a cost-effective and resource-efficient way from one accident or less, where "less" is a near-miss incident or a more exotic abnormal situation.

An adaptive system requires that routine monitoring provides a trigger to in-depth investigation. Whilst automation can certainly help the monitoring side of the equation by case-matching technologies, the role of human expertise to re-introduce the context variety in problem resolution, which was necessarily lost in that monitoring process, is essential to operational process adaptation and thus closure of the learning cycle.

In Chapter 6, the implementation of a SINS in a surgery department will be described and discussed in detail.

SINS Implementation: the PIERCE Project

6.1 Medical Centre C: PIERCE - learning from one Accident and less

The SINS concept was tested in a project that was set up within a cardiac surgery unit of a major hospital: the PIERCE project. PIERCE is an acronym and stands for 'Patient Incident Evaluation and Registration for Cardiopulmonary Environments'. This specialised surgery unit can be considered as an organisational entity that 'produces' cardiac surgery applied to selected patients, i.e. people with a cardiac defect who are evaluated as 'operable'. A description of this unit is provided below. The project provided opportunities to introduce low-threshold incident recording and organised learning under manageable circumstances. The system level of implementation was low (micro) compared to the national scale of the ARA project described in 2.3.1: the learning process was primarily at the same level as the occurrence of operational surprises. Of special concern was the high sensitivity to incident review in the medical sector, e.g. because of issues of medical liability.

Before the PIERCE project started, the unit was developing a medical complication registration system to record for each of their patients pre- and post-operative physiological problems and medico-technical counteractive interventions applied. These records did not contain any data about the operational circumstances of registered complications. The PIERCE project was designed to fill this gap. Cardiac surgery patients are intrinsically vulnerable, but should not get into further medical problems [Dijen 00] due to operational flaws in the work processes of the unit. A first step would be to introduce an operational surprise recording system for internal use. Therefore, the unit wanted an implementation of SINS, starting with the notification system and a learning agency. The following hypotheses were postulated:

Hypothesis 1: significant numbers of abnormal situations do occur within work processes of the cardiac surgery unit.

Hypothesis 2: these operational surprises can be successfully recorded for review without causing unacceptable pressure of time or concern over recording.

Hypothesis 3: beneath many of these situations lie systemic causes that can be identified in a systematic way, and of which some can be influenced by the management of the unit, and others can be passed on to other organisational units.

Hypothesis 4: the unit can effectively learn from these abnormal situations and change its theory-in-use.

In this chapter, the PIERCE project is discussed. The introduction of systemic incident review for learning lessons required great care, as the vulnerability of patients as well as clinical workers is very high. The main focus was directed to operational abnormal situations in all phases of the unit's work processes.

First the organisational unit is described. Then the project is outlined and discussed in detail. Finally lessons about implementation are drawn.

6.2 General Description of the Cardiac Surgery Unit (Medical Centre C)

Cardiac surgery is performed on patients with cardiovascular problems who are indicated for surgical intervention. Operations include placing bypasses, implantation of heart valve prostheses, and about 50 patients per year are treated for lung surgery. Patients come in by referral (scheduled), as emergency patient from another treatment (usually scheduled) or through the cardiac emergency aid unit (unscheduled). The unit in Medical Centre C does not perform heart transplants. Annually, about 900 patients are admitted for cardiac surgery.

The unit consists of a surgical staff, a secretariat, a ward providing normal pre-op and post-op care, as well as post-op medium care. An Intensive Care Unit (ICU) provides vital post-op intensive care. A specialised Coronary Care Unit (CCU) provides pre-op support on demand and operates the in-house telemetric monitoring system. The anaesthesia department and the operating theatre unit participate in the unit's surgical procedures, and vital services are drawn from the blood bank and the laboratory. Close working relations exist with the cardiology department. Residents and co-assistants support the surgical staff. Residents assist during surgery and are acting physician on the wards and in the outpatients' department. The nursing staff on the regular and medium care wards is specialised in cardiac surgery patients. The ICU has about 30% of its capacity available for the cardiac surgery unit. On average, a cardiac surgery patient stays up to 36 hours in the ICU before being transferred into the medium care unit for another 24 hours. The shortest overall stay for any successful cardiac surgery patient is 5-6 days from admission to discharge.

The wards function on a 24-hours/day basis with three shifts/day. Scheduled surgery is usually on weekdays only, i.e. Monday – Friday. Three operating rooms are fully equipped for cardiac surgery procedures. Per room, up to two operations may be scheduled per day. Senior surgical staff is usually on-call overnight and during the weekend. Regular meetings exist for medical staff [Sivro 97] and, separately, for nursing staff and for perfusionists, see also Table 6-1. Some of these sessions are set-up together with other disciplines, i.e. with the ICU-staff, respectively with cardiologists.

Table 6-1: *regular meetings in the cardiac surgery unit*

meeting	frequency	present	topic
morning session	daily	surgical staff + residents	all patients that require special attention + irregularities
heart team	daily	surgical staff + cardiologists	classify up to 25 patients: reject, cardiology, operate
capu ward	daily	nursing staff	all patients that require special attention + irregularities
heart team	weekly	surgical staff + cardiologists	assessment of problematic patients
Heart team	monthly	surgical staff + cardiologists	medical issues
surgery team	monthly	surgical staff + anaesthetists + ICU-staff	medical issues
staff	monthly	surgical staff	policy, general topics
ICU staff	daily	ICI staff	all patients that require special attention + irregularities
perfusion	quarterly	perfusionists + cardiac surgery unit manager	irregularities + technical issues

The wards function on a 24-hours/day basis with three shifts/day. Scheduled surgery is usually on weekdays only, i.e. Monday - Friday. Three operating rooms are fully equipped for cardiac surgery procedures. Per room, up to two operations may be scheduled per day. Senior surgical staff is usually on-call overnight and during the weekend. Regular meetings exist for medical staff [Sivro 97] and, separately, for nursing staff and for perfusionists, see also Table 6-1. Some of these sessions are set-up together with other disciplines, i.e. with the ICU-staff, respectively with cardiologists.

Incidents or complications may be discussed during these meetings, but this is typically done in isolation within the respective professional groups: doctors with doctors, nurses with nurses, residents with residents, and perfusionists with perfusionists.

6.2.1 Complications and Operational Surprises

The handling and review of so-called 'complications' had, up to initiation of the PIERCE project, focused on the medico-technical intervention to solve the physiological problem. Before the PIERCE project 'lessons learned' were neither recorded nor disseminated among relevant clinical workers. The possible linkage between the occurrence of complications and work process flaws in preceding stages of the cardiac surgery work processes had not been looked into systemically, despite the fact that this could lead to the identification of opportunities to prevent the development of certain types of complication. The existing complications database system recorded complications data of all patients that actually reached the operating table. If a patient died earlier or for some physiological condition had to be handed over to another discipline first, this was not recorded in the complication database. Nor was the effectiveness of corrective interventions recorded in the database.

The following working definition of the phenomenon of medical 'complication' was formulated early in the project:

A complication is an unplanned change in a patient's physiological system that usually requires intervention either to prevent irreversible harm to the patient, to neutralise reversible harm to the patient, or possibly to reduce the loss of availability of health care facilities.

Hence, a medical 'complication' is nothing less than a patient-related accident. The definition says no more about causation of complications than that the physiological system of a patient is directly involved in the harm process. The PIERCE project was designed to take up these aspects and apply the following risk management principles to them. The likelihood and harm potential of complications can be assessed to some extent in ways similar to that of other accident classes. Where the care providers can not (yet) influence causal factors underlying a complication, the associated risk of occurrence should be subject to a process of explicit acceptance. Table 6-2 lists complication types that are common in cardiac surgery.

Before PIERCE, data on operational surprises in work processes were hardly ever recorded for systemic review. It was felt that useful lessons could be learned if the prevailing focus on patients were to be complemented by a reflective view on abnormal situations in the full range of operational activities that members of the unit provide for care and treatment delivery.

Table 6-2: *complications common in cardiac surgery*

cardiac failure (as pump)
blood leakage
kidney function problem
infection
pneumothorax
atrial fibrillation
rethoracotomy
pacemaker problem
false aneurism
ulnary nerve compression

The range of these activities is broad. The surgical unit actively runs almost the whole set of processes that result in the cardiopulmonary surgery product. Other hospital units assist on demand, e.g. the staff of the operating theatres, anaesthesia teams, the coronary care unit and the clinical laboratory. The cardiopulmonary surgery process consists of distinguishable stages before, during and after the operation. These stages influence each other and have their own set of process characteristics and constraints. The mutual dependency of these stages is large. For instance, a surgical procedure can only start if a bed in the post-op ICU is available; a shortage of nursing staff or prolonged stays of patients reduces the availability of beds. Hence, the whole system is vulnerable to operational surprises, as the interdependent process stages are tightly coupled. Incidents occurring in the unit's processes affecting patients or others deserve to be detected and recorded for review and to answer the question what could be learned from such events. Thus, the focus in SINS is on operational disturbances in the cardiac surgery processes. See section 6.4 for a more detailed description of these processes.

6.3 PIERCE Project

The PIERCE project included the following work packages (wp):

1. the *modelling of the phases, steps and ownership of the cardiac surgery work processes* in order to simplify notification of operational surprises: see section 6.4;
2. a Reported Significant Observation (*RSO study*) [Eicher 76] to assess the general sensitivity to and perception of operational risks: see section 6.5;
3. a *safety culture measurement* (a baseline assessment in an early stage of the project and a progress measurement some time after introduction of the notification system SINS-1): see section 6.6;
4. the development of a *low threshold operational surprise recording system SINS-1*: see section 6.7.1;
5. specification and development of the *case base system SINS-2*;
6. redesign of the *ARA-PL system*, see 2.3.1, *for knowledge-based diagnosis* as SINS-3;
7. initiation and launching of a '*review team*' as a learning agency: see section 6.7.2;
8. *dissemination activities* in order to keep all those involved informed, in order to gain commitment.

The first phase included the development and launch of the operational surprise occurrence notification system SINS-1, and the setting-up of a learning agency initially as RSO-team, in a later stage renamed the Review Team. Thus, the early project stages were devoted to the work packages 1, 2, 3, 4, and 7. One of the residents was designated as a learning agent.

Work package 2 was a critical incident review study wherein members of the cardiac surgery unit recalled operational surprises with and without adverse consequences. We can link the work packages to the testing of the hypotheses formulated in section 0. The test of hypothesis 1 (significant numbers of abnormal situations) is obvious. Falsification would require the absence of any situation, observed or reported, that would be out of line either with the prevailing theory of action or with the system's goals. Hence, the number of incidents reported in or by means of work packages 2, 3, 4 and 7 would have to approach zero.

Operational surprises are recorded to some extent in the work packages 2, 3 and in SINS-1 after realisation of work package 4. The learning agency that is implemented by work package 7 also records cases for review. The testing of hypothesis 2 (surprise recording) mainly depends on the effective use of the recording system SINS-1 after its implementation.

Systemic causes underlying an operational surprise are designed to be discovered from case review. Case review is initially assigned to the learning agency in place, e.g. the RSO-team that is part of work package 2, respectively the review team installed by work package 7. The testing of hypothesis 3 (existence of systemic and influenceable causes) consists of the generation of lessons to learn, i.e. recommended measures to implement and the efficacy of implementation of such measures. The former is primarily done by the learning agency whereby the latter also depends on the effectiveness of the

communication about lessons learned with the unit's members, which is an essential part of work package 8. Thus, hypothesis 3 can be tested through work packages 7 and 8.

The process of organisational learning as depicted in Figure 4-1 involves the unit's members, who should signal the learning agency when an operational surprise is experienced. Ideally, hypothesis 4 (unit can learn) implies that the organisational learning process is in place and active. This can be observed by the initial surprise recording following work package 4, the generation of new lessons learned by the review team (wp 7), reuse of lessons drawn in the past (wp 5) and the subsequent adaptation of the prevailing theory-in-use (wp 8) by applying lessons learned. Work package 3 is meant to measure progress and blockages in the unit's learning processes.

Work package 1 is necessary for the development of a low threshold surprise recording system (wp 4) to allow simplicity in notification data, see section 5.2.2, and enables the formulation of another hypothesis that can be put to the test in work package 7:

Hypothesis 5: causal factors underlying an operational surprise in phase P_i can often be influenced more easily in a preceding phase P_j ($j < i$).

This hypothesis was added because of early experience in the project with review of the incidents often pointed out opportunities for improvement of control of operational risks earlier in time. Work package 3 includes elicitation of incident examples recollected by the interviewees. These examples should indicate measures taken to prevent recurrence. The review team (wp 7) will have to look for operational conditions that allow surprises to occur. Projection of the risk control measures or opportunities from the work packages 2, 3 and 7 onto the previous process phases coming out of work package 1 will prove or falsify hypothesis 5.

Work packages 2 and 3 were designed to capture the prevailing attitude towards 'safety' and identify pitfalls for implementation of a SINS within the cardiac surgery cluster.

In summary, the PIERCE project aims to record and learn from operational surprises in the cardiac surgery work processes. The focus is on the operation of these intentional processes by members of the unit. The identification and resolution of underlying systemic causal factors in the PPP-domains People - Plant - Procedures, see section 3.1.1 [Nertney 75, 87] is sought. These operational control factors must be distinguished from the physiological problems underlying the development of non-intentional instabilities in a patient's physiological system. The latter are usually called 'complications' and require adequate intervention. Nevertheless, causal factors of such a 'complication' might also be found outside the patient. Gaba calls the response to an identified complication 'crisis management' [Gaba 94].

6.4 Cardiac Surgery Process Modelling

The main purpose of modelling the organisational unit's processes is to facilitate the location of the operational context of surprises in one of the identifiable process stages. Starting point for modelling the cardiac surgery work processes is the individual *patient* with a cardiovascular disease. At a given time, the patient is accepted for cardiac surgery, thus entering the cardiac surgery system. Some time later, the patient leaves this system in one way or the other. Scheduled patients expect to have a fair chance of survival with an

increased quality of life. However, organisational factors and operational flaws can contribute to the mortality rate due to iatrogenic fatalities. For instance, a patient may die on a waiting list or through some lack of communication when crisis management intervention becomes imperative. The care providers need to detect emerging problems and have several resources for this purpose, like monitoring equipment, protocols for critical process stages, a test laboratory and administrative controls. To do this, the process stages are modelled from the perspective of the organisation's *raison d'être*, i.e. the individual patients.

6.4.1 Time-sequenced Process Phases

The product of cardiac surgery is realised in work processes by the broad cardiac surgery cluster supported by other hospital units, see Figure 6-1. The patient comes in for a surgical procedure, and must be prepared for this treatment, and needs to recuperate to some extent before discharge. Thus, from the perspective of patients the cardiac surgery processes are seen as time-sequential phases. The members of the cardiac surgery cluster are the active operators who manage risks in real-time while performing their cardiac surgery process tasks. Main features that discriminate one phase from another are process-ownership, operational risks and decision rules for transition to the next phase. Additional relevant distinguishing features of phase steps might be personnel:patient ratio, equipment:patient density, and dependency on support from other organisational units.

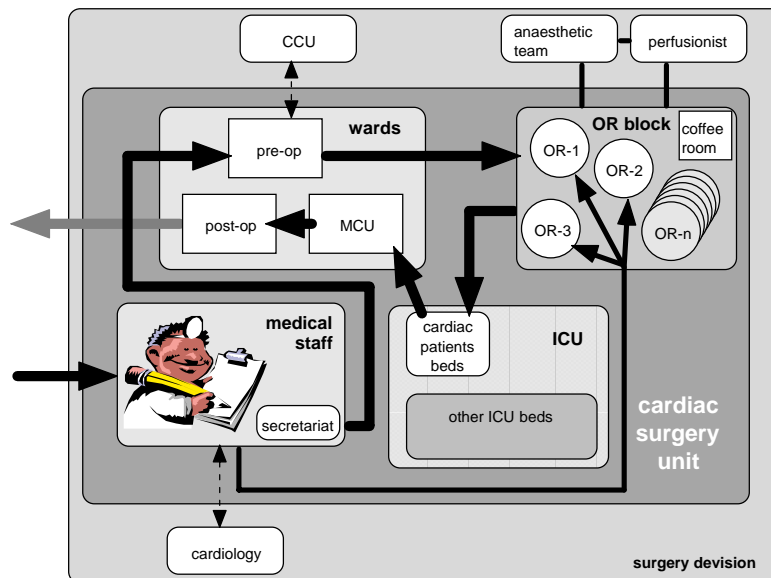


Figure 6-1: main entities in the cardiac surgery unit. Patients accepted for surgery are 'processed' in the white boxes within the cardiac surgery unit box (see the text box 'Throughput modelling' on page 77 for a process description)

From a patient's view the main time-sequenced process phases of the core processes of the cardiac surgery cluster are preceded by a stage of cardiovascular symptom development, Phase 0, see Appendix I. Phase 1 – intake – is the stage where the cardiac surgeon

becomes involved for the first time and the patient must be evaluated for the most appropriate treatment, whether this would be cardiac surgery or not. If surgery is indicated, the subsequent process phases become relevant, see Figure 6-2.

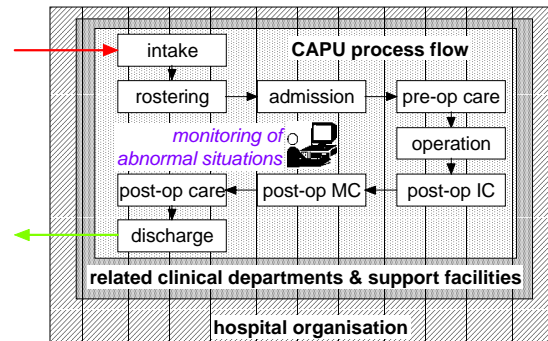


Figure 6-2: the cardiac surgery unit in a patient's perspective of subsequent process phases

The table in Appendix I summarises (in Dutch) the main phases and embedded critical steps, process owners and associated operational risks which are described above. This resulted in the simplified throughput model shown in Figure 6-2. Phase 2 – rostering – is the stage of administrative preparations to ensure that all the necessary resources and means are available and to prioritise patients on the waiting list according to criticality and the elapsed waiting time. This stage ends with an admission call to the patient. In Phase 3 – admission – the patient is hospitalised and prepared for the surgery in Phase 4. After completion of the surgical procedure, cardiac patients are transferred into Phase 5 – post-op ICU – and kept sedated for at least ten hours, thus being prepared for immediate corrective surgical intervention if required after diagnosis of an emerging 'complication' in the patient's physiological condition. As soon as possible, the patient is transferred from the post-op ICU to Phase 6 – post-op medium care unit (MCU), including the subsequent post-op care – until discharge in Phase 7. For each transition from one phase to another clear 'go-no go' decision-criteria or 'golden rules' could be identified, e.g. a patient in post-op ICU cannot be transferred to the MCU if breathing is supported by mechanical ventilation.

The phases and embedded critical phase-steps as well as the ownership transitions provide concrete links for domain experts to process contexts relevant to the reported operational surprises. In other words, simply indicating which phase a surprise occurs in triggers rich contextual knowledge by an expert, who can then recognise and understand very brief descriptions of the surprise, without the need for detailed description.

6.5 RSO Study

A Reported Significant Observation (RSO) study was set up in accordance with Eicher [76]. The RSO method is directly adapted from the 'critical incident technique' [Flanagan 54], whereby the words 'critical' and 'incident' have been replaced by 'significant' and 'observation'. It is an information gathering technique, which uses workers' participation to describe situations that they have personally witnessed involving good and bad practices, and safe and unsafe conditions. The questionnaire method – as opposed to 'interviewing' –

was selected to allow people to answer the questions at any time convenient to them and in peace and quiet, see Appendix L. It could also contrast with the interview method that would be applied in work package 3. The questionnaires were designed to capture risky operational surprises that either were managed adequately due to equipment, other devices, a relevant protocol and the proper use of these resources, or resulted in a harmful or loss-generating event. Each participant was asked to answer questions about two examples of 'near-miss' situations witnessed and recovered from and two examples of situations that actually went wrong. The objectives were to gain insight into operational risks, concrete hazardous practices and opportunities for improvement of management of these risks. An RSO-team consisting of three experienced surgeons, a resident, and an investigator was assigned for a quick analysis of imminent hazards, i.e. the left-hand route in the lower section of Figure 6-3, and for standard analysis of the returned observations, i.e. the right-hand route in the lower section of Figure 6-3, and for standard analysis of the returned observations, i.e. the right-hand route in the lower section of Figure 6-3,

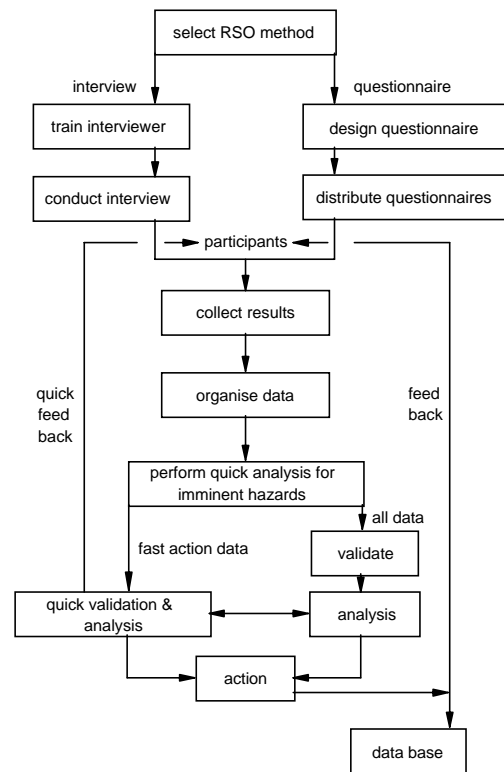


Figure 6-3: RSO process flow [Eicher 76]

This team could be expanded with other domain experts when needed. The blank sets of 2 * 2 RSO-forms were distributed to twelve experienced members representing all areas within the cardiac surgery cluster. Of these only two sets were returned despite several reminders. Both were basically answered adequately, but the observations were of negligible urgency. The RSO-team was therefore never activated.

At first sight this result seemed to falsify hypothesis 1. We did not find many RSOs. However, this seemed to contradict the impressions already gained from interviews and from the earlier hospital projects. Two explanations for this poor result are possible. One is that medical care providers are focussed on managing today's real-time problems as they occur and do not have time for retrospection and reflection; hence, it is hard to recall problem situations from the recent past. The other is that the line of questioning in the RSO-form may have induced defensive behaviour because of the details and assessment asked. It was therefore decided not to give up at this stage, but to consider the two returned sets of RSO-questionnaires as not sufficient to justify an assessment of hypothesis 1. The examples submitted could be interpreted when seen within the originating context of - in both cases - the ICU. They were not incidents where intervention could still be effective before harm was done.

6.6 Safety Culture Assessment

Partly as a result of the failure of the RSO-study a safety culture baseline assessment study was performed [Guldenmund 00]. This study was also designed to allow measurement of progress in organised learning within the cardiac surgery cluster. This study was focussed on issues concerned with individual perception and attitudes regarding incidents and the controllability of causes within their daily work. The expectation was that the introduction of a system for systemic notification and review of such incidents would lead to a shift in the perceptions and attitudes of personnel involved. The key questions for the PIERCE project were:

1. What kind of cultural barriers and organisational defences are to be expected that might cause difficulties for implementation of a systemic recording and analysis system (SINS) ?
2. Is there a change in the perception of accidents and near-miss incidents that can be attributed to the implementation of a systemic recording and analysis system (SINS) in the cardiac surgery cluster (incident perception concerns the nature, causes, and controllability of incidents and their occurrence)? If so, in what direction is this shift?

The study also served to develop methods for a related project on safety culture [op. cit.], which is not further discussed here. A validated questionnaire to assess safety culture in hospitals was not available. Hence, it was decided to present the perception issues by means of semistructured interviews of personnel from the different cluster units. An interview protocol was developed for the baseline assessment given the key questions above and notions about relevant issues indicated in literature, see Appendix K. The protocol covered the topics 'personal background', 'safety awareness', 'incidents', 'accountability', 'openness', 'communication' and 'work satisfaction' [Guldenmund 97].

The protocol was tested by means of a trial consisting of three interviews by two interviewers who independently worked out interview reports. Comparison of the reports showed that the interviews provided reliable answers to the questions posed. Fifteen people from the cardiac surgery process were selected on the basis of their familiarity with the work, their dedication and their willingness to express their opinions. Persons involved in the RSO study were excluded for several reasons. Firstly, both work packages would generate similar information concerning individual perception of risks, thus allowing reciprocal validation about risk perception. Furthermore, it was envisaged that work packages 2 and 3 could run simultaneously and demands on members in the cardiac surgery units would be too great. And finally, work packages 2 and 3 also aimed at raising awareness for the introduction of the notification system SINS. Each person volunteered to be interviewed. Table 6-3 depicts the distribution of interviewees over job categories and gender, while Table 6-4 shows the organisational units where the interviewees worked.

Table 6-3: *distribution of interviewees over gender and profession*

	Physician	resident	nurse	total
male	8	1	3	12
female		1	2	3
total	8	2	5	15

Table 6-4: *distribution of interviewees over organisational units*

cardiology	Thorax wards	CCU	ICU	OR	total
1	4	2	3	5	15

The influence of possible investigator biases was kept to a minimum by deploying different interviewers and by providing feedback to the interviewees so that they could check and modify the report on their interview.

6.6.1 Outcomes of the Baseline Assessment

A summary of the main findings is given in the following sections structured along the five main protocol topics. Regarding 'work satisfaction', all persons interviewed displayed and expressed satisfaction with their job in the cardiac surgery unit.

6.6.1.1 Safety Awareness

'Safety' is primarily associated with 'patient safety' and the prevention of iatrogenic harm; 'occupational safety' is strongly associated with infection control concerning patients and the individual worker. It is believed that safety awareness is gained and maintained by sheer experience; this is also thought to explain the big differences in safety awareness between different members of the personnel. Safety awareness is high, but has different meanings for different individual workers (varying from personal qualities to safety assurance based on checklist discipline). Protocols are made for normal treatment and care activities and thus contribute to the assurance of patient and worker safety; however, protocols do not exist for management of deviating processes⁴². 'Lack of safety' is not seen as a feature of process deviations: it just becomes apparent. 'Complications' are not seen as blameworthy, but are thought to be possibly avoidable in some cases.

6.6.1.2 Incidents

Incidents mentioned by interviewees are summarised in the text boxes below, of which the second one lists the most serious incident that they have witnessed. The baseline assessment exercise turned out to be a shortened RSO study. No single incident was mentioned by two persons. Personnel errors made in task execution are seen as the main incident cause. The compulsory form to report patient-related incidents is hardly used by members of the medical staff: it is too time-consuming and its usefulness is doubted; conversely, the nursing staff uses this form regularly, mainly to report fall incidents and medication errors.

Incidents are commonly seen as isolated events without underlying systemic causes. Suggested opportunities for prevention of incidents mentioned during the interviews were introduction or improvement of protocols, improved verification checking of critical tasks, and improvement of interpersonal communications at operational levels.

⁴² Although not clearly distinguished from other parts, the protocol for the thorax surgery wards [Sivro 97] contains a series of medico-technical crisis management scenarios. This protocol became available shortly after the baseline assessment interviews were held.

Incident examples that were mentioned during the interviews

- Lung puncture by the steel thread used to fix the breastbone (OR).
- Incorrect mixture of fluids applied for preparation of a heart-lung machine (OR).
- Bleeding due to the removal of pacemaker leads filled up the pericardium eventually leading to cardiac arrest; the removal was known to be risky (ward).
- A surgeon applied inappropriate tissue glue so that the tissue shrunk (OR).
- A nurse labelled some medicine with an incorrect colour code; later on, the anaesthetist administered the wrong drug selected by colour code (OR).
- Patient received blood of a wrong type and got, as a result, seriously sick (OR).
- During a replacement of infusion bags, the new infuse was not switch on or with an incorrect flow (OR, wards).
- A nurse administered wrong medication to a patient (ward).
- An opposing second opinion by a specialist from another discipline about a haematoma suspect was followed, but the initial diagnosis turned out to be correct one day later; a CT scan would have eliminated any doubt (ward).
- A patient who was restless due to some physiological problems could not be secured to bed because appropriate means could not be borrowed from another ward (ward).
- A restless patient fell out of bed (ward).
- A patient rang for over half an hour before a nurse responded to the call due to the configuration of the buzz system (ward).
- Frequently, a patient file is missing; it costs much time to retrieve the dossier (ward).
- A patient passed out in the bathroom before calling for acute help (ward).
- Medication for a specific patient arrived one day late (CCU).
- An anaesthetist submitted an unstable patient to the ICU without taking time for necessary briefing of the ICU staff (ICU).
- Understaffed, due to personnel illness, at night with restless patients, there was insufficient time to perform routine checks and problem anticipation (ICU).
- Patient monitors were adversely influenced when the network to which the monitors were connected was down; this phenomenon was noted when a nurse saw unrealistic values displayed on a bedside monitor (ICU).
- A patient monitor was still in the setting for the previous patient who was disconnected less than fifteen minutes before; such monitors reset automatically to default values fifteen minutes after the last use (MCU).

6.6.1.3 Accountability

There is a strongly developed habit among physicians not to criticise each other: you may be the next to be blamed. This attitude is seen as professional loyalty that should not be undermined. Surgical staff bears final responsibility for the cardiac surgery patients. In spite of this, staff members claim that sometimes residents act on their own account in a way that falls outside the final responsibility of the staff member. The accountability of nursing staff is anchored in legislation, i.e. the Professionals-In-Healthcare Act. Physicians function under the jurisdiction of the medical disciplinary tribunal. Individual commitment includes a strong tendency to take on responsibility for patients irrespective of their formal position within the cardiac surgery unit. Accountability and liability are usually only an explicit issue after complaints have been submitted. Operational responsibilities are not well described.

Incidents that were mentioned as 'worst case'-examples

- A secretary checked out an unattended snoring patient who was suffering from a cardiac arrest; the patient survived (ward).
- Physician forgot to adjust the settings of an insulin pump: fatal (ward).
- A defective backup defibrillator... a third one was available in time (OR).
- Wrong medication handed over to the surgeon and administered (OR).
- Incorrect assembly of a disposable tube set for a heart-lung machine was not noticed before application, air was pumped into the patient who survived (OR).
- An incorrectly labelled syringe containing a vasoactive instead of fentanyl was administered by the anaesthetist; the patient survived after strenuous crisis management (OR).
- An external pacemaker was replaced too slowly for a patient under telemetric surveillance outside the unit, who could not do without it; the patient survived, but passed out several times (CCU).
- A cardiac surgery patient in CCU care awaiting a bypass operation complained about pain in the breast, while the ECG showed symptoms of a heart infarct. After four hours the physician-in-charge arrived too late (CCU).
- A morphine dose was modified when the syringe was swapped and the nurse entered new settings (ICU).
- A patient who was submitted in a hypothermic condition while disconnected from mechanical ventilator equipment, failed to breath on his own and suffered a cardiac arrest; after reanimation the patient passed into a coma and died after several months (ICU).
- Without being informed, a nurse took over a patient for which all alarms were switched off (ICU).
- Due to a computer failure, essential patient data were unavailable for a long period of time (ICU).

6.6.1.4 Openness

Medical staff members are sceptical about and do not practise the compulsory notification of patient-related incidents to a review committee functioning at hospital level. This committee's image is much like that of a tribunal. Perfusionists discuss among themselves all heart-lung machine related incidents. Nursing staff submit (many) more compulsory patient-related incident forms than the medical staff, often after a team discussion about the incident and agreement about reporting to the review committee. Residents do not participate in complication review meetings of the medical staff. Individual learning by experience, i.e. the hard way from personal failures, is widely considered to be the most effective.

6.6.1.5 Communication

Staff members communicate about incidents and complications mainly through the meetings mentioned in Table 6-1. Outside those they talk to a limited number of direct colleagues. Nursing staff do not discuss incidents in meetings, with the exception of incidents that require notification to the compulsory patient-related incident review committee; however, all incident data should be documented in the patient's nursing file. Outside scheduled meetings, there is much communication about ongoing operational issues within the cardiac surgery cluster as well as with other disciplines acting as internal service providers.

6.6.1.6 Conclusions from the Baseline Assessment

The safety and prevention of iatrogenic harm to patients is a high priority issue for members of the cardiac surgery cluster. In their processes of providing care people hardly distinguish between incidents, complications or normally expected operational situations. Incidents are perceived as isolated events without underlying systemic causes and usually due to some personal error. The notion of organisational causal factors is absent or very weak. Individual learning from such incidents is limited to some adjustment of the prevailing protocol and temporary increased vigilance during task execution. With the possible exception of the nursing staff, incidents without adverse consequences for the patients are not reported. Retrospective contemplation of incidents experienced is completely absent. Communication about incidents is limited to the issues and the people that are of immediate relevance to the daily work processes. The operational surprises indicated in section 6.6.1.2 above show that incidents of all sorts do occur (hypothesis 1) and can neither be attributed to the last actor being the one to be blamed nor written off simply as a 'complication'. A closer look at most of the examples indicates – in some cases implicitly – possibly relevant causal factors that are controllable by the cardiac surgery unit or by other identifiable entities within the hospital organisation, at earlier stages in the process (hypothesis 5).

The baseline assessment provided useful compensation for the results lacking from the RSO study about incident perception and attitude towards reporting and learning. The interviews reveal that individual members of the cardiac surgery cluster find it difficult to think of recent abnormal situations that are perceived as incidents, but that they can all do so when prompted by an interviewer. The incidents recalled are then far from trivial, but people tend to blame themselves or other individuals rather than to diagnose such situations systemically. Individuals learn individually from incidents by experience. The existence of systemic causal factors underlying incidents is not yet recognised. These findings explain at least partly the RSO study results. These conclusions also indicate that the potential value of the PIERCE project is great, but that problems with implementation must be expected. The principles on which it is based contradict quite sharply the prevailing culture at important points. These will be the subject of extra effort, and further discussion later in this chapter.

6.6.2 Safety Culture Progress Measurement

After the safety culture baseline assessment discussed in section 6.6.1, the subsequent measurement of progress in safety culture was to be done about nine months after the introduction of the surprise notification system SINS-1, i.e. in the spring of 1999. At that time however, the implementation of SINS-1 was still incomplete, i.e. there was only one workstation in the OR without a tight coupling to the surgery report generator, which was unreliable due to networking modifications, and there was no database system manager. Thus, the SINS-1 system was hardly used, and not at all by staff of the wards. Improvement of the situation regarding SINS-1 by reactivation of the Review Team was attempted twice, but did not work out. Consequently, there was not much to measure in 2000. Many people who were directly involved in the PIERCE project and in the introduction of SINS-1 had changed job. A couple of interviews taken in spring 2000

confirmed the current invisibility of SINS-1 and the Review Team. Hence, it was decided not to proceed with the second safety culture measurement.

6.7 SINS Implementation

Against this background, the implementation of a Systemic Incident Notification System (SINS) for the cardiac surgery cluster started with the design and development of a low-threshold registration system SINS-1 to record observed operational patient-related surprises. It was designed to be complementary to the existing 'complication' register. By linking work process factors to physiological crises of patients, opportunities for prevention or improved control of risks to patients would be possible. The first approach was to define crisis management scenarios for cardiac surgery patients, similar to the scenarios that were identified for anaesthesia [Gaba 94]. It was thought that these could be used to build up a SINS-2 case base. However, it became clear that such crisis management scenarios are medico-technical in nature, aimed only at regaining control over the patient's physiological processes, regardless of operational circumstances. SINS-1 notifications on the other hand are about operational surprises, i.e. unanticipated changes in work processes. Thus, comparison of SINS-1 notification data with SINS-2 crisis management scenarios would be far from straightforward. One is cause, the other response. It makes more sense to design SINS-2 as an organisational memory of operational surprise experiences, wherein crisis management scenarios might be accommodated as validated solutions to specific operational problems. Integration with the complication register was considered, but rejected since this register is just a particular log of a patient's physiological condition, lacking any linkage with the relevant operational contexts. Furthermore, the register is updated centrally by the cluster's administrative staff, while the prime users of SINS were to be the operational clinical workers. In cardiac surgery work processes, patient-related surprises are handled as priority before anything else. That is to say, if a patient becomes unstable during care and treatment, priority is given to diagnosis of the immediate (physiological) condition and then to regain stability. Therefore, filling or consulting of an organisational memory system (SINS-2) by means of a surprise notification (via SINS-1) will not take place while there is a crisis to manage. In other words, the organisational memory (SINS-2) will at best be consulting afterwards when (relative) operational stability has been regained.

In the PIERCE project, the development of the SINS-2 and SINS-3 subsystems was postponed, because the development and implementation of the other main components, SINS-1 and the Review Team, was a priority. To this date SINS-2 has *not* been developed for PIERCE. In the SINS-1 - SINS-2 two-way interaction, the development of a taxonomy facilitating the structuring the description of validated case in the case base and speeding up the SINS-1-2 interaction requires a major investment which has not been forthcoming. Ichem (UK) has developed a simpler example of such a taxonomy for 'The Accident Database' and this took five years of work involving many domain experts [Powell-Price 98: private communication].

The development of the SINS-2 component is resource demanding and, equally important, would require availability of domain expertise for building the case base. Furthermore, as the main function of SINS-2 is to be an accessible organisational memory, the processes

of organisational learning needed to be started-up to begin with in order to generate lessons experienced in the past and worthwhile to store into the case base memory.

The SINS-3 module was available from previous projects as a prototype embedding expert knowledge in a way that can be employed by non-expert users, see sections 2.2.4, 2.3.1-2.3.2. This prototype rule-based expert system needed redesign so that interaction between SINS-1, and SINS-2, with SINS-3 would become feasible. The nature of SINS-3, an in-depth analysis of organisational causal factors usually as part of a double-loop learning process, justified substitution of a SINS-3 module by a qualified investigator as a first step. This is what was done for PIERCE.

A Review Team was set up – as a learning agency – to review surprises notified and to store validated cases in the SINS-2 case base. A resident was assigned to co-ordinate Review Team activities and to prepare case reviews: to become the learning agent or investigator for the agency.

In the following sections the development and implementation of the SINS-1 module and Review Team are described and discussed in greater detail.

6.7.1 SINS-1 - Low Threshold Notification

Eventually, any person involved with the cardiac surgery cluster should be able to submit a surprise notification message using a general-purpose workstation. However, residents are the only clinical workers who are actively engaged in practically all throughput phases and, therefore, are familiar with the work process phases and their contexts in the different units of the cluster. Hence, residents were selected as the prime user group for the trial of the SINS-1 notification system. A one-stop approach for reporting operational problems of all sorts was designed to make it easier to stimulate the use of SINS-1. This is very attractive when, otherwise, one observed event or condition would require several notification reports to different people or bodies. Electronic forms allow entry of data items only once, where they are common to different notification reports.

The specifications for the electronic notification forms in SINS-1 were derived from the following considerations. The system should be maintainable within the span of control of the cardiac surgery unit, i.e. be built with off-the-shelf software tools and supervised by the unit's system manager. Entry of data about operational surprises should be easy, i.e. the form ought to be kept simple and the workstation must be nearby. It was imperative that the recorded data could be compared with similar experiences in the past, i.e. in the SINS-2 module. The one-stop policy meant that the compulsory form for notification of a patient-related incident, the so-called FOBO-form, must be integrated. The SINS-1 system had also to be able to identify a potentially serious threat during entry, independent from case comparison in SINS-2 or a Review Team assessment. In such a serious case it was felt that it might be useful for event review if more detailed incident data could be recorded immediately in addition to standard items. In order to monitor operational surprises, the cardiac surgery cluster manager needed to be informed – as a standard operating procedure – about reported surprises by the day following on the notification. An in-built report generator was felt to be very useful for this purpose. Reports generated for other bodies could be sent out (automatically) via Intranet. User access had to be controlled, e.g. by means of a log-on procedure.

6.7.1.1 SINS-1 Functional Specification

The considerations above and operational constraints within the cardiac surgery cluster resulted in the following set of functional requirements for the SINS-1 notification system. Data entry should be done through workstations networked with a central SINS-database server. This was chosen because notification by means of paper forms would require an administrative stage of data transfer into a database without adding value to the notification report. The turnaround time for filling in the form and the administrative post-processing would be significantly more than 24 hours and thus delay feedback to the work floor. Workstations were planned in the OR area, in the cardiac surgery cluster wards, in the ICU, in the residents' room, and for senior staff members.

Smart screen forms ensured that only those forms and items that were relevant were displayed and avoided duplicate data entry. Details and common categories of problems were predefined when possible (roll menus, lists, radio buttons, check boxes, etc.) and functional, e.g. to mark the operational phase as relevant context and indicate the critical activity at the time of the detection of the surprise. The requirement that the SINS-1 system has to be able to identify a potentially serious threat during notification data entry was met by implementation of a risk graph algorithm [IEC 98, DIN 94] tuned to the processes in the hospital. This tuning can be done by peer review of notifications that demonstrate a significant discrepancy between the severity assessment by the reporting person and the result of the risk graph evaluation.

Any kind of operational surprises – whether or not of mixed nature – have to be accommodated in the notification forms. Therefore, the standard notification form consists of a set of free text items to address besides selective items such as severity assessment, corrective intervention and its effectiveness, situation, surprise and circumstantial factors. The unit's members were requested to notify any operational surprise without bothering about the category of the notification, e.g. an observed occupational risk, a patient-related incident or the malfunctioning of a medical device. Instead, the notifying person indicates that the notification involves one or more of the non-exclusive generic object categories: oneself, team members, patients, devices, installations or protocols.

As a consequence of the one-stop policy the initial idea was to integrate the FOBO-forms (compulsory patient-related incident notification) as well as the complication register. Work process surprises like occupational safety, health and environment (SHE) problems, technical problems of devices, installations or other systems were picked up during the initial review of the first design of the screen forms.

Priority was given to the coupling of SINS-1 with the computer-based surgery report generator program in the operating rooms block, which had to be re-engineered for this purpose. In the OR application, tight coupling of SINS-1 notification with the delivery of the (electronic) surgery report resulted in 100% coverage of the surgical procedures, because these reports are made during and shortly after the finishing of each cardiac surgery procedure. The concise OR report forms were developed primarily to speed up notification if there was no operational surprise from a surgical procedure to report. This 'short' form contains a list of surprises that occur more or less frequently during cardiac surgery procedures. The reporting person only has to click one or more of these 'incident scenarios' instead of writing free text lines. Although the number of surgical procedures

within a specified time span is also known without the coupling of the two software systems, it was expected that coupling would lower the threshold to report an operational anomaly.

It was felt important that the reporting person should have access to his own notification records to correct or update data. The reporting person may wish to update or complete notification data on second thoughts, without being allowed to modify submission from other persons. Access control is therefore needed.

Screen forms should be self-explaining whenever possible to avoid the need for on-line support. Context-specific help pages that were crucial for proper judgements of knowledge model items in the ARA project, see 2.3.1, were hardly read by (skilled) users, which adversely affected the quality of the diagnostic outcomes. In a notification system, the reliance on help pages for adequate data submission would either raise the threshold for people to report an operational surprise or reduce the quality of data submitted. These were therefore avoided where possible and the screen designs used were presented to residents, senior staff members and nursing staff for comment and to test their clarity.

Notified data should be stored in a (raw) database. There is no reason to eliminate the raw data when a matching case is retrieved from a SINS-2 case base. The raw data can still be used to produce reports for external purposes and to test the correctness of SINS-2 case matching by peer review.

New notification submissions should be printed automatically overnight. The idea here is that the cluster manager is informed about notified operational surprises not later than by the beginning of the next day.

A (SINS) application manager was needed to initiate fast corrective action if problems should arise with usage of the SINS-1 system for whatever reason, and to maintain the SINS application, including the database. The SINS-1 notification system must be on-line, if people are to be asked to submit any operational surprise experienced: flaws in availability would have a significant adverse impact on people's willingness to report.



Figure 6-4: SINS-1 start-up screen

6.7.1.2 Realisation

The screen forms were designed for regular surprise notification, concise surprise notification linked to the surgery report generator in the OR unit and, in addition, for

generation of FOBO reports. The routing for entry of notification data is depicted in Figure 6-5. After a start-up procedure, the user must decide which form applies.

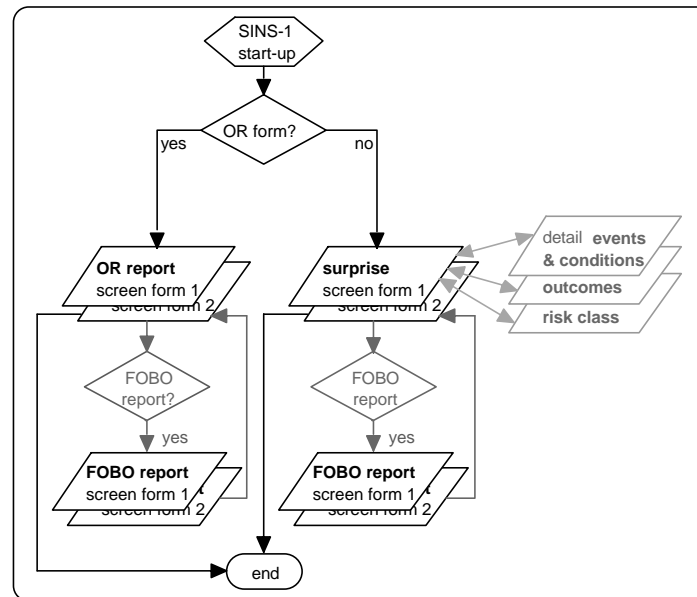


Figure 6-5: screen form navigation in SINS-1

Three modules of screen forms are implemented: the concise OR forms that are used in conjunction with the surgery report generator, the surprise forms that allow data submission in greater detail, and the FOBO forms that apply for patient-related incidents. The OR report forms can be completed in just three mouse clicks if no surprise is reported from a surgical procedure. The grey lines represent optional routes.

The SINS-1 system was implemented in a network configuration with four workstations and a database server and launched in April 1998. The initial group of users consisted of the residents of the unit, some members of the senior staff and a head nurse. The residents usually work in all phases of the cardiac surgery process and can also easily be approached by the learning agent. Eventually, it was planned that other professional groups would also be challenged to submit directly their operational surprise experiences.

The learning agent introduced the SINS-1 notification system in April 1998 to his fellow residents and subsequently stimulated them to use the SINS-1 in the OR area, as well as in the other entities within the unit, including the ICU in so far as it concerned the cardiac surgery patients. Six months after introduction of SINS-1, the Review Team was installed and operational surprises, reported in SINS-1, could be reviewed for drawing lessons and generation of recommendations to adapt prevailing theories in use theories-in-use, see section 6.7.2.

In eighteen months 168 notifications were submitted to SINS-1 of which just 23 in 1999, see Table 6-5 below. These figures reveal firstly that the integration with the outdated surgery report generator had not yet succeeded, because this would have led to 900 notifications. This DOS application required a major redesign effort in order to realise the

tight connection with the Windows application SINS-1. In addition, the learning agent left the organisation for a job elsewhere. His replacement did not come until half a year later, and by that time the discipline to notify operational surprises had faded out. Without SINS-2, the SINS-1 subsystem could only be used for data collection. The reports in the SINS-1 database show that several reported incidents did deserve to be reviewed with priority by the review team. A daily notification summary report to the cluster manager might have identified these cases as their notification was submitted, but this aspect was not implemented systemically.

Table 6-5: *SINS-1 notifications per form category*

	OR report	Notification	FOBO
1998: 8.5 months	79	54	10
1999	1	22	0

6.7.1.3 Assessment of, and Conclusions concerning SINS-1

Per April 2000, the unit assigned a 'quality manager' for 1 day/week who has taken on the role of leaning agent. The preceding year, 1999, had shown that voluntary agents are too easily absorbed by other duties. The lack of continuity in the functioning of the learning agent that became apparent with the departure of the first one contributed to a large extent to the poor notification practice in the second year. It also stalled the functioning of the Review Team so that feedback about lessons learned from case review could not be communicated to the members of the cardiac surgery unit. Such a feedback would almost certainly have served as incentive to submit notifications of operational surprises experienced.

The presumption concerning the need for a fast notification procedure using the OR report forms, where the standard situation would be to use the more extended set of notification forms, may have to be revised. The use of the concise set is not restricted to the OR stage and makes more sense if repetitive, but relatively minor operational surprises need to be reported. Furthermore, *the optional form for details of events and conditions*, see Figure 6-5, turned out to be *superfluous*: the efficacy of a review preparation by a skilled learning agent outweighed by far the benefits of challenging a reporting person to do a sort of event reconstruction using the details form. This form has not been used since the SINS-1 system came into service. In addition, the concise form was also used outside the OR area. In these cases, the questions for severity assessment and the adverse outcome specification from the standard form were omitted, because they were left out of the concise form. Thus, the use of the concise form instead of the standard one implied the loss of an opportunity to establish an objectified operational risk classification standard.

The tight link between SINS-1 and surgery report generator in the OR area would have led to a 100% rate of surgical procedures monitored. Albeit a 100% coverage of processes monitored seems ideal, because exposure measures can then be defined and data validated to it, it is not necessary for surgical procedures, as the precise number of operations is known at the end of the day. So, the SINS-1 system did not have to 'wait' for the OR-report generator as long as surprises experienced were notified.

Rather than to have two similar but still different sets of notification forms within SINS-1, i.e. the standard form plus the concise OR set, it is justified to reduce these to just one set, based on the concise OR form. The standard form has too many features, which are partly superfluous, whereas the concise form suits its basic notification signalling function adequately. The additional features from the standard form for risk classification and loss specification can be transferred to the concise form. Thus, the standard form can be made obsolete. If we take the forms for risk class and adverse outcomes and transfer them to the concise report form set, the standard form set can be removed completely. This would simplify the notification even more as is depicted in Figure 6-6. The FOBO form features remain intact given the principle that one should not submit the same data twice for different notification purposes.

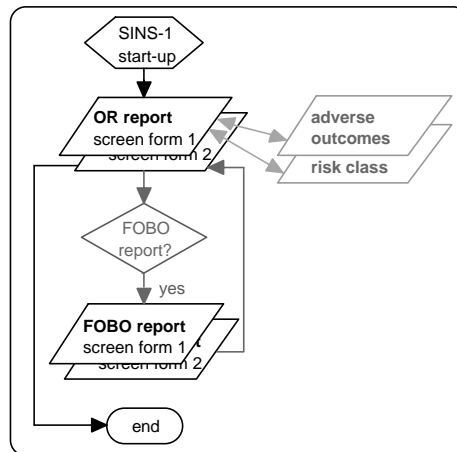


Figure 6-6: proposed configuration of screen forms in SINS-1 based on the PIERCE experience

Although the list of scenarios in the concise form is derived from the surgery phase in Figure 6-2, this form may also be used in any of the other throughput phases. The predefined scenarios can be adapted for other throughput stages and can be expanded temporarily to trace a specific type of surprise for a specific time span if desired. This approach of 'notification topic campaigns' has been demonstrated to be successful in relation to the complication registration system. For a voluntary notification system SINS-1, similar attention awareness tactics might be effective as well.

6.7.2 Review Team

A Review Team was formed to become the cardiac surgery unit's learning agency. The core members, a senior staff member, a head nurse and an intensivist, had a profound combined knowledge of the cardiac surgery processes in most of the stages distinguished in the throughput model of Figure 6-2. The Review Team was co-ordinated by the learning agent who – during the start-up – happened to be a resident and, hence, had a wide overview of the work processes and the transition stages. When needed, it was possible to invite a domain expert from another discipline, like anaesthesiology or cardiology, to participate in the review of a specific operational surprise. I, as experienced incident investigator, supported the domain experts of the Review Team.

The main mission of the Review Team (RT) was to review and learn from operational surprises in such a way that the cardiac surgery unit could improve its performance with respect to control of operational risks. Therefore, the Review Team had to assess new types of operational surprises as notified via the SINS-1 system and transform these into valid cases for the SINS-2 case base. Furthermore, the RT might perform or commission a SINS-3 type of in-depth analysis when needed. And finally, the RT were to work out example cases of learning from operational surprises to illustrate the merits of surprise notification to all members of the cardiac surgery unit. Sources for cases to review were the SINS-1 notification system, the daily morning sessions, see Table 6-1, and observations reported to the learning agent. Proper understanding of surprises reported and productive evaluation of options for improvement of operational risk control requires appreciation of the applicable operational context and the prevailing theories in use. By its composition, the Review Team was well suited for this task.

6.7.2.1 Case Review Process

A protocol 'Review of Abnormal Situations', see Appendix J, was developed to describe and structure the work by the Review Team. The 'throughput' model of Figure 6-2 was adopted as a context reference framework for case review. The learning agent was to prepare the sessions of the Review Team, especially concerning fact finding and dossier tracing. A new case was first classified either as (new) problem to solve, or as a problem known to exist, but for which a feasible solution is outside the direct span of control of the cardiac surgery unit. For the latter class of problems, a decision had to be made to initiate a double-loop learning process. A case presented for review was initially classified as a problem to analyse (PA) or a case for which a decision about implementation of a known solution (DA) is to be made at some time. In accordance with the classification, the Review Team could invoke procedural steps as described in Kepner [81].

Typically, the Review Team has to set a trigger level value when a case is classified as a decision analysis (DA) case, meaning that for this specific type the incidence must reach the trigger level before further action is initiated: see also section 5.2.6.

A problem analysis (PA) case was first reconstructed applying the Events and Causal Factor charting analysis technique [Buys 95], because it combines a systemic, time-sequenced ordering of facts on events and enabling conditions, i.e. underlying causal factors, in a graphical representation, see Figure 6-7. During a RT session, discrepancies in event perception among RT members were communicated objectively with the chart as the common reference. The method is scalable to the magnitude of the case and is easy to learn. E&CF analysis can be used to test the adequacy of the sequence of events, be scoped to delve into systemic, organisational (root) causes, or extended into verifying the adequacy of supporting evidence. This E&CF reconstruction of a PA-case has to be prepared by the learning agent, who for this purpose – as an investigator – might have to trace back events and conditions relevant to an explanation of the occurrence of the operational surprise experienced. This might involve going back to the notifier or others in the unit for additional information. Patient-related surprises more often than not required that patient records and data from the hospital information system be retrieved to identify case specific facts and conditions that contributed to the surprise. These files could reveal parts of the tacit prevailing theory-in-use and could also indicate flaws in the equivalent

espoused theory. The reconstruction of the surprise event was restricted to the facts describing relevant deviations from normally expected acts and conditions, as well as to intervention decisions in response. The resulting chart formed the basis for the problem analysis during a session by the Review Team.

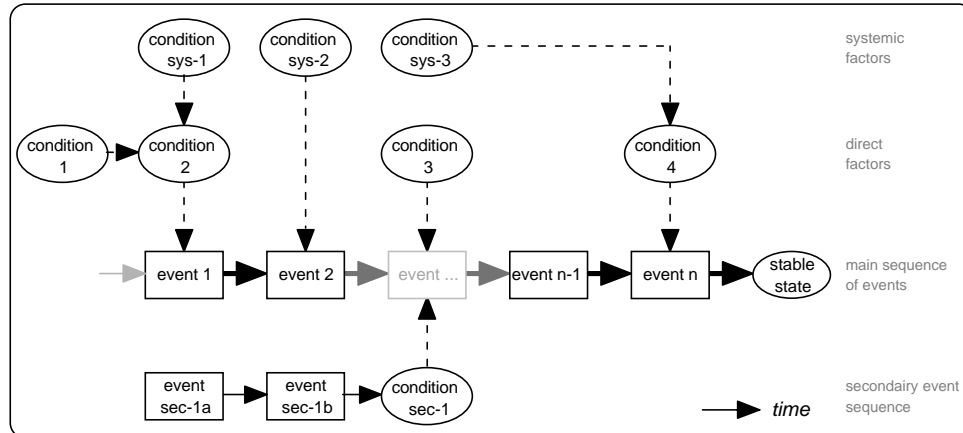


Figure 6-7: general format for Events & Causal Factor charts based on [Buys 95]

A case reconstruction can be called satisfactory when the primary sequence of events is consistent and fits with secondary events and other conditional factors identified. In the PIERCE project, case reconstruction concerned operational states and conditions that were counterproductive regarding the goals of the unit or were not anticipated in the prevailing theory-in-use. The starting point for E&CF reconstruction was the reported surprise. Then E&CF reconstruction proceeded forwards in time until the operational deviation was brought back under control, and backward in time until the surprise-initiating event. The operational surprise description focused on the sequence(s) of events that deviated from the anticipated process and the specific operational conditions that enabled the 'deviating' events to occur. For instance, the rupture of a surgical glove during an operation is counterproductive, although anticipated, whereas a cardiac arrest in the post-op ward might well be unanticipated for a specific patient. A post-op physiological state of hypothermia (surprise) on delivery of a patient to the ICU may go back to a standard operating warming-up procedure in the OR.

6.7.2.2 Learning by Case Review

The problem analysis focused on the enabling conditions that were linked to events in the primary sequence. An enabling condition is a contributing factor that has been identified during the case reconstruction inquiry and has been linked with a particular event from the primary event sequence. Such a contributing factor itself may result from a secondary event sequence. Solutions to the problem that was contained in the operational surprise break the chain of primary events, preferably before harm is done. In principle, each enabling condition identified indicates an option to break the chain of primary events. However, there is no need to try to eliminate all of these enabling conditions in order to assure prevention of reoccurrence. Hence, the Review Team had first to assess where the primary chain of events could be broken and which of these possibilities offered the best

option. Then it could assess which of the enabling conditions preceding the 'sequence breaking' event provided the most effective solution. If this solution concerned a systemic factor underlying an enabling condition, it might be hard to implement the solution in the short term, as it probably required a double-loop learning process. However, other enabling conditions might offer partial solutions that were easier to implement and, when combined, provided adequate control of the problem analysed. The question must always be whether a stop gap partial solution will then stand in the way of a more permanent solution, since it removes some of the urgency. Staff may stop pressing for more and the adequate becomes the enemy of the good. The weighing of the pros and cons of more or less systemic solutions to an operational problem identified draw upon profound operational domain expertise of the RT members. The outcome of this analysis process was a set of feasible measures to implement within the unit's span of operational control, and requiring modification of applicable theories of action in the throughput phase where the operational surprise was observed or in a preceding stage. Appendix M contains a worked-out review of case 6 from Table 6-6 below. The Review Team could analyse on average two cases per session of 1 - 1.5 hours.

Table 6-6: *types of operational surprises reviewed by the Review Team*

No.	Description	Type	Detection phase	Lessons	Implementation phase
1	bleeding post-op aneurism	PA	post-op care	active aneurism policy as s.o.p.	post-op care
2	pre-op stress	PA	pre-op care	adjust heparin policy	pre-op care
3*	patient died after 100 days	PA	MCU	<i>many + in all phases</i>	all
4	post-op hypothermia	PA	ICU	revise hypothermia policy	OR, ICU
5	use of β -blocker	DA	MCU	ensure that current protocol is known	all
6	skin needle visible on X-ray picture	PA	post-op care	Include skin needles in counting protocol of OR-team; retrain residents and staff in picture reading	OR; general
7	trouble to have X-ray picture made in the ward	DA	MCU	communicate with ambulant X-ray service provider	MCU
8	pre-op aneurism	PA	Admission	active aneurism policy as s.o.p.	pre-op care
9	10 days delay before surgery was performed	PA	pre-op care, OR	assure operational readiness when other disciplines must join	rostering (staff)
10	beds parked in ICU-corridor (access from OR)	DA	ICU	beds removed to traffic free section	ICU
11	surgical glove rupture	DA	OR	when often, claim better quality gloves	IncT = 100
12	internal bleeding after removal of a femoral catheter	PA	ICU	no measure: extraordinary set of physiological conditions and operational circumstances	-

In its active initial period of existence, the Review Team selected twelve cases for review, three of which came out of the SINS-1 notification system. Eleven of these were actually reviewed in five sessions, one large case (marked with *) was ready for review. Table 6-6 lists the sorts of surprise reviewed. Preparations included the drafting of an E&CF representation of the operational surprise. Drawing packages were applied for the drafting of E&CF chart. At the end of the day it was concluded that drafting this way is too time-

consuming even for an experienced learning agent. Instead, Post-It notes were used to set up the E&CF chart on a wall. The wall display aids the discussion in the Review Team.

Other sources for new cases like the morning sessions, daily SINS-1 summaries of new notifications on the desk of the unit manager or corridor chat, remained under-utilised. The daily summary report generator was not yet implemented in SINS-1, while the first and the last source put extra demands on the presence of the learning agent in order to translate surprise signals into surprise descriptions. For instance, in the daily morning sessions, residents from the night shift inform staff members of the morning shift about the status of critical patients and discuss the management of crises that occurred overnight. Whether or not such a crisis was unanticipated or the response handling inadequate can only be sensed during the sessions by an observer, e.g. the learning agent. But normally he was present as a resident and had his hands too full to do this. When he left, this sort of opportunity was impossible to capitalise on without a learning agent.

Results from surprise reviews were presented informally by RT members to the accountable manager whenever possible, and to the location manager of the cardiac surgery unit. This informal approach proved to be very effective in some cases, e.g. in case 6 of Table 6-6 relating to skin needles where the systemic solution option fell outside the unit's formal span of control: the solution was still implemented within 24 hours. The recommendations and preference choices made by the Review Team were appreciated and accepted by those involved in the cardiac surgery processes for their applicability in current practices. The wider dissemination of review outcomes was not seen as a Review Team task: this was felt to be a responsibility of the location manager. However, it was concluded that appropriate communication channels needed to be developed, for instance, as a "Review of the Month" Bulletin for distribution within the unit.

6.7.2.3 Review Team Continuity

The departure of the first learning agent demonstrated the relevance of the role of co-ordinator of the Review Team. The team was practically stalled immediately, because cases for review had to be prepared and usually required inquiries. The Review Team did not have the time to do this. An attempt to continue the RT work after a gap of six months stranded again, because the new learning agent was overloaded with various time-conflicting tasks. Despite twelve months of inactivity, however, because of the positive experiences, the members of the Review Team were still willing to continue the RT work. This willingness of Review Team members to act as organisational learning agency was remarkable. It was based on the recognition of the value of the objective scrutiny of cases using the E&CF analysis method in conjunction with the throughput model of the unit's operations. Also, the case reviews provided the members convincing insights into the relevance of operational constraints to, and opportunities for improved control over operational risks.

Of the four tasks of the Review Team stated above, i.e. case review, preparation of validated cases, in-depth investigation and generation of example cases for wider dissemination, only the case review and, to a lesser extent, the example case generation were actively taken up in the first meetings. The production of validated cases for the SINS-2 case base not only required the existence of the SINS-2 Case-Based Reasoning

system, but also a template for case description. This case base template is one of the issues for follow-up work. No single operational surprise occurrence was identified during the project that would have demanded a SINS-3 type of in-depth investigation. The implementation of the SINS-1 notification system also had not been going long enough to reach the incidence trigger level set for the surgical gloves, for which a SINS-3 type of investigation could have been justified. Consequently, the Review Team has not yet been involved with an in-depth analysis of a specific (type of) operational surprise to identify underlying organisational, i.e. systemic causal factors underlying the unit's risk control problems.

6.7.2.4 Review Team reviewed: Discussion

Organised learning at the level of the cardiac surgery unit requires a learning agency that is well informed about the daily work processes and process constraints. The evaluation and selection of options for concrete measures and their implementation would have been less effective if the RT team members had come from elsewhere, outside the unit. Single loop learning characterised the Review Team's main mode of functioning in the (few) operational surprises reviewed.

An unexpected result was the key role of the RT co-ordinator as inquirer on behalf of the Review Team. The search for facts that either are or should be traceable in the patient's files may be time-consuming, but is crucial for a proper case reconstruction in E&CF format. It would be a waste of expensive resources to leave this preparatory work to RT members. The assignment of the RT co-ordinator role to a member of the unit also expresses the willingness of the organisational unit to organise its learning processes. In this way, the RT co-ordinator can function as a primary learning agent, a sort of local safety co-ordinator.

The importance of training the learning agent is another conclusion. Because the E&CF analysis was the Review Team's preferred method, the learning agent needs to learn to capture facts collected from various sources that account for events and conditions and display them in the E&CF chart representation. For this purpose, a subset of the E&CF charting procedure was taken to design the Systemic Incident Description Technique (SIDT), see text box below. This should be the basis for training.

6.8 Discussion

The testing in the PIERCE project of the SINS-concept described in Chapter 5 was limited to the notification system SINS-1 and the Review Team as a learning agency. The problem of "what comes first, the chicken or the egg?" was inherent in the concept concerning the SINS-2 case base as organisational memory: the memory must be filled first before case retrieval is possible. The project did not get further than notification and Review Team review since the resources were not present to develop SINS-2. Hence, the stage was not reached where SINS-2 matching could further lighten SINS-1 notification.

Systemic Incident Description Technique SIDT

The purpose of SIDT is to capture evidence before ordering events and enabling conditions into an E&CF chart. Linking of events with each other and with conditions is omitted in SIDT, although some ordering in time is recommended.

Incidents can be reconstructed in terms of (time-sequenced) events together with the conditions linked to the events that enabled the event to occur the way it did. Methods like Events & Causal Factors Charting, Energy Trace & Barrier Analysis, Fault Tree Analysis or Tripod- β do this in different ways for different purposes. The depiction of an unanticipated deviation in a work process, i.e. an operational surprise, may be done by classifying relevant facts as an event or as a condition enabling an event. An event is an action by an actor, possibly upon an object, or the (corrective) response to it.

An event in an SIDT account of an operational surprise should not reflect an unrelated normal operational action. Other relevant, although secondary events may be identified that end up in an enabling condition of the main surprise event sequence. A condition 'stamps' the event it relates to and makes it possible for that specific event to occur as it did. The nice feature of conditions is that they are – at least partly – systemic, that is to say, embedded in the set-up or design of the work processes in which the surprises occur. Systemic conditions can be influenced proactively. See below for an example of a SIDT account of an accident.

Example of an SIDT description of an accident

A patient dies after an acute cardiac rhythm disturbance in the post-surgery ward. The patient was connected to an external pacemaker until a resident could prepare the bed-bound patient for transfer to the sending hospital.

Structured in terms of subsequent events and allowing conditions, this abnormal situation can be rewritten in the following way:

Condition 1: a resident was assigned to move a post-surgery bed-bound cardiac patient in order to prepare this patient for transfer to the sending hospital.

Condition 2: the doctor had prescribed an external pacemaker for correction of instabilities in cardiac rhythm

Condition 3: the resident was not informed about the fact that the patient was connected to the external pacemaker

Condition 4: the resident was unaware of the reasons why the unusual provision of the external pacemaker at the bedside was in place

Event 1: while moving the bed, the resident disconnected the patient from the external pacemaker => task performance error = allowing condition 5

Condition 5: the patient's heart function is no longer controlled: disturbances in cardiac rhythm can propagate freely

Condition 6: the resident was inexperienced in the cardiac surgery ward

Condition 7: the performance error of the resident was neither noticed nor corrected in time

Event 2: the patient suffered a serious cardiac rhythm irregularity and fainted

Condition 8: subsequently, the acute heart condition was not recognised quickly enough as such by nursing staff

Event 3: ...and the patient died.

Condition 9: the so-called FOBO protocol requires notification of incidents to patients: this incident was assessed as an avoidable, but fatal accident.

Describing an incident in this way helps to focus on opportunities for prevention of unacceptable development of emerging abnormal situations rather than to look for the scapegoat.

In the example above, the identified conditions 3-4 and 6-8 can be linked with generic causal factors like communication (of critical process information), personnel performance discrepancy, supervision, and vigilance regarding emerging abnormal situations. Assurance of any one of the underlying causal factors indicated above would most likely have prevented or blocked the fatal course of events.

The implementation of SINS-1 was slowed down by various causes, such as the design of the one-stop notification forms, the implementation of the server - clients network concept and the twinning of SINS-1 with the surgery report generator. The first generation users, i.e. residents, needed some training by the learning agent. Some staff members also learned to submit surprise notifications using SINS-1. The concise form proved to be sufficient for the realisation of the basic function of signalling the organisational learning agency, i.e. the Review team, that another operational surprise was detected, provided that, without the SINS-2 subsystem, the SINS-1 database was inspected regularly. The extended – standard – forms were less user-friendly, because free text was required and because of additional features for data elicitation. For instance, case review preparation by means of the detail form was just a superfluous element. However, it must be said that the key role of the learning agent who co-ordinates the Review Team, was not recognised before the Review Team was actually launched almost half a year after SINS-1 became operational. The other important function of SINS-1 was that of a process monitor for the unit manager. For this purpose, the daily reports were to land on his desk, but this did not happen: the report feature was not realised after the launch of SINS-1. Inspection of the SINS-1 database demonstrated that several notifications deserved to be prepared for RT-review, but this did not happen. Such an inspection can best be assigned to the learning agent who then informs the Review Team about the nature of the cases selected.

The *Review Team functioned as an organisational learning agency* that provided adequate compensation for the loss of context in the SINS-1 notifications: recommendations made were adopted easily, thus leading to the conclusion that the lessons were indeed learned in the few cases reviewed. This result indicates that the impact as learning agency on the cardiac surgery unit's theories-in-use most likely would have been less if the RT members had not been members of the unit as well. The vulnerability of the Review Team due to its dependence on the Review Team co-ordinator as a primary learning agent came out as a surprise. As a result, a 'quality manager' has been assigned to assure a structural presence of a learning agent.

The SINS-2 subsystem, although not implemented in PIERCE, is a crucial part of the full SINS package. It should function as an organisational memory of operational surprises experienced in the past and of responses to them in order to regain process stability. The value added by the Review Team to a SINS-2 case base concerns the root cause classification of enabling conditions and the assessment of adjustments to be made in the relevant theories-in-use. These conditions, classified using a scheme of causal factor categories as shown in Appendix N, should be part of the case stored in the case base, and provide generic clues for resolution of unwanted operational surprises in future. They can link on this basis with a SINS-3 expert system. The PIERCE project showed that real experts commissioned by the learning agency may well replace the SINS-3 subsystem. The incident trigger *IncT* in SINS-2 forms an objective mechanism – especially regarding apparently minor operational surprises – for managing the deployment of scarce investigative resources when a SINS-3 review is needed, i.e. double-loop learning mode is required. All this did not happen in the PIERCE study. It was set up, but not triggered in PIERCE, so no adequate test of its value could be made.

The Review Team's products as a learning agency were recommendations or lessons that were communicated to the (sub-) units involved, such as proposed changes in practice and protocols, and worked out examples for general awareness raising. In principle, they

should have produced validated input to a SINS-2 subsystem, but again this did not happen, as SINS-2 was not developed. These may have a format as illustrated by the generic cases in Appendix O. The propagation of the worked out examples among the cardiac surgery unit members should not rest with the Review Team, but with unit's management, possibly supported by the learning agent (RT co-ordinator).

A second series of interviews after three years revealed that the communication among professional care providers still seems to be on a strict 'need-to-know' basis. The members of the Review Team did not even informally discuss lessons learned with direct colleagues. In one case, a direct colleague of a RT-member was not aware that his mate was involved in the Review team. This 'laissez faire – laissez aller' principle indicates a strong *inhibitor* for organisational learning

6.9 Conclusions

A major finding from the PIERCE project is that the SINS-concept can work. Bits and pieces of the system for organisational learning from small-scale operational anomalies drawn from the whole history of projects up until then fell into place like a jig-saw puzzle. The SINS-concept [Koorneef 98] is described in detail in Chapter 5. The concept provides key functions for organisational learning from abnormal situations ranging from a large number of instances to less than a single accident occurrence, e.g. a near-miss. The implementation of the SINS-concept in the surgical unit generated many lessons on what was feasible and important and what components needed tuning. A major advantage of the SINS configuration is that initial notification of an operational anomaly is merely a signal for initiating the organisational learning system: "Ouch! An outcome of the work process is not as anticipated! Have a look." The project showed that the notification form may remain surprisingly simple, provided that other components of the learning system are in place. Not the least is the conclusion that the organisational unit must have a willingness to learn to start with and that organisational learning must be organised. It does not just happen.

Concerning the hypotheses postulated in 6.1 and 6.3, the following conclusions can be drawn from the PIERCE project. During the PIERCE project, the notifications submitted to SINS-1, the interviews from the baseline assessment described in 6.6 and cases directly reported to the learning agent confirm the first two hypotheses, namely that

- a) significant numbers of abnormal situations do occur within work processes of the cardiac surgery unit, and
- b) these operational surprises can be successfully recorded for review without causing unacceptable pressure of time or concern over recording.

The RT reviews demonstrated systemic factors underlying reported operational surprises and the E&CF reconstruction of a case also show to a greater or lesser extent underlying systemic causal factors. The RT recommendations focused on the factors that could be influenced by the unit's management and the actual implementation of recommendations, which is in accordance with Hypothesis 3 that states that systemic factors exist, some of which can be influenced by the unit.

The effectiveness of learning from operational surprises by the unit is demonstrated partially by the Review Team after case review. For instance, the skin needle case led, because of its clarity, within 24 hours to modification of the prevailing theory of action applicable to the OR nursing staff and is still remembered as a lesson learned. However, the efficiency needs to be improved especially during preparation of a case for review, and in future also by interaction during entry of notification data between SINS-1 and SINS-2. Hence, hypothesis 4 - stating that the unit can effectively learn from these abnormal situations and change its theory-in-use – has not been falsified yet. However, the second series of safety culture interviews indicates a strong inhibitor for organisational learning when even direct colleagues of RT members were not informed about Review Team activities. We might say that the receptiveness of members of the organisation is a key factor for successful implementation of organisational learning.

In practically every case analysed with E&CF Charting it was possible to project enabling conditions on the time-sequential throughput model of the cardiac surgery unit. For instance, the risk of misinterpretation of X-ray photos in post-op intensive care in the 'needle example', is more systemically covered by counting the skin needle before the thorax is closed in the OR phase. This is positive evidence for hypothesis 5 stating that causal factors underlying an operational surprise in phase P_j can often be influenced more easily in a preceding phase P_j ($j < i$).

The (safety culture) baseline assessment made clear that cultural blockages exist in the different professional groups that constitute the cardiac surgery processes. Systemic reflection on incidents among staff hardly exists and not beyond the own professional group. The second series of interviews show that these blockages still remain.

PIERCE showed in many ways how the SINS implementation could work, but also that there were still aspects, notably of culture and resources, which were not in place to make it a complete success. This shows yet again how hard it is to implement fully successful learning systems. The next chapter will address this.

In Chapter 7, the progress made up to now will be critiqued and evaluated by means of a Viable System Model (VSM) diagnosis. The VSM is a model explicitly developed to deal with the informational aspects of maintaining organisations able to deliver satisfactory performance in changing environments. Although this was found too late to influence the vast majority of the projects described in this and the earlier chapters, it seemed to offer a generic framework for understanding the successes and failures of those projects and generalising the results. The VSM can be deployed to diagnose the problems in the organisational learning processes in a viable organisational system.

The VSM View...

7.1 Introduction

The incident recording and analysis systems (IDA, ISA, SINS) all intended to help inform organisations and each was meant to become embedded in a particular organisation (Rotterdam harbour, individual and groups of companies, hospitals). The two main objectives in this chapter are, firstly, to critique tools and systems like ISA and SINS as applied in organisations, and, secondly, to generalise SINS to other organisations. During the course of the research described in the earlier chapters I came across, as part of a review of the systems literature, the work of Stafford Beer. Although this was found too late to influence the vast majority of the projects described in the earlier chapters, it seemed to offer a generic framework for understanding the successes and failures of those projects and generalising the results.

Both objectives of this chapter stated above can be served by the 'Viable System Model' (VSM), a model explicitly developed to deal with the informational aspects of maintaining organisations able to deliver satisfactory performance in changing environments. The VSM was developed by Beer [Beer 79, 81] as an application of cybernetics principles to the control of organisations. Cybernetics is claimed to be of particular relevance to problems of control:

"In defining cybernetics in the first instance, the great Norbert Wiener used the famous phrase, 'the science of communication and control in the animal and the machine'. In so doing, he was trying to emphasise two discoveries. *The first was that communication and control were virtually synonymous. Regulation is an informational process.*" [Beer, 1983, page 1 (*emphasis added*)]

In keeping with systems science generally the aim of cybernetics is to elucidate the general principles of regulation found in all systems. Beer continues:

"The second discovery was that the classical dichotomy, inherited from the ancient Greeks, between the animate and inanimate was delusory. *There are invariant laws of regulation that apply, like gravity, to everything.*" [Ibid. (*emphasis added*)]

The VSM is designed to apply these 'invariant principles' to any organisation (that is, all organisations that are viable⁴³ in principle). Therefore, the VSM offers a means to

⁴³ The Oxford English Dictionary defines the term '*viable*' as 'able to maintain existence'.

understand the design and implementation of SINS, etc., and a means of generalising this understanding to other viable systems. Furthermore, if other general approaches (such as that proposed by Argyris) can be reconciled with the VSM, we will also have a choice of two compatible approaches to generalising the SINS approach and the other findings of this research.

In the previous chapters, the search for effective learning from small-scale incidents has led to the concept of SINS within the framework of organisational learning. The implementation of a SINS in a hospital unit revealed obstacles of different kinds. A deeper diagnosis of such implementation problems is needed and can be facilitated if the dynamics of information flows for process control are looked into systematically.

An effective implementation of a system for organisational learning from operational surprises, e.g. SINS, must therefore (by design) form an integral part of a viable system-in-focus. SINS is intended to be a generic concept of essential functions for organised and cost-effective learning from incidents. Hence, the systemic translation from one organisational application context into another requires a transformation method.

The Viable System Model (VSM), see Figure 7-1 below [Beer 79, 81, 85], was designed as a powerful help for diagnosing an organisation for its viability. It can also function as a transformation scheme to adapt SINS for other application domains and other implementation levels. The usability of VSM for evaluation and propagation of SINS is discussed in this chapter, posing the following questions:

1. Can descriptions and problems from previous chapters be accommodated in the VSM view, i.e. can the VSM diagnosis of a system for incident monitoring and learning categorise and further illuminate issues arrived at the hard way in the projects discussed in previous chapters?
2. Does the VSM organise the conclusions drawn so far?
3. What does VSM tell us about the work and findings described in the preceding chapters, especially regarding the SINS concept. Does the VSM diagnosis reveal issues of design and implementation of a SINS not identified yet in the previous work?
4. Does VSM provide means to translate the SINS concept as implemented in a hospital context into another domain of organisational activity?

7.2 What has VSM to offer?

The VSM, see Figure 7-1, is a model built from axioms and laws and represents functions that are essential in the operation and management of a viable system. It is used to diagnose (organisational) systems for viability problems and to help specify improvements in the functions assuring viability in the so-called System-in-Focus (SiF).

The VSM model is recursive with distinct sorts of couplings within and between recursion levels. A system-in-focus is situated at one recursion level, e.g. a work team. The purpose stated for the SiF is critical when determining neighbouring recursions up- and downwards, in this case a department and an individual team member respectively. One specific SiF may have many recursions in different dimensions depending on the purpose

selected. For instance, if we take a person as the SiF, than neighbouring recursions will be different for different purposes of the persons' functioning, e.g. being the head of a family, a politician, a narrator or a bus driver: each of these may be valid for one and the same person. See Appendix P for a brief introduction to the VSM.




Figure 7-1: *the viable system (after Beer, 1985)*

An essential notion in VSM is the acknowledgement of Ashby's law of Requisite Variety [Ashby 56]:

"Only variety can absorb variety." [Ashby 56]

It therefore has mechanisms for attenuating and increasing variety in different parts of the system.

The VSM is built up from five separate system functions – Systems One to Five – and six different categories of transmission channels at any recursion level. Figure 7-1 depicts the recursion of a SiF, the next lower recursion in each of the Systems One, and is itself a System One of the next recursion upwards. By adequate attenuation and amplification of variety contained in messages through these channels, the SiF will do what it is meant to do, but will also maintain its viability with respect to the SiF's *raison d'être*. A pragmatic insight from Beer's work is that not even the manager of an operational unit can grasp the reality of that unit in full, although the manager actually manages the process. Beer calls this basic phenomenon the 'muddy box' view held by the manager [Beer 79].

The different functions of the Systems One to Five may all be realised by a single actor, e.g. in a successful one-man company. This person not only does all the work (System One) including coping with its variations and deviations (System Two), but also manages his own work (System Three), explores the outside world for future business prospects (System Four) and reconciles problems, when System Four and Three interactions lose homeostasis, by 'informing' Systems Three-Four about 'right' or 'wrong' in the discrepancy observed (System Five). In larger organisations all of these may be done by different actors. Thus, the VSM is a functional model and *not* a structural model of successful organisations that tells how to draw an organisation chart depicting the allocation of formal responsibilities.

VSM provides a compact, but comprehensive framework for analysis of information flows essential for operational control and viability of a System-in-Focus. This makes it fit for use as a tool in diagnosing SINS problems and, ultimately, in designing better SINSs.

7.3 VSM Diagnosis of operational Surprises

The VSM model is recursive and does not depict an organisational chart. In order to use it, one must define the system in focus (SiF), noting its recursions one level up as well as one level down. As Beer points out, the viability of a system SiF is constrained. A human being is a viable system, as he or she is able to maintain a separate existence and also functions within a family, community, etc., but cannot survive in a vacuum:

"An organisation is viable if it can survive in a particular environment." [Beer 85, p. 1]

The VSM model of the SiF (recursion 1) always contains the Systems One that actually are the *raison d'être* of the SiF. These are the parts, which carry out the function, e.g. individual actors in an operating theatre team (SiF). It connects to a 'parent' system – one recursion level up – in which the SiF is a System One: the CAPU in the case of the OK team as SiF, see Figure 7-2.

Accidents, incidents and other operational surprises occur within a System One of the SiF and represent – potentially – an operational instability that might affect other Systems One of the SiF. System Three manages the work of the various Systems One. System Two is the anti-oscillatory component of the SiF that links System Three with the Systems One independent of the command, bargaining and accountability channels. Its sole purpose is damping oscillatory behaviour between the Systems One. System Two therefore plays a vital role in on-line risk control. System Two is dependent on the 'senior management' that comprises Systems Three, Four and Five, and is connected to the local regulatory

elements that link a System One's operation with its management. An example of a System Two is the roster of lessons for a semester: students, teachers and school management know exactly when to go where for a specific course without the need of repetitive communication through command and resource bargaining channels of the respective Systems Three. The roster itself does not generate the lessons, but it damps

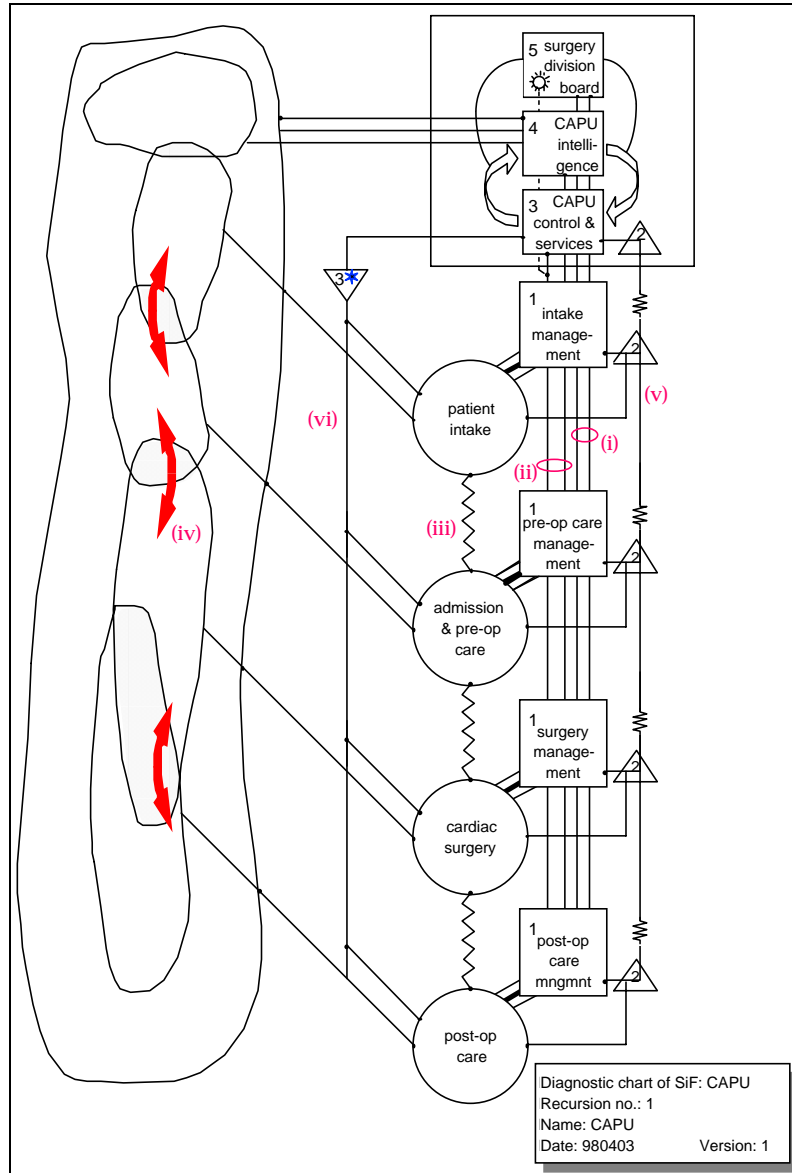


Figure 7-2: the Viable System Model of the cardiopulmonary surgery unit (CAPU) as the System in Focus (SiF), i.e. at recursion level 1. Note that the CAPU is defined by what it does: perform cardiac surgery.

much oscillation when students and teachers swap in order to go to the next course room. It is in System Two that much of the work of a SINS is located.

Supported by the System Two stabilising functions, System Three manages "inside-and-now" the Systems One. System Four is focussed on "outside-and-future" in order to explore the needs and opportunities for changes in the Systems One operations, i.e. the course of development of the SiF. System Five makes the high-level policy choices for the SiF. This allows System Five – as a huge variety attenuator – to settle arguments between Systems Three and Four and, thus, provide homeostasis in the System Four-Three interactions.

7.3.1 Communication Channels in the VSM

The VSM identifies six types of communication channels [Beer 85, p. 83], depicted in Figure 7-2:

- (i) the command channel (corporate and regulatory intervention: top-down)
- (ii) the resource bargaining and accountability channels (two-way)
- (iii) the operational (squiggly-line) linkages between the Systems One (two-way)
- (iv) the intersects of environments of the Systems One, i.e. the linkages of System One elements through shared environments (e.g. geography, markets, consumer-base; two-way))
- (v) the anti-oscillatory linkage (two-way)
- (vi) the System Three-Star (sporadic) audit channel (elicitation of bottom-up data on top-down initiative by System Three)

Outside direct management control channels exist to transmit alert or alarm signals, which Beer refers to as 'algedonic' channels. The term 'algedonic' is constructed from the Greek words for pain and pleasure. Through this channel, a distress message 'ouch, this hurts' may be propagated from within a System One without any intermediate contents processing up to System Five of the SiF, if not to a System Five at a higher recursion. Whistle blowers in organisations are extreme examples of such algedonic communications.

Channel (iii) represents the operational interactions between Systems One. It contains any product that is transferred from one operational unit to another one. In the case of the CAPU, channel (iii) typically contains cardiac surgery patients, as a patient moves from one phase into the next one. The patient enters and leaves the CAPU system at the next recursion up (recursion 0).

The impact of an operational disturbance in one of the Systems One is propagated mainly through channels (iii) and (v), and possibly also through the environment. For instance, an unscheduled delay in the start-up of a surgical procedure implies for a subsequent patient the re-scheduling of the operation to another day. If the damping by System Two is inadequate, System Three receives such a message through System Two. The System Three may need to peek into the troublesome System One through a System Three-Star audit in order to generate variety that is sufficient for problem analysis and resolution. Then, System Three uses the command and bargaining channels to restore (indirectly) homeostasis in and between the Systems One, e.g. by changing the priorities for the scheduled operations based on new assessments of urgency.

7.3.2 Anti-oscillatory Behaviour

How does this work? Start again with an operational surprise within a System One. The local regulator of System One, which connects to System Two, receives a signal from the process through transducers, e.g. a patient monitoring device that activates an alarm. This signal is classified before it is passed to System One's manager and - after filtering - is relayed to other Systems One [Beer 81, p. 170]. Table 7-1 shows how the patient monitor alarm at a medium care unit (MCU) invokes System Two, and/or System Three when an emergency call protocol exists within a cluster of wards. The information from the local regulator, i.e. the alarm of the MCU, has a lower variety than the operational situation monitored, i.e. that the patient in trouble may have a fibrillating heart condition or that the signal may be a false alarm. The System One manager's response to the situation is of strategic or tactical nature and fed back to the operational process and, also after filtering, to the regulatory centre System Two of the SiF. For instance, System One management may decide that nursing staff will (have to) intervene without being qualified in the case of a fibrillating heart condition if a doctor arrives does not arrive within a specific time. In general, the anti-oscillatory states that are possible are shown in Table 7-2. From this table and the VSM model we see that at either level of regulatory control dynamic criteria apply to the transmission to System Two and through to senior management (System Three) of - filtered - operational instability notification messages.

Table 7-1: *anti-oscillatory mechanisms in terms of the VSM – example of a patient monitoring alarm in the MCU*

System challenged by with the internal instability	Outcome	Mechanism
System One local regulatory centre (MCU ward)	Succeeds	Alarm is rightly classified as false, or stop by spontaneous recovery, or local intervention is feasible and adequate
	Fails	Problem requires intervention by a doctor who is not around or comes too late.
System Two regulatory centre (all wards): alarm beep protocol states that at the beep every doctors drops the task at hand and runs to the alarm-generating location... the first arriving doctor assists in the emergency, the others return to their posts.	Succeeds	At least one doctor arrives timely and intervenes successfully.
	Fails gradually	Doctors are not available or arrive too late, but System One intervenes successfully without being authorised. System Three may have to change resources (assign a doctor to the MCU, or authorise staff members to intervene after an 'upgrade' training).
	Fails completely (IncT = 0)	Doctors arrive too late or not at all, and System One is not able to intervene successfully. System Three needs to, for instance, supply resources to the System(s) One, e.g. more qualified staff or better diagnostic means, and/or needs to by modify the System Two alarm protocol to help short-staffed night and weekend shifts.

Note that System Three actions are (permanently) influenced by System Four's exploration of future directions for the SiF, see also Figure 7-1. Thus, a System Three intervention triggered by System Two may speed up the evolution of the SiF into the future, by permitting the whole viable system to shift its ways of working. For example, new and more reliable equipment may be introduced to replace current apparatus that regularly breaks down, producing Systems One and Two oscillations, e.g. a new genera-

tion of patient monitors that alarm more reliably and have distinct alarm signals for different classes of potential problems.

Table 7-2: *anti-oscillatory mechanisms in terms of the VSM*

System challenged by with the internal instability	Outcome	Mechanism
System One local regulatory centre	Succeeds	The most common situation. A weak signal may, in addition, be relayed to local regulators of other Systems One or to the SiF's System Two, but does not need elicitation.
	Fails	Regulatory action by System Two of the SiF is needed.
System Two regulatory centre	Succeeds	System Two damps oscillatory behaviour among Systems One.
	Fails gradually	Action by System Three of the SiF (through the command or bargaining channels) may be required to put System One back into business or to eliminate a specific oscillation trigger. For this purpose, System Two may be designed to count a number of incident occurrences before a System Three action is invoked.
	Fails completely (IncT = 0)	Action by System Three of the SiF (through the command or bargaining channels) is essential to put System One back into business. If System Three cannot provide a resolution by generating requisite variety, the SiF is in oscillation which must be resolved by the parent system at the next recursion up where our SiF is a System One.

We see here that a local regulator filters messages about an operational surprise that has occurred, when relaying data to local regulators of other Systems One or to the SiF's regulatory centre. The sensitivity with which the monitoring signal is transmitted from the operational process to the local regulator must be taken into consideration in a re-design of System Two. For instance, the 'synapse' connection in this link may have a high threshold limit below which no signal is transmitted [Beer 81, p. 170]. Similarly, the filter between the local regulator and System Two may be too narrow, so that no signals come through. Thus, the unexpected death of a patient may not be signalled beyond the local regulator of System One when this operational outcome is written off as a medical complication rather than an operational surprise.

Systemic solutions for System One instability problems take the form of intervention rules or policy explanations – through the command channel (i) – or changes in resource programmes through channel (ii). For instance, repetitive failure by doctors to arrive in time at the alarm-generating bedside may trigger System Three to decide that a specific category of patients-at-risk must stay in the ICU rather than in the MCU to ensure adequate and timely intervention when required. Alternatively, System Three may decide that some members of the MCU nursing staff are sent on a qualification programme, e.g. to become qualified to defibrillate a patient with a fibrillation heart condition, so that adequate resources become permanently available within the MCU.

Information on local operational disruptions and the damping of oscillations throughout the SiF is disseminated through the local regulators of System Two, and might be shared with other organisational systems through the System Four-Three interactions with the environment. A monitoring device that generates many false alarms (disruption in a

System One) may be marked for early replacement and, for the time being, call for a 'double-check-the-alarm' protocol. Other health care systems may have experienced similar 'false alarm' problems with the monitor and become informed through a shared environment accessed by System Four. Lessons available at System Three may also be propagated through the next higher recursion of which the SiF itself is one of the Systems One. For instance, the (limited) effectiveness of the alarm beep protocol may eventually lead to modification of legal requirements for allocation of professional responsibilities, e.g. a doctor's task that is transferred to a nurse qualified for this specific task.

Systemic causes underlying operational surprises in Systems One of the SiF appear in VSM as a mismatch in variety absorption within the SiF and also in the interaction with the next recursion upwards. Systemic problem solution implies a policy change imposed by System Three through the command channel (i), an adjustment of resource bargaining through channel (ii), and/or a re-design of the anti-oscillatory mechanisms in System Two. System Three uses system Three-Star audits to understand the causes of the oscillations. These are designed to collect rich data of suitable variety to re-establish control, in contrast to the attenuated signals reaching it via System Two. This distinction between signals via these two lines usefully distinguishes in-depth problem solving from routine monitoring of performance indicators. The latter may trigger the former if patterns of problems are found. System Three is also influenced by its interactions with System Four that analyses developments in the outside world.

At this stage, we conclude that the VSM allows analysis of operational instabilities, such as accidents or other operational surprises. Learning within VSM takes place in the form of adjustments of System Two mechanisms and System Three management decisions to modify policy, resources and accountability of one or more operational elements. From the VSM model we can also see that all sorts of organisational learning process blockages are possible. For instance, too much signal filtering in System Two, System Two not effectively connected to local regulators, System Two mechanisms overloading the System Two-Three link, System Three making no or inadequate decisions including failure to activate a System Three-Star audit or inquiry, messages being propagated through inappropriate channels, etc.

7.4 Concept of Organisational Learning and VSM

The VSM can also accommodate the concept of organisational learning as discussed in Chapter 4 and depicted in Figure 4-1 (reprinted below). This says that the organisation only learns through its learning agency that is configured to function in single- and double loop modes. A learning loop ends when the results of adjustments in the operations of the organisational unit become apparent: match or another mismatch.

The organisational unit in Figure 7-3 may be seen as a System One of the System in Focus, see Figure 7-2. The OL-'agency' functions are distributed over System Two (classification of disturbances, memory) and senior management (Systems Three-Four-Five). System Three is prompted by the System Two anti-oscillatory regulator to intervene through resources bargaining or policy propagation, possibly after a sporadic System Three-Star audit into Systems One. System Three is influenced by its interaction

with System Four that might have identified experiences with and solutions to similar operational problems elsewhere.

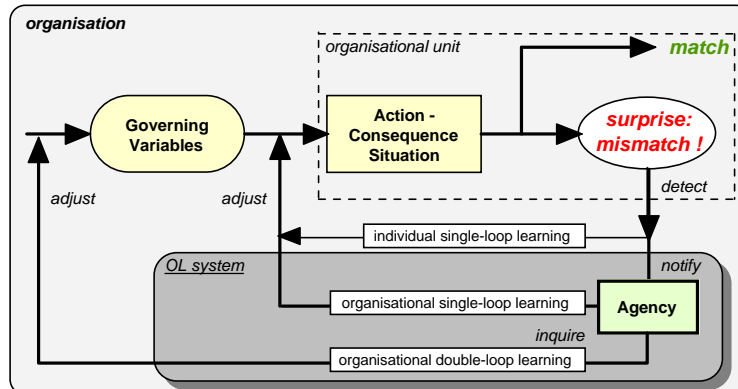


Figure 7-3: Argyris' concept of Organisational Learning (OL) - OL system marked

In the VSM view, individual learning takes place at the recursion level where individuals are the Systems One of an organisational unit, e.g. a ward or the ICU (see text box below on recursive views...). Individuals may learn informally from each other by interactions via channels (iii) and (iv). Organisational learning requires the presence of an agency that is designed to identify changes needed in the channels (ii) and (i) regarding resource bargaining and policy respectively to resolve oscillation problems among Systems One. This lifts the learning to the higher recursion levels of organisational units or the whole organisation and may also spread them to the neighbouring Systems One at each recursion level (to other units or organisations).

On recursive views

The cardiac surgery unit CAPU, see chapter 6, can be taken as an example of a 'System-in-Focus'. The location manager performs System Three roles. The distinct wards, the medical staff and the OR-unit are Systems One of the SiF (see also Appendix Q). An individual nurse within a specific ward operates as a System One at one recursion downwards. His or her physiological systems are functioning as Systems One at the next lower recursion, i.e. two recursions down from the SiF. **Zero Learning** takes place at this level: the perceptive subsystem of the individual observes changes in his environment that are evaluated by his System Two as a surprise. It is passed on to his internal System Three (zero learning) or filtered out [Beer 81]. The settings of the System Two filters inside the individual are subject to the safety climate or culture prevailing in the system one recursion upwards, e.g. the ward wherein the individual nurse is a System One. On the basis of these settings the individual's System Three also decides to **notify** – or not – the surprise experienced, by signalling System Two of the next recursion upwards, e.g. that of the ward. The filter settings in this System Two (of a System One of the SiF) can and should be designed by the organisation and the same applies to settings for System Two of the SiF.

Table 7-3 is a concept-translation table with elements from the organisational learning (OL) concept and their transformation in VSM terms. Some of the OL elements are ambiguous in VSM terms, meaning that such an OL element represents several distinct VSM functions. This means that a VSM-analysis can be more incisive than an OL-analysis. The light-shaded lower part of Table 7-3 below addresses terms that are essential in OL, but are not depicted in Figure 7-3 above.

Table 7-3: translation of the Organisational Learning (Argyris) concept into VSM (Beer)

Basic OL components	VSM analogy
Organisational unit	A System One of a SiF, i.e. the management unit + the operational processes that it manages.
Mismatch	System One instability/process disruption
Agency	System Two (classification) + System Three-Four interaction + System Three-Star audit
Individual learning	System One that learns to restore operational stability without effectively signalling System Two (agency) of the SiF: see also text box above "on recursive views".
Single loop learning	System One – local regulatory loop (System Two)
Double loop learning	System Two – Three (-Star) – System One [System Three in homeostasis with System Four System Five]; may include System Three – System One interaction one recursion higher
Detect	See text box "on recursive views..." above: System Two input synapse fires at individual level (2 recursions down) [unless automatic detection is invoked]; or signal through the 'algedonic' channel to System Five !
Notify	System Two input fires at System One level through message from individual who had a 'detect' experience [unless automatic notification follows detection]
Inquire	System Three-Star audit; or System Two verification of incoming System One signals
Adjust	Restore operational stability through System Two-One interaction - single loop - or by System Three intervention through channels (ii) and (i) -double loop
Action-Consequence situation	Operational state resulting from a System One process
Governing variables	System One constraints set by resources bargains in channel (ii) + by policy requirements via channel (i) and corporate ethos in System Two
Learning agent	System Three-Star function of agency
Deuterolearning	Learning how to learn: a) by System Three-System Three-Star of parent system of SiF: 1 recursion up; b) by System Four-Three interaction whereby System Four explores learning by other organisational systems
Organisational memory	System Four self-reference + System Two anti-oscillatory mechanisms where there is a "lessons learned" case base, it relieves the System Two workload of the learning agency by storing possible interventions to damping oscillation and by remembering what messages have already been sent to System Three.
Defensive learning	a) System Two damping too strongly, b) System Five homeostasis of System Three-Four interaction suppressing System Four-Three interaction, c) inadequate use of communication channels, e.g. (i) instead of (ii) or via System Two, d) System Two not carefully designed, and/or e) System Two imposed by System Three rather than designed with Systems One involved

Argyris' concept of Theory of Action as a composite of Espoused Theory and Theory-in-Use cannot be mapped into VSM this way. It would weaken both the impact of Argyris' view and the VSM ideas to include this. However, that the mapping does not embrace this dimension of Argyris' theory is not a problem to the equivalence between Argyris and VSM. In fact it is a positive matter – some of the denaturing of vertical channels in VSM can be explained using Argyris mechanisms.

The mapping of organisational memory is actually inadequate. The VSM is built from Ashbian ideas [Ashby 56] – much of which imply or explicitly require memory in the regulatory components (e.g. homeostats). Beer notes that in the VSM homeostats "requisite variety applies in three distinct ways: to the blocks of variety homeostatically related, to the channels carrying information between them, and to the transducers relaying information across boundaries" [Beer 89]. E.g. the devotee of classical music may only

listen to a radio station that broadcasts concert recordings, using HIFI audio equipment in a acoustically appropriate space, in order to enjoy the richness of a symphony when he or she is not to a live concert. In other words, wherever there is a line on the VSM, memory is implied at one or both ends of it, see also Figure 7-1. Thus, it is possible that certain transducers and channels involve memory like the vacuum tubes of Alan Turing's computer, Colossus [Sale 96]. The strength of VSM is that although some areas may appear more important than others, memory cannot be reduced to such specific areas – the operation of the whole system is needed.

We conclude at this stage that the OL components, learning loops and essential variables can all be transferred into VSM terms. The different learning loops (OL) can be precisely pinpointed in VSM both within a recursion as well as between recursions and beyond (via the environments). The OL concept does *not* have sufficient variety to accommodate all of the corresponding VSM functions. The VSM depicts an infinite number of learning interactions within a SiF as well as between recursions, which makes the VSM look complex. The OL concept is simple and provides less profound insight, but is easier to communicate with organisations that want to improve the organisational learning processes. It should be retained for that purpose. However, being a truly recursive model, the VSM allows us to draw up an endless number of recursions in order to map the interactions of learning processes between recursions and within a recursion, e.g. the SiF. Although such a VSM depiction looks complex, it provides more profound and more precise insights into the issues of productive and defensive organisational learning. Therefore, the SINS concept will be diagnosed for its viability in the next section using the VSM.

7.5 SINS in VSM

The SINS-concept described in Chapter 5 and in part implemented in the PIERCE project, see Chapter 6, can be translated in the following way into VSM terms, see also Appendix 7-1. SINS has been (partially) implemented at the recursion of the CAPU as the System-in-Focus. Operational surprises are usually observed by workers within one of the Systems One of the CAPU (see also the text box on "On recursive views" above). Hence, the 'work process' box in Figure 7-4 represents the various Systems One of the SiF: wards, ICU, medical staff, and OR unit.

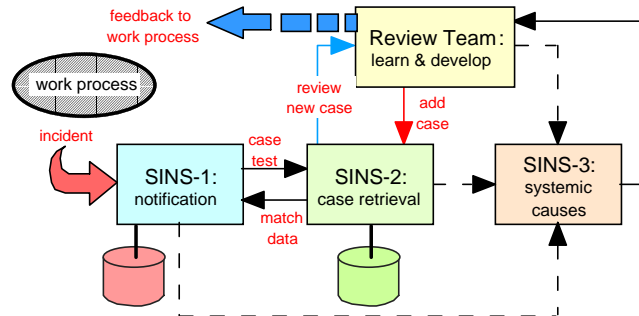


Figure 7-4: SINS concept [Figure 5-2]

SINS-1 (notification) and SINS-2 (case base) are System Two elements. Their prime function is to record operational disturbances and interventions to regain control in and between the Systems One, when an undesired oscillatory behaviour of a System One has been detected and notified. In CAPU, the various Systems One are tightly coupled through channel (iii). This follows from the fact that the main objects for transformation by operational activity are the patients who must move from one System One to a subsequent one in order to realise the prime purpose of the CAPU. The CAPU's *raison d'être* is to deliver cardiac surgery to patients who will gain in "quality of life" from this. A member of the CAPU who submits a notification message using SINS-1 is, in turn, a System One at recursion (CAPU) – 1.

The SINS-2 case base is part of the CAPU's organisational memory that is designed to store – in an accessible way – the experiences with operational disturbances. This should be seen as part of the System Two. Consultation of this organisational memory may help to resolve the notified operational problem more quickly or more effectively.

SINS-3 is not automatically invoked by System Two, but requires a System Three or a Review Team decision (both are possible). SINS-3 becomes a System Three-Star function when peeking into System One operations, which is necessary to identify concrete systemic causes underlying a System One disturbance or a repeating pattern of disturbances. Such systemic causes can be found at this same recursion of the SiF, i.e. CAPU, as well as at metasystemic recursions (upwards).

The Review Team, as the organisational learning agency, is a complex component with several distinct functions:

1. the RT performs a System Three-Star audit when reviewing and validating a case. The need for such a role, of peeking into Systems One, is reduced when members of the Review Team also come from these Systems One: they already embody (almost enough) requisite variety needed to absorb the variety related to the disturbances in Systems One. They already understand a lot about the case and its possible causes and solutions. However, they still need Three-Star audits at intervals to ensure that this prior knowledge does not blind them with their own prejudices.
2. the RT is a designer/constructor of System Two, when it stores a new case in the SINS-2 case base. An evaluation scheme for classification of surprise-enabling conditions in terms of cause categories, for instance the one shown in Appendix N, is just another filter feature in the System Two memory.
3. the RT influences the System Two filters – through the policy channel (i) also on lower recursions of the SiF – that mirror the corporate ethos about operational risks. For example, by demonstrating the use and relevance of lessons learned, individuals within the Systems One of the SiF may modify their own System Two filters (recursion – 2) so that it increases their sensitivity to detect operational surprises and to decide to notify them.

The feedback to work process action is equally complex since System Three, the here-and-now manager of the SiF, must employ both the channels for resource bargaining (ii) and policy propagation (i). The PIERCE project revealed that feedback information from the Review team informally passed through channels (iii) and (iv) to neighbouring Systems One could be very effective as well. An explanation for the phenomenon might be that management at the lower recursion (-1) can sometimes implement changes in

operations within their span of regulatory control, without the need for metasystemic intervention by or through System Three of the SiF. This would seem to be more easily achievable in these cases by peer, than by hierarchical pressure, perhaps a typical phenomenon of a professional bureaucracy [Mintzberg 83] such as a hospital.

Table 7-4: translation of the SINS concept into VSM (Beer)

Basic SINS components	VSM analogy
Work Process	A System One of the SiF, e.g. the pre-op ward of the CAPU
Incident	Operational surprise, i.e. disturbance inside a System One
SINS-1 notification	Input synapse for System Two [initial crisis management is taken care of via local regulator of a System One]
SINS-1 database	Organisational memory situated in System Two (latent, indirect)
Notifier (curved arrow)	Activator of System Two input synapse signalling incident occurrence possibly from within System One operations
SINS-2 case base	a) System Two (classification)..., b) organisational memory contained in System Two (anti-oscillatory!): known problems can be resolved quicker, and/or c) ...possibly shared with other recursions / other organisational systems through System Three-System Four
SINS-3 in-depth (root cause analysis = RCA)	System Three-Star function (RCA) and/or System Three-Four reflection
Feedback to work process	Implementation of adjustments to System One by the local regulator, through System Two, and/or through System Three(-Four) channels (ii) or (i). Influencing System Two settings of individuals (2 recursions downwards) through ethos (System Two of SiF), and/or channel (iii) interactions between Systems One and via shared environments [channels (iv)].
Review Team (RT)	System Three-Four/System Three-Star review of a System One instability problem/oscillatory behaviour, and validation of System Two (case base) memory
Learning agent	a) RT co-ordinator (see above), b) System Two classifier, and/or c) System Three-Star investigator
Lessons	a) changes in regulatory commands (i) and resource and accountability bargains (ii) to resolve System One oscillation, and/or b) improved System Two (damping) components... lessons learned by RT or stored in organisational memory: System Two case base
Throughput model of work process	Systems One connected in sequence through operational linkages (iii)

Table 7-4 is a translation table with elements from the SINS-concept and their transformation in VSM terms. The light-shaded lower part of Table 7-4 addresses terms that are essential in SINS, but are not depicted in Figure 7-4 above.

Note that organisational memory of operational surprises and their solutions forms a set of system state classification criteria. If an operational irregularity is known in memory, the corrective action or intervention is also known and available for re-use. Thus, an incident case base, if consulted, has an anti-oscillatory function.

In Chapter 6, the learning agent was found to be necessary for the learning agency: without preparations for a case review the precious resources of the (loss-of-variety compensating) domain experts in the agency were not made available or were wasted. The learning agent furthermore classifies signals that have been passed on by local regulators for this purpose, e.g. when selecting from operational surprises that were notified to SINS-1. He therefore sits in the channels (v) and (vi) to facilitate their working.

In VSM lessons learned are also propagated through informal links between different System Ones and through (shared) environments i.e. through channels (iii) and (iv). Learning by the SiF in this way takes place less directly than the organised processing of its own operational surprises through System One-Two-Three-Three Star. It requires an organised pathway through System Three to System Four, which identifies useful lessons elsewhere. This is active external learning. Alternatively, the link can be a channel (iv), e.g. two people chatting at a party.

System Two regulatory mechanisms in CAPU include corporate and professional ethos [Beer 95], a waiting list and surgery rosters, classification schemes (complication, incidents), and some of the regular meetings, see Table 7-5.

Table 7-5: *regular meetings in the cardiac surgery unit (VSM view)*

Meeting	Frequency	present	Topic of the agenda	VSM SiF system	VSM channel
morning session	daily	surgical staff + residents	all patients that require special attention + irregularities	System 3 System 2	(ii) (v)
heart team (d)	daily	surgical staff + cardiologists	classify up to 25 patients: reject, cardiology, operate	System 2	(v)
hand-over meeting	daily	nursing staff	all patients that require special attention + irregularities	System 3 and 2 at one recursion down	(ii) and (v) recursion SiF - 1
heart team (w)	weekly	surgical staff + cardiologists	assessment of problematic patients (state-of-the-art probing)	System 3-4	(ii), (iv)
heart team (m)	monthly	surgical staff + cardiologists	medical issues	System 4	(iv)
surgery platform	monthly	surgical staff + anaesthetists + ICU-staff	medical issues	System 3-4	(ii), (iv)
staff	monthly	surgical staff	policy, general topics	System 4	(iv)
ICU staff	daily	ICU staff	all patients that require special attention + irregularities	System 3 and 2 at one recursion down	(ii) and (v) recursion SiF - 1
perfusion platform	quarterly	perfusionists + cardiac surgery unit manager	irregularities + technical issues	System 2 System 4	(v), (iv)

7.5.1 Discussion VSM – SINS

We conclude at this stage that the key functions in the SINS concept – detection, memory, learning agency and systemic cause analysis – can be projected onto the VSM model. As an organisational memory for operational surprises the SINS-2 case base is part of System Two and anchored to the SiF recursion (CAPU level, see Figure 7-2). This is an important new element in the learning system, which is not so explicit in VSM, but can be located there. The SINS concept as represented in Figure 7-4 does not show the links with other organisational levels, whereas the VSM reveals the various communication routes (within a recursion as well as between recursions) needed for learning by a viable system. For example, once invoked, the SINS-3 function usually requires (metasystemic) solutions

provided by senior management (System Three-Four-Five) if not from a higher recursion. These correspond with double loop organisational learning and/or deuterolearning (according to Argyris). An example from the PIERCE project may illustrate the above: see the text box below.

The skin needle case

While closing up a patient in the OR, a skin needle had dropped unnoticed into the pleura-cavity. Three X-ray pictures taken as protocol steps clearly showed this needle, but four physicians on duty did not see it. The location manager, who as a quality assurance action regularly reviews X-ray pictures of CAPU patients, spotted the needle during such a picture review. The patient had to return to the OR five days after the cardiac surgery in order to have the needle removed. The manager ordered that medical staff must apply a trained method for systematic examination of X-ray pictures of their patients.

The OR process represents a System One of the SiF: CAPU. The location manager is mainly functioning as a System Three. The X-ray pictures are made by radiologists who provide services to the CAPU. The four doctors operate as Systems One one recursion down in one of the Systems One of the SiF: ICU, MCU, regular post-op ward. The counting protocol in the OR is a local regulatory System Two component. The protocol did not prescribe the counting of skin needles, because they are not used inside the patient's body. The operational surprise was identified when the patient had already been moved through channel (iii) into a third System One of the SiF: the operational disruption occurred in the OR and from there the patient was transferred to the ICU, the MCU and to the regular post-op ward. The signal came from System Three of the SiF – instead of System Two – when System Three was performing a periodic System-Three Star audit by reviewing an arbitrary selection of patient files. The location manager as System Three used the policy and command channel (i) to propagate the message to doctors of CAPU to apply the method of X-ray picture examination that was trained some time ago. The Review Team was informed by System Three about the surprise detection and performed a System Three-Star audit of the occurrence of the initiating operational surprise: the unnoticed dropping of one skin needle.

From this review, two interventions were considered for their effectiveness. The first one was to include the skin needle in the counting protocol of the OR-nursing staff. This measure implies the adjustment of a local regulatory System Two component, i.e. the counting protocol in the OR.

The second one was to assure that medical staff applies and improves its skills to read X-ray pictures (the location manager's solution). This requires the anchoring of systematic X-ray picture review as a policy and protocol item in the Systems Two of the individual doctors being Systems One at one recursion down from the SiF. Hence, multiple lower recursion Systems One must be reached from a higher recursion through command and resource bargaining/accountability channels (i, respectively ii).

It can be seen that the first option is to be preferred: less complex to implement and easier to assure. In practice, members of the Review Team – also members of Systems One of the SiF – used channels (iv) and (iii) to inform senior management of the System One involved about the simple solution. The preferred solution, i.e. adjustment of the counting protocol as a System Two element of the OR unit at one recursion down from the SiF (CAPU), was implemented the next day after the review.

So, we see in this example that the lack of variety expressed in one or more channels, e.g. (i) and (ii), has been compensated through other channels that were available: (iv) and (iii).

7.6 VSM Diagnosis of Lessons Learned

In the previous section we have seen that The Viable System Model enables us to diagnose a specific organisation for its capabilities of organisational learning from

operational surprises and to assess the configuration of an incident review system like SINS. In this section earlier projects and lessons learned from these projects are diagnosed briefly using the VSM in order to arrive at overall conclusions from this thesis concerning the viability of "organised learning from small-scale incidents" systems.

7.6.1 Incident Data Processing in VSM Perspective

Chapter Two describes the early projects carried out to collect data on accidents for learning purposes. Recognition of harm mechanisms or scenarios, identification of systemic factors and patterns of such factors that related to a type of scenario were leading issues in section 2.2. The projects in section 2.3 aimed at cost-effective one-off analysis of small-scale incidents within a setting of organisational learning. Diagnosing these project results for viability using VSM leads to the following considerations, see also the text box on "VSM components in early projects" below. The next paragraphs re-examine conclusions drawn in Chapter 2.

Data on accidents, etc. that are collected without the purpose of restoring operational instabilities due to the surprise occurrence in and among Systems One have no local function. They go to higher recursions through the resource bargaining and accountability channel (ii), usually as performance indicators like FAFR, LTI frequency rate, etc., and are used to trigger System Three-Star interventions or other System Three actions. As such it is not surprising that Systems One regard them as irrelevant bureaucracy at best, and signals to bring down trouble on their heads at worst.

Large incident data sets (as required or aimed for in the early projects) are part of organisational memory used at (several) recursions up from the recursions where the System One disturbances occur. They cannot have direct relevance for anti-oscillatory damping behaviour of the lower recursion level System Two. Data comes from the lower recursion through the (vertical) bargaining and accountability channels (ii) whereby considerable variety loss in the data about System One disturbances is inevitable. This means that all that is left once they reach higher recursions, are performance indicators. The projects proved that great variety cannot be transmitted easily to higher levels. The higher recursions must initiate a System Three-Star audit to restore the variety.

However, when System One-instability problems are caused by higher recursions or by the environment, learning at several recursions higher than the Systems One where the operational surprises occur might be effective. For instance, operational problems that relate to people who lack specific training, are likely caused by an inadequate personnel selection and/or by an insufficient training programme, both causal factors being systemic and rooted in a higher recursion. The anti-oscillatory System Two must be designed to filter out data that is not relevant to the higher recursion. The higher recursion System Three may have to counteract through its resource bargaining channel (ii) and/or the policy and command channel (i). In particular cases, System One signals may be passed on directly to a higher recursion management system. Illustrative examples concerning the reporting of medical device-related problems are discussed in section 7.6.1.1 below.

VSM components in early projects (two examples)

Diagnosing incident monitoring systems using VSM starts with the identification of the Systems One where the surprise information comes from.

In any project, the Systems One and the System(s)-in-Focus must first be known. Then System Three must be identified. System Three can decide to solve problems of oscillatory behaviour of Systems One within span-of-control limits, including by calling for higher recursion intervention. Then System Two can be identified, linking Systems One with System Three to achieve oscillation control. The examples presented below illustrate that Systems One exist in different dimensions.

The DOHO project was curious as the notification data source was the injured person picked up at the gate in an ambulance on its way to the nearest hospital. These persons were a lower recursion System One where the production unit would be a System One. There was no direct contact in the DOHO project with the Systems One of the SiF.

case: DOHO

SiF = Rotterdam harbour
 Systems One = harbour enterprises (in distinct sectors)
 System Three = not clearly identified: authorities (Rotterdam Harbour, health and labour inspectorate)
 System Two = regional OSH services
 Recursion – 1 = production units within a harbour company
 Recursion – 2 = victim in a production unit
 Type of signal from System One: .. injury accident in company *x* of sector *y*
 Type of System Three response:.... N/A ! No System Three intervention was anticipated.
 Aggregated data: for research purposes.

The medical device reporting system demonstrates that surprise monitoring can be set up "irrespective" of the organisational systems or persons where problems are observed and which therefore are the Systems One in a viable MDR system.

case: Medical Device Reporting

SiF = national health care system
 Systems One = health professionals, user facilities; consumers (submission directly by internet; filtered by MDR regulations (see [CDRH 96]), not by health care institutional recursions!)
 System Three = FDA
 System Two = MAUDE (memory) fed by MedWatch (CDHR)
 Type of signal from System One: .. serious problem with device application
 Type of System Three response:.... alert all or selective Systems One; update System Two memory; withdraw device from market; revise PMA regulation; revise System Two (MDR regulation)
 System Three-Star activity:..... investigate how a specific device has been manufactured, or is being applied by a particular user.

The VSM diagnosis reveals that the DOHO project failed, for one reason because there was no System Three in the System One-Two-Three anti-oscillatory control loops. The medical devices example, however, indicates a potentially successful learning system, despite the fact that Systems Three and Two are at macro level: Systems Three, Two and One are all clearly identifiable and configured well.

The VSM view implies that an operational instability in a System One needs to be brought back under control - through System Two/System Three - even when the origin is beyond System Three's span of control as in the case of a systemic medical device deficiency. System Three can, within limits, change resources and System Two in order to improve operational stability despite a defective device, e.g. by increasing surveillance, pending a better solution, such as provision of a device of a better design and make. In order to

achieve the better solution, the device problem message needs to be conveyed to a higher or adjacent recursion System Two. Now the details concerning System One are irrelevant for the higher recursion System Two. The lower recursion System Three is only involved in managing here-and-now System One operations within the resources and means available (including critical devices with embedded flaws). The other recursions have other decisions to make, e.g. whether to buy different devices or put pressure on the manufacturer to improve quality. For that they need different information.

Following this line of reasoning, we must rephrase conclusion 4 of Chapter 2 about the *purpose* of data collection and analysis:

The main purpose of data collection and analysis should be focussed on local learning about anti-oscillatory control and/or on learning about systemic problems that may impact on System One operations, but can only be solved at a higher recursion. The two purposes are not necessarily incompatible, but require different data to be transmitted to different locations (systems and recursions).

VSM modelling shows that *causal models* can have several functions. Where they are incorporated in data collection schemes they are System Two filter design features. A causal model used in this way is a variety attenuator that may simplify the data coming from a System One which is needed by System Two to function effectively.

However, when the causal model is a model of systemic higher recursion functions (e.g. the MORT-based K-model in Company B, ARA) rather than of System One process deviation scenarios (DOHO, Company A), the model is a System Three-Star component that System Three deploys to identify root causes and solutions for the System One problems. The conclusions in Chapter Two concerning the use of a knowledge-model underlying data collection were only projected in a System Two setting rather than as a System Three-Star component. Those conclusions therefore need expanding.

The use of *smart forms* for surprise notification is a System Two function. When combined with *embedded knowledge* on higher recursion functions, it helps System Three to perform a System Three-Star audit in an objectifying way without additional resources. In other words, it combines an information-seeking audit with a simple notification procedure. However, System Three-Star audits need to be sporadic, and such smart forms could raise notification thresholds in the Systems One – System Two linkages. System Two and System Three-Star pose contradictory requirements to variety in the data transmission. Smart forms can be a compromise between these. They allow temporary modifications in the System Two filters in order to collect details on a specific class of System One surprises before – or possibly instead of – invoking a System Three-Star audit concerning this class of operational disturbances. Drawing conclusions from the System Three-Star audit is a System Three function rather than a System One. For that reason, the conclusion drawn in Chapter 2 that "the use of *smart forms* with an *embedded knowledge* body objectifies the elicitation of event and evidence data", makes sense for System Three looking into System One disturbances.

The *detection* of an operational surprise by an individual member of an operational unit occurs at two recursions lower than the recursion where the disturbance occurs in a System One. The individual, being a System One at the next recursion down, functions as 'senior management' of his sensory perception systems (lower recursion Systems One). This lower recursion senior management may decide to signal System Two at the next

recursion up by issuing a notification message. The individual's sensitivity for an emerging operational surprise and his drive to signal System Two is influenced by higher recursions through channels (i) and (ii), the environment interactions through channel (iv), and the design of System Two synapses and filters linked to channel (v). For instance, a high administrative burden that comes with a notification submission increases the threshold in the system One-Two synapse. A policy stating that repetitive involvement in reported operational surprises will be noted in one's personal record, will reduce the individual's sensitivity to report an operational surprise when it might be observed. A reduction of an individual's sensitivity to notice an operational surprise should also be expected when the surprises occur regularly during routine operations, e.g. as in the case of the surgical glove ruptures discussed in Chapter 6.

Such influences can be designed at a higher recursion. The ethos or culture of the organisation propagated through channel (i) and reflected in System Two has a major impact through System Five on the individual's System Three decision – 2 recursions down – to signal System Two about an operational surprise detected. In other words, the individual's perception of the prevailing safety climate can be influenced by higher recursions through different channels with their distinct transduction features, especially channels (i), (ii), (v), but also (iv) and (iii). Performance indicators set in the accountability channel (ii) may include items on 'lessons learned' at the recursion of the SiF, which requires a System Two and channel (v) design that is useful for Systems One. One must be careful with the way incidents are accounted for in channel (ii). The inclusion of reportable surprise occurrences in the set of performance indicators about System One in the accountability channel (ii) may have the adverse effect of increasing the resistance of the System One – System Two synapse. Therefore, fewer operational surprises in Systems One will be signalled to System Two, thus 'killing' the possibility of learning from the System One incidents that are small enough to conceal. For example, in a state-controlled electricity utility the reportable occurrences dropped by 50% in the first month and even more in the subsequent months after management decided to judge the performance of the plant on the reduction of reportable occurrences which could be achieved [Ives 91]. On the other hand, a 'case of the week' propagated among Systems One through channels (iii) or (iv) may increase the individual's trust in the System Two-Three(-Star) interactions concerning handling of operational problems and, hence, the individual's willingness to signal System Two. Interaction between different SiFs through shared environments may spread the positive attitude towards surprise reporting. Thus, the VSM view makes the issue pointed out in the conclusion in Chapter 2 that "the *detection* of an operational anomaly relates to individual sensitivity to notice the *unanticipated surprise*-experience..." much more specific and suggests ways in which it can be coped with.

People who detect an operational surprise can be found in different observer positions: inside the System One that produces the surprise, in another System One of the same recursion which discovers or suffers consequences of the surprise, or residing in the environment of the System One concerned. In the latter case, senior management may be alerted through the shared environment channel (iv) instead of by System Two, probably at a higher recursion. As many of these alerts as possible need to be incorporated in the reporting system, which may need to have several access points and communication channels to accommodate them. The idea behind this is that *at least one* of the different 'surprise' observers signals the surprise occurrence to a relevant System Two. As will be

discussed in section 7.6.1.1 for the case of medical devices reporting, different Systems Two may be relevant to one System One, i.e. that of the own and/or higher recursion of the default System-in-Focus and/or in another dimension. If surprise memory is shared throughout different recursions and/or dimensions of viable systems, the surprise signalling through different access points and communication channel by different observers provides opportunities for optimisation. The various projects illustrate different solutions to this need for redundant reporting channels.

The administrative nature of *thresholds* for notifying operational surprises are caused by System Two design and (incorrect) allocation of data flows (between Systems One-Two-Three). The conclusion in Chapter 2 that "the actual *reporting* and *recording* of detected operational anomalies relates to the *notification threshold*..." mixes System Two design issues with System Three-Star activity. This may be clarified in the following way. Let us take the case that System Three intervention is not invoked by System Two and, therefore, no System Three-Star audits is initiated to generate adequate variety for solving an oscillatory problem through System Three intervention. This could mean that System Two has adequate variety to damp all oscillations that might be generated among the Systems One of the SiF. Or it is because signals coming from Systems One are attenuated so strongly by System Two that no oscillatory behaviour is noticeable until the SiF collapses (being no longer a viable system). The latter situation must be avoided. In the first case, System Two has been designed perfectly and only needs maintenance. However, such a perfect System Two is very unlikely in practice. Therefore, the VSM model implies that an optimal System Two is one which is designed to invoke System Three intervention every now and then, just to show it is functioning. It does this in the case that oscillatory behaviour cannot be damped sufficiently by System Two. A refinement of this would be that System Two invokes System Three intervention when repeatedly occurring System One disturbances are recognised as an annoyance problem that deserves to be resolved permanently. The incentive for the Systems One to signal System Two consistently may be great. Thus, we may postulate that the threshold in the System One – System Two synapse is modulated by the visibility and the efficacy of System Three/Three-Star activity in response to System Two anti-oscillatory behaviour: the less visible this is, the higher the threshold will be. The example of Medical Centre A also demonstrates the impact of ethos or management vigour on threshold lowering in a setting where higher recursions were not considered.

In VSM, the decisions about how to *describe* an operational surprise are issues relating to the desired variety attenuation between the disturbance generating System One and the anti-oscillatory System Two. Too much variety coming in from System One, e.g. in a extensive message describing the surprise, may block System Two (data overload) and will also increase the threshold in the System One – System Two synapse, eventually reducing the number of System One disturbance signals. On the other hand, low variety in the System One surprise message may mean that anti-oscillatory control remains ineffective. In that case, System Three must be triggered by the crude performance indicators from System Two to intervene through channels (i) and (ii), which costs extra resources and complications. Designing organisational memory into System Two can solve this dilemma. It stores lessons learned about anti-oscillatory intervention by System One management into this System Two memory for reuse when a similar System One disturbance occurs. The System Two capability to respond to variety then increases over time. Only signals that really need System Three intervention get through this filter. On

the other hand System Two may also attenuate too strongly the surprise messages coming in from Systems One so that System Three will not be alerted. This is the case when a notification message is ignored in System Two, when certain messages are tabooed (e.g. patient may only be harmed by a complication, not by an accident) or when scapegoat seeking is the prime objective. The requirements concerning format and contents of a description of an operational surprise in a System One that is reported to System Two depend on the design of all these attenuators and memory functions in System Two: the less well defined, the more System Three/Three-Star needs to be invoked. In Chapter 2 we concluded that "the anomaly *description* is a message conveyer between reporter and learning agency about an operational surprise: the less the agency is familiar with the normal operational practices, the higher the demands on completeness of descriptive data regarding the normal process constraints". The agency is a System Three-Star resource for System Three. From VSM we now know that demands on the completeness of a surprise description can be divided according to the variety absorption capacity by System Two and the additional capacity (to be) provided by System Three. We had already noted that more complete descriptions in messages from System One to System Two increases the threshold in the System One – Two synapse.

The above have been general points about information system design made clear by the new insights of VSM. The next three sections look in more detail at two projects and one further model from a VSM perspective.

7.6.1.1 Example 1: scrutinising Medical Device Reporting (2.2.6)

As pointed out above, in particular cases, System One signals may be passed on directly to a higher recursion management system. An extreme example is found in the Medical Device Reporting system set up on behalf of the US Food and Drug Agency (FDA). Users of medical devices are expected to report problems experienced directly to the national notification system (MedWatch) without having to put this message through intermediate, institutional transmission channels [CDRH 96]. The variety lost in this communication short-cut consists of the particulars of the recursion in which the signalling System One works, e.g. the health care institution. However, the purpose of the reporting system is confined to the device itself, its design, manufacture and use, and these particulars are not therefore so important. What is vital is the technological features of the specific device that set off the operational surprise in System One. These can be linked with problems reported through a device-focussed System Three-Star audit. So, in VSM terms we see the following in the USA. The nation-wide health care system is the SiF. Medical device problems experienced in health care institutions or by consumers are reported via the MedWatch notification programme to the Manufacturers and User Facility Device Experience (MAUDE) case base (System Two, including memory), see [CDRH 00]. The FDA functions as System Three. System Ones may be health care organisations as well as private persons. The System Three intervention is focussed on product market approval, not on running health care organisations. The MAUDE database and FDA regulation on market approval are System Two components at the SiF recursion.

The heart valve prosthesis [CEC 90] and Therac-25 [Leveson 93] cases also reveal that the anti-oscillatory function of System Two was (and needed to be) bypassed by data and messages transmitted through shared environment channels (iv), before appropriate higher

recursion Systems Three started to respond [CEC 90; Leveson 93]. In the case of the heart valves, patient organisations and critical doctors and lawyers had to join forces – through shared environments – for a forensic investigation of the manufacture of and communication about the suspected heart valve. Formal channels (ii) and/or (i) linking Systems One with the appropriate System Three, e.g. regulatory body for health care, were overloaded with deceptive data from the manufacturer. The series of fatal radiation accidents with the Therac-25 demonstrates among many other things the incapability of a competent regulatory body, i.e. the FDA, to identify the technical problems that caused the accidents despite many resourceful attempts. Once the users from different institutions joined forces, a break through was realised relatively fast. Both cases illustrate that a medical device problem monitoring system must be positioned at a recursion that is beyond the health care producing institutions, and certainly needs to be independent from individual manufacturers.

7.6.1.2 Example 2: diagnosing the ARA Knowledge Network

Of all projects described in Chapter 2, the ARA project resulted in the most complete set-up of organised learning from accidents at work, see section 2.3.1 and also Figure 2-4. In Figure 2-4, repeated below as Figure 7-5, the companies are the Systems One of a nationwide organisational learning system. System Three is the User Panel, which decides on incremental development of the ISA (expert) system. We are therefore dealing here with a management function limited to the management of an information system (ISA-PL), not to the management of the risks in the companies using the ISA system. The ISA system was embedded at the top level of each of the participating companies and deployed by the SHE department. The expert panel has a System Three-Star role and sporadically audits the data that is submitted to the raw data base that is part of the surprise memory designed in System Two. System Four is a combination of members of the user and expert panels who keep an eye on future developments and needs. A typical System One disturbance signal contains the message that the ISA resource provided to the Systems One is less than adequate.

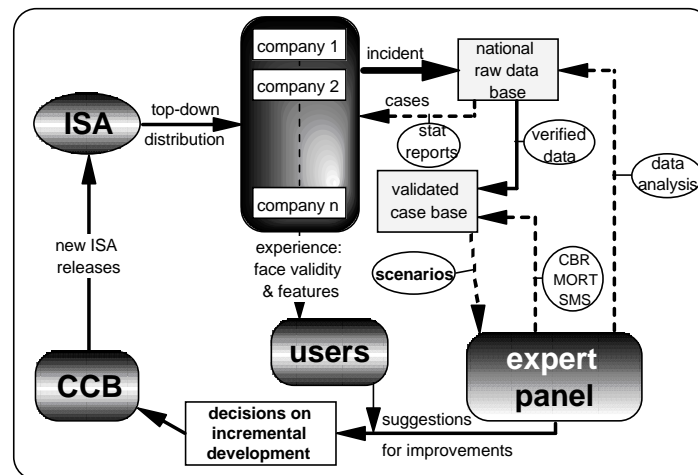


Figure 7-5: the ARA Knowledge Network at national level [see 2.3.1]

By designing a validated part of the memory function in System Two it is possible to filter out from the signalled disturbance all but the ones which are unknown. Such a validated case base contains lessons learned by a System One that can be retrieved by any System One without needing to involve System Three. So far so good! The main problems are twofold:

- 1) the Systems One did not see themselves as a nation-wide organisational learning system (SiF) and did not act as one. The nation-wide system only existed in the organisational landscape that forms the scene of System Four - System Three interaction (castles in the air).
- 2) the ISA-PL system produces generically stated analysis outcomes about operational surprises that occur at least one recursion down from the Systems One of the SiF, i.e. at the workplaces of the companies.

The generic formulation makes it hard for a company to interpret the case in its own system to recreate the variety to make the generic analysis live in their terms. The details in the cases did give context, but it was industry specific. The cases found in the validated case base in System Two may, therefore, only be useful for Systems One of the SiF that share environments, e.g. operating in the same branch of industry and hence having similar workplace operations or sharing the same senior management at a higher recursion. This undermines the desire to learn across industries in a truly nation-wide learning system. In addition, the ISA resource was not positioned adequately within the Systems One of the SiF.

As diagnosed above, the ISA system was implemented at the lower recursion as a System Two component, but it should have been positioned as a System Three-Star tool. The System Two use of ISA with hundreds of cases notified this way would most likely have degraded the System Three-Star sharpness as a result. However, the same ISA tool was judged to be effective when deployed as a System Three-Star device by an external auditor, like a labour inspector or a trade union representative from outside the System One concerned. Thus, the VSM-diagnosis of the ARA approach pinpoints some key issues regarding System Two design and ISA implementation within Systems One that were not acknowledged during the ARA project. These can be enlightening when designing learning systems organised by higher recursions in future.

7.6.1.3 Example 3: review of the (modified) NMMS Model

The NMMS-view deployed in Chapter 3 resulted in proposed modifications to the configuration of modules in a Near-Miss Management System, see Figure 7-6 below. The initial NMMS depicted in Figure 3-1 was *not* designed to be a System Two for the Systems One of a SiF. It should rather be seen as a performance measurement system in the accountability channel (ii). Based on System Three 's judgement of the Systems One near-miss performance, it was designed to invoke a System Three/ Three-Star intervention in step 2. The original NMMS was designed to identify systemic problems through computation of classification data in surprise notifications from lower recursions arrived at by means of a System Three-Star audit in step 3. Thus, step 1 of the initial NMMS was only loosely embedded in System Two. The initial notification requirement of "aiming at a complete, valid reporting of all near-miss situations detectable by employees" (see section 3.1.2) was already a System Three-Star-like support for System Three. Such a

'completeness' criterion denies the filtering function of System Two and threatens to overload System Three, and/or demotivates Systems One reporting if no immediate action is taken. The possible conflicts between System Two and System Three-Star have been discussed above (section 7.6.1).

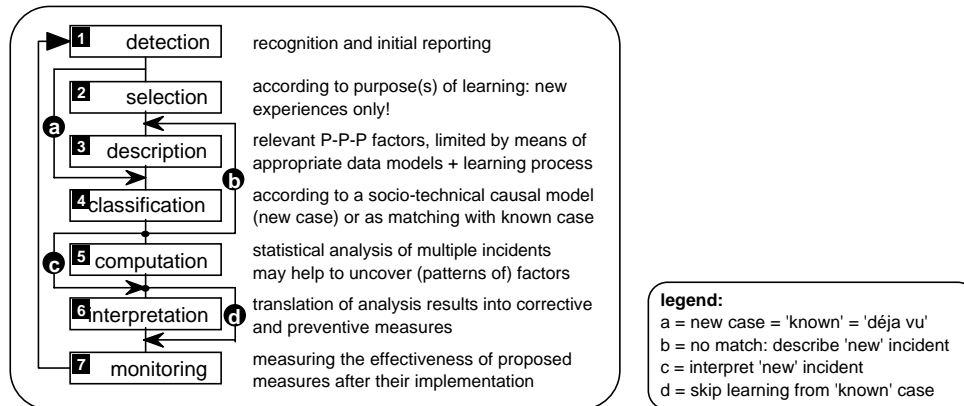


Figure 7-6: restructured modules for Near-Miss Management Systems [from Figure 3-5]

The modified NMMS model structures distinct steps of, and flows in a process of learning from operational surprises. The detection step **1** comes from a System One (department) at one recursion below that of the SiF (organisation) where the NMMS is situated. For the signalling System One, System Three is part of the department's senior management. The variety reduction in the notification message coming from System One and going into the NMMS would be twofold, once when going through System Two and again through channel (v) upwards. The NMMS might invoke the System Three of the SiF that, in turn, needs to perform a System Three-Star audit in order to generate adequate variety to solve problems identified one recursion down. In this case, the NMMS will become a resource-demanding anti-oscillatory mechanism. A basic improvement can be realised by positioning the NMMS at the lower recursion of the Systems One concerned. The classification data of cases assessed in step **4** and the computation function in step **5** may then be transferred to System Two of the next recursion upwards. Thus, for a 'new' case, the interpretation step **6** is at the NMMS recursion of the System One surprise, where it should be. Steps **3** and **4** ('new' case) are typically System Three-Star activities.

Where computation (step **5**) of the classified case data results in the pinpointing of systemic problem areas (repeat cases), these must be looked into, e.g. by a System Three-Star audit from the higher recursion. This looks at the System One of the lower recursion, so that the higher recursion System Three can work out a metasystemic intervention (step **6**) aimed at one or more recursions downwards. This division of the NMMS into two parts, dealing with new and recurrent cases, can be recognised from section 7.3 and from Chapter 6. The monitoring module (step **7**) is a part of System Two at the lower recursion whose linkages with System Two of the higher recursion must be designed carefully.

Thus, the VSM diagnosis of the (modified) NMMS leads to the conclusion that the NMMS should be positioned at the recursion of the Systems One where the incidents occur. However, step **5** and possibly step **7** of such an NMMS are allocated at the next higher recursion.

7.6.2 Specification of a Surprise Monitoring System for Organisational Learning

This chapter has shown that the Viable System Model appears to be a powerful tool to diagnose systems that are meant to enable learning from unwanted deviations. Organisational Learning theory added the notion of organisational memory, which is vital to retain and reuse the lessons learned. In VSM such a memory of experiences with undesired System One deviations can only fit into System Two, the anti-oscillatory component that relieves System Three from information overload coming from operational surprise signals. However, as pointed out in section 7.4, System Three also contains a memory function. This is focussed on the intended functioning of the Systems One. It contains espoused theories for Systems One about correct ways of carrying out tasks and established exceptions to them and it sets their governing variables. Thus, System Three memory is normative for Systems One, whereas System Two memory is descriptive. The anti-oscillatory System Two in VSM is linked to the inside-and-now management System Three. Ultimately, System Three will adapt its normative memory using also the insights from the System Two memory. This double-loop learning is based on the (System Three-Star) learning from incidents and propagates that learning through the policy channel (i), the resource bargaining channel (ii), and/or by modifying filters in System Two. However, a System Three memory is not sufficient. For (near-) accidents occurring in a System One, the System Two memory is crucial: without a memory, learning by the organisation is just not possible. This insight is a cornerstone for organised learning from small-scale incidents.

Based on all these insights, it is possible to use the VSM to specify the design requirements for operational surprise control systems. First, one must identify the Systems One where the surprises will eventually occur. This is usually an organisational (production) unit, but it may also be an individual functioning in a specific context, e.g. a health professional or a patient applying a medical device. Then we need to pinpoint *a* or *the* relevant System Three, because, without a System Three, organisational learning from operational deviations in Systems One is simply not possible. There is nobody to coordinate or take organisational action. The type of lessons to be learned from the System One surprise must then be made explicit, because this has a major impact on the choice of perspective of the System One recursions.

7.6.2.1 Example 1 revisited: specifying Requirements for a viable medical Device Reporting System

The medical device problem may clarify this point. In a System One, e.g. a hospital ward, a functional device failure, such as a malfunction of a defibrillator, disrupts the intentional care delivery process. System One management will try to restore operational homeostasis in any case, e.g. by using a back-up device to save the patient and end the operational emergency mode. Several Systems Three can be identified. Obviously, the health care institution that embeds the System One must realise health care services and, therefore, System Three is focussed on the functioning of Systems One. If System Three is the hospital department manager that supervises the distinct organisational units within that

department, the type of lessons that can be learned concern training, inspection and maybe change in purchasing policy. These lessons aim at improving operational control by Systems One in case of failure of the specific kind of device. This System Three manager is not embedded in a recursion of the manufacturing system where he/she can bargain for or require improvement in the design or manufacture of the device. Nor has this manager the means to recognise a recurring device problem that might be seen early at global level. The manufacturing and marketing of medical devices is no goal of the individual health care institution. The System Three that can influence design, manufacture and selling of a medical device is outside the health care provider's organisational system. Systemic solutions to operational device problems include the design of the device and market approval regimes.

Operational problems with devices may indicate a need to modify the device or to withdraw the product from the market, especially in the case of critical devices (e.g. classes II + III in European and FDA's classification schemes [EC 93, FDA 92]). The actual severity of problems with a particular device could be assessed if the incidence rate were known, but individual institutions as Systems One usually have low incidence rates. Combining incidence data at a much larger scale, e.g. world-wide, amplifies the signals coming from device-related System One surprises. So, where might be a System Three that can resolve device-related problems systemically?

An evident answer would be: with the device manufacturer! The manufacturer can decide to alert the user, or to modify, if not withdraw, a device in response to signals coming in through appropriate Systems Two and System Four. The manufacturer can influence Systems One by means of their products through a resource bargaining channel (ii) as well as through shared environments (channel (iv)). If this were done device by device, the users as Systems One would need to interact with a System Two and a System Three for each device manufacturer. Furthermore, each System Three designs filters and attenuators for its System Two–Systems One anti-oscillatory linkages and channel (i) and (ii) communications that differ from manufacturer to manufacturer and from product to product. The BS-cc heart valve case [CEC 90] illustrates such manufacturer-designed and owned System Two features. In this case it collected numerous, and revealing statistics on valve failures but responded with "Dear doctor" letters via channel (ii) to calm down professionals. This gave it a notorious reputation in the BS-cc heart valve case. This System Three information system design choice relies heavily on the integrity of individual companies, which is not a priori justified (as Pfizer showed in the BS-cc case [CEC 90], see also the proven deception of the public by the tobacco industries [Glantz 98]).

Another System Three option might be a governmental regulator, for example the FDA. Such a regulating body imposes (pre-) market approval (PMA) rules for medical devices. A manufacturer must comply with these rules in order to stay in business. The System Two linkage between this System Three and the Systems One that experience problems with the application of devices might be designed as a surprise notification system that is accessible by any System One that uses a medical device. Again System Three designs filters and attenuators for this System Two and its channels (i) and (ii), but these will – most likely – be independent from device producing manufacturers and may be questioned at a higher recursion, e.g. by a parliament. An intervention by this System Three may directly affect Systems One by the command channel (i), e.g. a "STOP using

this device" ruling, or more likely through the accountability channel (ii), as well as through the shared environment (channel (iv)).

The design of the memory function in System Two needs to be considered carefully from the point of view of lessons that can be learned this way (change PMA requirements, conditional withdrawal of a market approval, and/or alert users) as do the System Three trigger mechanisms, the System One-System Two attenuator and the modulation of the System One-System Two synapse. Manufacturers have interest in System Two designs that have strong anti-oscillatory characteristics suppressing information to System Three on problems. In this way they can sell products as much as possible. On the other hand, a regulatory body like the FDA as System Three has a great interest in a well-designed System Two for anti-oscillatory control of devices-related disturbances in Systems One that, for instance, has effective 'early warning' features without creating too many false alarms, but still sufficiently relieves – by filtering out problems related to the user or health care institution – System Three in its interactions with users (Systems One) via channels (ii), (i) and (vi).

Thus, we may conclude that medical devices reporting systems that rely on adequate response by individual manufacturers will be less effective from the point of view of Systems One that experience undesired operational surprises than one that is 'run' by a System Three that – in another recursion perspective – is also a System Three of the individual manufacturers concerned. An optimal design of a medical device reporting system would be one in which an independent regulatory body (or agency on behalf of it) is the System Three and designs the System Two that is directly accessible to (numerous) Systems Ones. The design of System Two must include a memory on problems and solutions applied that supports System One managers to recover from an operational surprise without needing to invoke a System Three intervention. In addition (not instead) signals coming in from Systems One can be passed on as a service to a manufacturer's System Two.

The US MDR system in the light of the VSM requirements

If we evaluate the current (example of the) Medical Device Reporting programme in the USA using the specifications derived above, we make the following observations. System Three (FDA) is well placed. Systems One can submit surprise signals via MedWatch directly through Internet. System Two is Less Than Adequate: it very difficult for users to retrieve useful data from MAUDE (organisational memory), and it is also unclear how and when System Three is invoked by System Two. The System One – System Two signal template [CDRH 96] is not suited for the available types of System Three interventions. Therefore, System Three may need to initiate a System Three-Star audit too often or fail to respond effectively.

7.7 Discussion

The questions posed in the introduction of this chapter (p. 120) can now be answered.

1. Can descriptions and problems from previous chapters be accommodated in the VSM view, i.e. can the VSM diagnosis of a system for incident monitoring and learning categorise and further illuminate issues arrived at the hard way in the projects discussed in previous chapters?

The basic answer is affirmative. The main message from the VSM is that incident monitoring and handling is an essential feature in a viable system, and monitoring for learning is part of the anti-oscillatory System Two. System Two must be designed carefully and must include data filters for messages coming from the operational level about disruptions. The variety attenuation in the message by System Two must always be huge. Any attempt at an extensive or complete incident reporting scheme in System Two could contribute to a System Two 'breakdown' rather than to a flying start for a System Three-Star audit (see also below at questions 3 and 4). Table 7-1 and Table 7-2 above demonstrate that the concepts of organisational learning and of SINS fit with the VSM. Problems experienced with the implementation of SINS can also be explained in VSM terms. The more precise allocation of SINS-functions by applying the VSM for the design of SINS could have prevented the development of the extended notification form in SINS-1, because of the essential notion of variety attenuation in System Two. The extended form in SINS-1 was conceived as a System Three-Star tool, but was wrongly placed in a System Two setting: in VSM terms this does not make any sense; the implementation project resulted in the same conclusion the hard way.

2. Does the VSM organise the conclusions drawn so far?

Once again the answer is positive. VSM helps to allocate organisational learning issues regarding learning from operational surprises to topics that link to distinct Systems (One to Five), transmission channels and recursions. The limitations of the (2-dimensional) OL and SINS models are resolved in the VSM. The diagnostic features of the VSM also distinguish between transmission channels and oscillation control functions. The following example demonstrates the organising features of the VSM.

As discussed in Chapter 6, it was conceived that all notification reports submitted to SINS-1 would be sent to the location manager on a daily basis. In VSM-terms we may say that in this way a System Two function was light-heartedly transferred to System Three (i.e. the location manager of CAPU, see Figure 7-2), because the location manager was supposed to select from the daily reports notification signals that required System Three intervention. From the VSM perspective, however, we expect that System Two should only generate signals to senior management about operational surprises when System Three involvement is needed for counteracting oscillatory interactions between the Systems One. The VSM model does provide a direct link between the regulatory centre (System Two) of a System One operation and its management unit. Therefore, it does make sense to convey notification messages from a System One operation submitted to SINS-1 to the respective System One managers daily, but not to the location manager at System Three. The System One managers may have to account for their operational surprises to System Three through channel (ii), but will do that in another way, not with full notifications. System Three should only be bothered by System Two about operational instability problems when notification criteria are met. The incidence trigger in SINS-2 is a design attribute of System Two for invoking System Three intervention.

3. What does VSM tell us about the work and findings described in the preceding chapters, especially regarding the SINS concept. Does the VSM diagnosis reveal

issues of design and implementation of a SINS not identified yet in the previous work?

The interconnectivity as well as the spread of learning functions over SiF Systems One to Five and over the different channels (i) to (vi) relevant to organisational learning is a strong feature of VSM. Learning from operational surprises at one recursion level is moderated by System Three either after being triggered by System Two, or only after using a System Three-Star audit. Alternatively, System Three may receive new insights from System Three-Four-Five interactions deploying new tactical and strategic information from outside the SiF, or even outside the organisation. This may occur after intervention through channel (i) or a System Three-Star audit from a higher recursion and/or another dimension. Learning by System Three directly from System One signals and associated messages using System Three-Star audits is less effective if conditions for surprise recurrence remain, and may obstruct other System Three activities. Thus, the VSM model clearly locates the part of the memory function of the SiF dealing with unwanted disturbances at System Two as far as this part of the organisation's memory contains previous, i.e. known experiences. The anti-oscillatory function of System Two is enhanced by such a memory if a System One gets access to this memory and finds a solution for an operational surprise that it is facing. It can cope without the need to involve System Three for intervention. This insight confirms the key role of SINS-2 and the case base in the SINS concept. However, the VSM diagnosis also points to the importance of the position and role of senior management (System Three-Four-Five) which is still only implicit in the SINS concept so far. The VSM shows that the systemic solution of System One problems for System Three requires requisite variety that can be achieved through a System Three-Star audit, or by a RT that already embodies the requisite variety. The latter might be more cost-effective than the former when it comes to the small-scale incidents, since a full System Three-Star audit takes time and disrupts work.

The Review Team (RT) combined tasks of several VSM functions: a System Three-Star learning agency when it reviewed new cases; and a System Three task to maintain System Two memory when storing validated cases into the SINS-2 case base memory. By enhancing the memory function of SINS-2, the RT may have to be invoked less frequently. From here, the conclusion can be – as a new insight – that the memory of System Two can act as an anti-oscillatory component in organisational (single loop) learning, since operational surprises that have been experienced before can be damped considerably if the SINS-2 case base is well designed and loaded. Needless to say that System Two must be designed to prevent too low as well as too strong damping of the oscillatory behaviour of Systems One.

Maybe the best illustration of clarity produced by VSM diagnosis is that a System One usually exists in different dimensions. For example, a person may be at the same time a member of his family, active in a trade union or the local golf club, and a System One within his organisational unit of a company. This notion of multi-dimensionality of viable systems unifies the – at first sight – separate regimes for notification of operational surprises within organisations and for notification of product-related problems: see section 7.6.2.1 above. Implications for SINS are that classes and routing of incoming surprise signals through SINS-1 needs to be defined more precisely in order to improve overall learning effectiveness. Furthermore, the design of the System Two memory of operational

surprises may be optimised by sharing parts of this memory with other recursions in the same and/or in different dimensions. For example, a System One signal concerning a critical device failure may be stored at the same time in an organisation's data base and in a nation-wide device reporting system so that other users elsewhere can be informed directly or share operational surprise control experiences through this shared memory.⁴⁴

Striking is also the insight provided by the VSM into the meaning and positioning of 'audits' compared to what is usually thought. The VSM tells us that many so-called 'audits' are actually System Two monitoring schemata; this includes methods like Tripod- δ [Roggeveen 99] and Tripod- β [Tripod 99] that for this purpose requires multiple accidents. Tripod-Delta, the 'audit' tool is based on a large set of questions covering the General Failure Type (GFT) categories. This set of about 1200 questions is generated for and tailored to the company that is being audited. For the audit a sub-set of questions is used and distributed among the workers in the organisational units. By developing the questions specifically for the company, taking the work processes closely into account, and by distribution of the questionnaires to all workers, the Tripod-Delta audit comes rather close a System Three-Star audit and absorbs the variety in the responses by workers more meaningfully than can be realised with audits tools that are based on generic question sets. Many responses to the questions that are rooted in the work processes justify a high confidence in the resulting GFT-profile. But the profile itself is a strong variety attenuator. For a subsequent audit some years later, another subset from the same question set will be used. This and changing operational conditions may make it difficult to interpret the new GFT profile. The accident investigation tool Tripod-Beta is complementary to the audit tool. The latent failures found in one-off Tripod- β analysis are evaluated for their association with GFT categories (now called BRFs: basic risk factors). The resulting BRF-profiles of several accidents need to be accumulated and then compared with the already determined GFT-profile: in due course, the aggregated BRF-profile should match the GFT-profile. Therefore, such methods do not increase managerial variety (in fact they are variety attenuators, whereas a System Three-Star audit amplifies System Three variety temporarily).

4. Does VSM provide means to translate the SINS concept as implemented in a hospital context into another domain of organisational activity?

For an organised learning system like SINS, the SiF needs to be made explicit, components for monitoring operational instabilities by System Two must be designed and feedback routes be identified. A many to one model like VSM helps to interpret features from one SINS implementation into VSM terms and then project the VSM version into another organisational reality. Such a translation procedure helps to maintain the integrity of the transformed SINS as implementation flaws are scrutinised by applying the intermediate VSM-mode in this process.

⁴⁴ On April 13, 2000, the Academic Hospital Maastricht reported about a software flaw in a blood transfusion control system that allowed the entry of a blood type that differs from the first time entry for a patient. Thus, an unnoticed data entry error for a blood request resulted in the death of a patient. When the hospital reported this accident to the health inspectorate, it learned that practically the same had happened several months earlier in another teaching hospital. Although the software system is used by about 50% of all hospitals, this information and the software modification solution applied after the first accident was not disseminated with any priority by or under supervision of the health authorities.

The VSM helps to sort out appropriate communication channels and the blockages in them. Actual obstacles for organisational learning may be due to failing or missing system synapses, inappropriate use of communication channels (e.g. propagating a lesson learned through the command channel (i) instead of channels (ii) and (v)). System Two may be overcompensating, i.e. damping too much and thus also killing learning capabilities of the system, for instance by means of a tacit 'safety culture' which tells individual members that hazards are always under control, so that safety problems do not get raised. The implicit message is that individuals or teams may be punished if signals become apparent that risks are not controlled. System Four-Three learning may be blocked because of a weakly developed System Four. The next higher recursion System Two may not transfer incident signals from the SiF. An operational problem message to the local regulatory centre coming from the System One operation may be limited to a one bit 'something happened' signal that requires full attention by System Three and a System Three-Star audit. Two recursions down, the System Two elements like 'ethos' or 'culture' are part of the individual's way of thinking and a powerful blockage or stimulus to learning. In other words, the VSM provides diagnostic features to identify obstacles in productive learning as well as to define essential features and mechanisms for learning from operational disruptions by a viable system.

The translation of the SINS concept from a hospital environment into another domain may be facilitated by using the VSM model as a transition scheme. VSM System One to Five entities and their connections can be identified more easily than with the SINS model which may be too generic for this purpose. This is particularly vital in complex systems, such as hospitals, but should be useful also in less complex ones.

The 'vertical' translation of SINS is also an issue within a particular domain of activity. When a SINS is to be transformed for use one recursion up or down, the distinct SINS components will have to be re-designed as well. Again VSM can help draw the parallels. For instance, the application of the SINS concept down one level, for anaesthesia during surgical procedures, might lead to a simple push-button notification system SINS-1 coupled with an automatic anaesthesia log. In the case of observation of a crisis in the patient's physiological system by the anaesthetist, he or she can mark the beginning of the observed surprise by a simple activation of a push button and again when the crisis is brought under control. After the operation, the anaesthetist may invoke a System Three-Star audit and a System Three-Four exploration in order to identify a solution to the problem analysed. A simple set of three push buttons linked to the automated anaesthesia log might be an adequate SINS-1 translation for this specific downward recursion.

Hence, we conclude here that the VSM provides means to translate the SINS concept into other domains of application or for other recursion levels.

This thesis has enhanced the VSM as diagnostic tool by adding the insights from organisational learning principles. The necessity of an organisational memory enriches the understanding of the VSM. In his writings Beer has only addressed the metasystemic 'organisational landscape' [Beer 79] memory up to now. However, System Three must have a memory about how and why to run System One operations. Similarly, the anti-oscillatory System Two is best suited to accommodate the real-time memory needed to retain experiences and operational surprises in Systems One and associated lessons about controlling them so that System Three does not have to be invoked too much.

The insights gained can be summarised in recipes for diagnosing or designing a System for Organisational Learning from operational Surprises, see text box below.

Recipe for Designing a System for Learning from operational Surprises using VSM

- select a System One;
- identify a System Three in the relevant dimension [consider the necessity to accommodate the design of the System for Learning from operational Surprises for different dimensions]
- define the type of lessons to be learned by System Three [for each dimension considered];
- design the System Two memory
- design System Two filters and attenuators
- assess opportunities for (sporadic) System Three-Star interventions

Recipe for Diagnosing a System for Learning from operational Surprises using VSM

- look for relevant Systems One;
- look for memory in System Two;
- identify the System Three that links to System Two and Systems One (verify the appropriate dimension for System One-Two-Three)
- assess the potential influence of System Three on Systems One;
- if memory in System Two is inadequate then look for memory on surprises in System Three;
- evaluate scope and incidence of System Three-Star audits invoked by System Three
- analyse and evaluate the allocation of "operational surprise" messages in the vertical channels (v), (ii) and (i).

7.8 Conclusion

The VSM has demonstrated that it is very useful for specification of systems for organised learning from small-scale operational surprises. Compared with the VSM, Argyris' model of organisational learning is less coherent as it does not provide adequate variety, but it does facilitate communicating and 'selling' the SINS message of organisational learning.

For the actual translation of the SINS-concept into a specific organisational context, the SINS provider still needs to apply a many-one transformation scheme like the VSM. However, it is not appropriate to sell VSM instead of SINS, because then the buyer must understand the VSM language before the SINS message can be understood. As the reader of this chapter will have experienced, this is no simple task.

SINS is conceived as a generic model that represents and configures functions that are essential for productive learning from unwanted operational surprises. Argyris's organisational learning model is reflected in the SINS concept as functions and components in the learning loops: notification, agency, single ('SINS-2') and double loop learning modes ('SINS-3'), lessons to learn and feedback into the working processes concerned. The SINS implementation project inside a specialised surgical unit of a hospital was at a recursion level where organisational entities form the Systems One. The SINS concept can also apply to other recursions, but the realisation of the SINS elements will differ. At a lower recursion, simple push buttons may replace electronic notification forms (System Two), but in-depth analysis might need to become domain specific (System Four exploration) rather than organisational (System Three-Star audits).

At a higher recursion level, the purpose of learning and the recursion dimension is crucial for the positioning of the learning agency. For instance, a complaint submission system for patients may be anchored at hospital level, where the members of the learning agency (System Three-Star function) come from the clinical departments. However, a medical device reporting system needs a SINS-1 that is accessible from within all health care institutions, and a learning agency that can assess operational risks related to technological artefacts while being independent from suppliers. As discussed in 7.6.2.1, a possible System Three for such a System Three-Star may be a consumer organisation or a political authority. SINS-2 is then the memory of lessons learned concerning the intentional or unintentional use of devices.

Although the SINS concept is generic, the actual contents of SINS-elements are specific for the application domain and the recursion of implementation in an organisational system. SINS only works for organisational systems. The VSM provides the meta-language that is necessary to link and specify generic SINS-functions to organisational elements and interactions of the system where SINS is to be implemented.

Conclusions and Future Work Issues

The main objective of this thesis was to find ways and means to learn effectively from daily – small-scale – accidents and near-miss situations. For this purpose, data on accidents and other unwanted operational surprises needed to be collected. In the early projects, the emphasis was laid on the capturing of notification data in order to create substantial data sets that could be analysed statistically. The IDA data-recording tool with its electronic forms was developed and deployed to lighten the administrative burden of filling in forms and subsequently processing the forms. Data acquisition was defined as the problem and monopolised the thinking at the expense of the prime question of what the surprise data are meant for and how they will be processed after they have been collected.

The early projects showed that large sets of data on incident occurrences are not suited for learning about systemic causal factors for reasons of data model limitations and the time frame needed for collecting data of a sufficient number of cases. The use of causal models describing harm process scenarios, embedded in the data models for incident capturing, still required too many recorded surprises to be usable in a reasonable time frame. Furthermore, the proximal causal factors, that directly contribute to damage, do not constitute the lessons that are most useful for an organisation to learn. An organisation's interest in learning from incidents lies in factors earlier in the accident process, whose control leads to improved control of operational risks in intended processes and, consequently, fewer losses. The relevance of time-sequenced operational conditions in relation to safety and control barriers was also elaborated in several projects in health care institutions. The serious limitations of learning from incidents by means of statistical analysis led to the concept of deploying knowledge about systemic causal factors underlying accidents as the basis for incident data capturing. Such knowledge was available in the generic Management Oversight and Risk Tree (MORT) diagram that consists of organisational or 'safety management system' (SMS) functions relevant to assure loss prevention by adequate operational control. Application of this model in surprise data recording using the ISA tool demonstrated the feasibility of one-off analysis of small-scale incidents. The ISA tool was transformed into an expert system to do this. This led to the discovery that thorough situational knowledge is demanded of the user in order to realise time-efficiency and convert generic recommendations into concrete ones. The type of lessons to be learned concern the process shaping conditions that the organisation can control. The cumulation of data from different cases provides an

additional opportunity to identify (latent) problems related to repeated incidents or to specific SMS functions.

This approach was not sufficient for the recording and analysis of product-related incidents where many of the systemic problems are beyond the control of product users. The heart valve prosthesis case showed that learning is essential from the very first case when a safety-critical product is involved in an operational surprise and opportunities to regain control without irreversible damage are practically absent. The type of lessons to be learned from product-related incidents focus on risk control opportunities in earlier phases of the product's life cycle, all the way back to the stage of conceptualisation. The wide spectrum of safety-critical products and their intended and/or anticipated use requires a post-market surveillance system to identify problems that relate to a particular product where first-time learning may be necessary: cumulation of notification data from many user organisations helps. So, surprises in the operational process that involve product failure require two sorts of learning. The first recovers the process despite the product failure. The second must feed the surprise back up the chain to the product designer, manufacturer and purchaser. The question which remains open is how to design and co-ordinate these two learning processes.

The issue of organised learning from incidents was first addressed explicitly in the ARA project where different companies deployed the same MORT-based knowledge model for accident recording and analysis. The set-up of a nation-wide learning system was drafted, but could not be launched. However, it was concluded that learning beyond individual companies might be possible if companies were in the same line of business, i.e. have comparable operational conditions. The acknowledgement of the relevance of 'context' for interpretation of incident notification messages as well as for translation of generic recommendations into site-specific measures was a major conclusion coming out of this macro-level project. Although the 'context' issue seems trivial it is not. The attempt in the Near-Miss Management System (NMMS) model [Schaaf 91b] to capture context by means of comprehensive incident descriptions does not solve the keyhole-view problem that a model of reality represents and, in addition, discourages surprise observers from submitting a notification, because of the recording burden.

The projects gave rise to legitimate doubts about the feasibility of this top-down type of approach. Large numbers of recorded cases are needed and, consequently, many harm processes have to occur before data analysis can start to become meaningful. Factors found through such data analysis cannot be interpreted into operational meaning relevant to the organisations being monitored, unless a common background is shared. This may be found, for instance, where different companies run similar technological processes, or belong to the same mother corporation. In the same way, technological systems from different product brands may share life cycles with similar risk control characteristics.

The principles of organisational learning as described by Argyris pinpoint the need to actively organise the learning processes of an organisation that wants to learn. Learning does not take place unless the learning process loop is closed. An organisation may learn through the people who work for it, but only if one or more of the members are assigned to be a member of a learning agency to learn for it. Single- and double-loop learning starts with 'zero learning' [Bateson 72] when an individual notices an unexpected operational condition ('surprise!') and may or may not inform the learning agency. The prevailing

culture within the organisation has a strong impact on 'zero learning'. It may inhibit the recognition of a surprise condition, and subsequent notification. Once reported, the organisational unit where the surprise occurs may be able to resolve the surprise condition within its regulatory limits (single loop learning). If it cannot, higher management must decide about changing objectives, norms and values under which the unit functions (double-loop learning). The concept of 'theory of action' [Argyris 96] provides insight into the limited impact of the explicit 'espoused' theories about how things should be done compared to the tacit 'theories-in-use' about how things are really done. In defensive organisations the theory-in-use is in conflict with the prevailing 'espoused' theory and there is no way to put this difference on the learning agenda. This concept of 'theory of action' emphasises the limitations of top-down approaches when setting up organisational learning, as only the 'espoused' theories can be propagated.

The Systemic Incident Notification System (SINS)-concept is a synthesis of lessons learned through the projects described. In essence, it fits within Argyris' model of organisational learning with key positions for the learning agency (Review Team = RT) and the organisational memory (SINS-2). The members of the Review Team were also members of the organisational units where the surprises occurred. Thus, notifications could be kept relatively simple, because the Review Team was able to interpret the message within its operational settings. Resources for the RT were limited in the hospital case as its members already had a high workload. This is always likely to be the case. For this reason, the organisational memory is of particular importance. Lessons, once learned, need to be stored and made available when appropriate. The more lessons learned, the simpler the notification messages (using SINS-1) can be if appropriate cases in memory are retrieved effectively and efficiently. The SINS-concept also provides a key mechanism for cost-effective organisational learning from small-scale incidents. The known surprises (and countermeasures) stored in memory are classified by the learning agency in risk control categories. Classes vary between "unacceptable: immediate one-off in-depth investigation and system solution required" and "ignore: no action needed". An incidence-trigger value *IncT* can be assigned to each validated event in memory. A new notification is then matched with cases in memory: when the event frequency reaches the *IncT* value, the Review Team as learning agency is warned and can decide to invoke a double-loop learning process to find a systemic solution. To support this, the RT may deploy the expert system on organisational causal factors (SINS-3). SINS-1 and the Review Team were implemented in the most recent of projects, whilst SINS-3 had already been tested in previous projects. The cumulation of data in SINS includes the recurrence of a particular type of incident (linked to *IncT*), as well as the causal codes assigned by the Review Team to systemic conditions in validated cases. This results in root cause profiles similar to those in PRISMA [Vuuren 98] and TRIPOD [Groeneweg 98]. The concept of the incidence trigger *IncT* allows process management to direct scarce investigative resources to those incidents that matter most.

The initial notification can be kept simple when a basic process model is applied in the data model: when a surprise is linked to a process stage and the learning agency consists of members from the surprise-generating units, the agency only needs half a word to understand the surprise message. Thus, a basic, sequential business process model provides an effective descriptor of a surprise notification, also in situations where different process phases occur simultaneously in one geographical area. Systemic

solutions to the surprise may be more feasible when looked for and implemented in a preceding process phase. The last project demonstrated this in several cases.

The realisation of organisational memory (SINS-2) within SINS was explored to some extent. At micro-level, such as a production unit, it may be realised using Case-Based Reasoning (CBR) technology. The key issues related to this memory function are twofold. First, cases in memory must have been validated within an identifiable context and assessed for criticality, e.g. by means of the setting of *IncT*. They must also be given a generic classification to some extent in order to match incoming notified cases. Development of this classification is an area for further work. Second, an incoming notification message must be compared with cases in memory. This is always possible in CBR, e.g. using 'nearest-neighbour' testing. In order to make this effective, however, dynamic case matching during data entry is necessary in order to shorten the notification procedure in the case of a repeating, known surprise. Thus, the interaction between the reporter using SINS-1 and the organisational memory SINS-2 must be designed bi-directionally. Early case matching shortens the entry of notification data and a retrieved case includes lessons learned in similar occasions in the past. The different technological principles for CBR systems [Kolodner 93, Bareiss 89] for building a cost-effective system for learning from incidents have not yet been assessed. The dynamic interaction between SINS-1 and SINS-2 is, therefore, also a key area of future work.

Only when an incoming surprise message cannot be matched with a case in memory, or when a match also activates the incidence trigger (*IncT*), does the learning agency need to be invoked. The SINS-3 component then may also be invoked. It is envisaged as an expert system containing knowledge on systemic causal factors to support the learning agency in a double-loop learning process. It can be replaced by experts in the methodological domain, if they are available, as may be done when an in-depth investigation is initiated from a higher system level. As basis for SINS-3, the ARA-application (ISA) is available. A complete implementation of the MORT diagram in such an expert system and the validation of the diagnostic module, however, form another potential topic for future work.

The SINS-concept does not accommodate the concept of 'deuterolearning' [Bateson 72] or 'learning how to learn', which is essential for (organisational) learning [Argyris 96]. Argyris' model of organisational learning (OL) and the SINS-concept are both unclear about the role of management in organisational learning processes. However, in the field trial of SINS-1 plus Review Team, it became clear that a learning agency needs a coordinator or learning agent who prepares meetings of the agency. Such a learning agent is a critical factor in an organisational learning system. Another critical factor is the propagation of lessons learned by the learning agency for the organisation. These lessons must go back to the organisational unit(s) concerned along routes that the OL-model depicts as single- or double closed loops. The manager is important in this respect, but the communication channels to unit management are at least equally important in single-loop learning. The allocation of tasks and contents regarding the propagation of lessons learned was only explored incidentally in this thesis. How these roles should be optimally fulfilled is another area for future work.

8.1 VSM Diagnosis

As the SINS-concept was only partially tested in practice, and the implementation was subject to inhibitors of various kinds, I felt it valuable to diagnose the problems in using the SINS-concept by means of Beer's Viable System Model (VSM). The VSM is powerful, but it is complex and hence hard to use directly as a meta-language with organisations that want to improve their organisational learning system(s). However, the VSM can be deployed to diagnose the organisational learning processes in a viable organisational system. The recursive features of the VSM not only call for clear identification of the system-in-focus (SiF) at the recursion that is to be diagnosed, but also for the identification of the 'hosting' system at the next recursion upwards that embeds the SiF. Typical for VSM and unlike organisational charts, a particular SiF at a specific recursion is embedded in many different systems at the next recursion upwards. Hence, an operational disturbance in the SiF may have to be passed on to different organisational learning systems. This insight, which is an inherent feature of the VSM, finally reconciled the apparently distinct systems for organised learning from operational disturbances and from product-related user problems.

The VSM diagnosis of the SINS-concept combined with Argyris' principles of Organisational Learning, the modified NMMS model and the practical attempts in previous projects to organise learning from small-scale incidents, resulted in a recipe for designing a System for Organisational Learning from Operational Surprises (SOLOS) and for diagnosing problems in an existing SOLOS.

The VSM enables a more precise description of some of the future work issues. The first issue concerns organisational memory to capture and retain *unwanted* operational experiences, like accidents, near-miss situations and other unplanned process distortions in order to reuse lessons learned in the past or to initiate double-loop learning. Such a memory must link to operational entities where the surprises occur and be accessible for retrieval of earlier surprise experiences in a way that is useful for these operational entities. Furthermore, it should alert the operational management of the surprise when appropriate, e.g. when there is not a similar surprise experience to recall or when some intervention trigger is activated. This organisational memory must be assigned, designed and maintained.

The most important consideration for locating such an organisational memory is the positioning of the management which has the power (and knowledge) to make decisions in the case of either a new surprise or the need for double-loop intervention. Such a management always exists within the particular organisational system that controls the daily processes in which operational surprises occur, but relevant management may also reside far outside, depending on the span of control that is needed to solve a systemic problem underlying the surprise experienced.

These different *dimensions* of organisational learning systems linked to a particular (work) process are critical for the effectiveness or failure of organisational memory. *Case-Based Reasoning* could be applied to build an organisational memory as part of a particular organisational system for learning from unwanted operational surprises. Both these are topics for further work.

The *variety* contained in the *messages* in a system for organisational learning from unwanted operational surprises needs to be modulated and transformed while the message is being conveyed between organisational memory and the surprise-generating (organisational) units. What data in the message can be left out and what needs to be added or compensated between a message generator (e.g. a surprise notifier) and a message recipient (e.g. an organisational memory containing earlier cases and lessons learned), varies. For instance, the inherent loss of situational context in a surprise notification must be compensated before the receiver starts interpretation of the message. This is less difficult when the recipient is knowledgeable about the prevailing operational situations. For organisational learning from small-scale incidents, it is crucial to minimise variety in notification messages and compensate efficiently for this loss of variety⁴⁵, e.g. by means of a carefully assigned learning agency. Both issues need elaboration in future work.

Lessons learned in one organisational unit may be valuable for other units, or comparable entities in other organisations (*horizontally*, e.g. between all cardiac surgery units in The Netherlands), but requirements for message contents are far from obvious. The *coding* of the relevant situational *context* of the surprise(s) that resulted in the learning of a particular lesson needs closer attention in order to avoid trying to reuse a lesson in an unsuitable situation. This issue will be crucial in the design of the organisational memory as discussed above, and the projects up to now only give initial indicators of how this must be done.

Besides the reuse and sharing of lessons between comparable organisational units, there is also the transmission of surprise data from one organisational level upwards within and beyond the organisation, e.g. to corporate, national or international level. These surprise data conveyed should also be fit for *purpose*. OL purposes might be high-level control adjustment of lower level operations, e.g. as demonstrated in Chile in terms of the national economy [Beer 81], or to achieve learning and dissemination of lessons faster than is possible horizontally (e.g. about product-related hazards where manufacturers and regulators need to be reached). Usually, the purpose of upstream flows of surprise data is just to generate performance indicators for use in setting priorities in policy and resource allocation. For each of these purposes, the design of data *filters* in order to put messages across rather than just data deserves further exploration.

Furthermore, the recipes sketched in section 7.7 could be developed into validated *instruments* for the specification and assessment of organisational anomaly monitoring systems.

In general, the VSM is very useful for the analysis of information flows through different types of communications channels that exist in viable organisational systems. Organisational learning capability is negatively influenced when the variety of messages is not handled adequately. The management in charge may suffer from data overflow, as also may the communication channels, which inevitably have limited capacity. Organisational memory must be (made) accessible for retrieval of earlier surprise experiences and lessons

⁴⁵ On the one hand, we have seen that a message used for notification of an operational surprise cannot contain the variety of the situational context of the surprise occurrence. Ultimately the variety of a surprise notification message may be reduced to two alternatives: 'surprise!' or 'okay'. If a notification message can be matched with a case base in which the cases stem from the same or a similar operational situation, then a low-variety partial description may already result in a match without the need to provide more specific details. In this way, the variety in the notification message is amplified through the case matching. The context is embedded in the case base.

learned by organisational units that experience and notify operational surprises. These and - using VSM - other identifiable explicit links between organisational learning and (safety) culture need further exploration to understand how to remove or design out barriers to learning about improving control of operational risks.

8.2 Epilogue

Looking backward, I have to admit that none of the projects described in this thesis were a complete success. The different projects reveal that organisational learning from operational surprises is complex and difficult to design, not to mention to implement. As such, all these projects presented surprises in the intended process of learning how to learn from small-scale incidents and each contributed to the lessons to learn about learning from incidents. The cultures in different organisations obviously had their impact on the success of projects. Some were and are highly defensive. The key premise remains that an organisation must want to learn from its operational surprises or it will not do so, no matter what effort is invested by researchers, consultants or even the staff in a given part of the organisation. It is an interesting question for further study whether a well-designed organisational learning system for this purpose may help the organisation to propagate a renewed ethos or safety culture down to the lowest meaningful recursions.

This thesis started with organised collection of accident data and ended with organisational learning from incidents. A major shift was the change from the top-down to the bottom-up approach, acknowledging the basic fact that unwanted operational surprises occur at the work floor level of organisational units and end users. Although systemic causal factors may be far beyond their direct control, these people and their desire to learn are the keys to success. By starting from the level of operational disturbances, it is now possible – using the VSM-based recipes – to identify relevant management and anti-oscillatory regulatory systems. Thus, organised learning from small-scale incidents really has become feasible now we have a more profound understanding of the specific nature of the learning processes that at first sight seem to focus on the unwanted rather than the intended outputs of organised activities. Although we now know how to design the organisational learning, the obvious future work is to put this knowledge to the test by going from A – Z in design and implementation.

References

- [Aamodt 94] Aamodt A., Plaza, E. *Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches*. In: AICom - Artificial Intelligence Communications, Vol. 7, No. 1, 1994
- [ASRS 98] Aviation Safety Reporting System. *ASRS Program briefing*. Available at URL <http://asrs.arc.nasa.gov/overview.htm>
- [ARSR 00] *Aviation Safety Reporting System*. URL: <http://www-afo.arc.nasa.gov/ASRS/ASRS.html>
- [Argyris 74] Argyris C., Schön D.A. *Theory in Practice: increasing professional Effectiveness*. Jossey-Bass, San Francisco, 1974
- [Argyris 82] Argyris C. *How learning and Reasoning Processes affect Organisational Change*. In: Goodman P.S. & Associates (ed.). *Change in organisations: New Perspectives on Theory, Research and Practice*. Jossey-Bass, 1982. (reprinted as Chapter I in [Argyris 92])
- [Argyris 92] Argyris C. *On Organizational Learning*. Blackwell, Cambridge/Oxford, 1992
- [Argyris 96] Argyris C., Schön D.A. *Organizational learning II; theory, method, and practice*. Amsterdam, Addison-Wesley, 1996
- [Ashby 56] Ashby W.R. *An Introduction to Cybernetics*. Chapman and Hall, London, 1956
- [Ashby 60] Ashby W.R. *Design for a Brain*. Wiley, New York, 1960
- [Bareiss 89] Bareiss R. *Exemplar-based knowledge acquisition: A unified approach to concept representation, classification, and learning*. Boston, Academic Press, 1989
- [Bateson 72] Bateson G. *Steps to an Ecology of Mind*. Chandler Publishing Co. San Francisco, 1972
- [Beer 79] Beer S. *The Heart of Enterprise*. John Wiley & Sons, Chichester, 1979
- [Beer 81] Beer S. *Brain of the Firm*. 2nd edition. John Wiley & Sons, Chichester, 1981
- [Beer 83] Beer, S. *Introduction: questions of quest*. In: *Cybernetics Theory and Applications*. Edited by Trapp, R., London, Hemisphere, 1983
- [Beer 85] Beer S. *Diagnosing the system for organisations*. John Wiley & Sons, Chichester, 1985
- [Beer 89] Beer S. *The viable system model: its provenance, development, methodology and pathology*. In: *The viable system model: interpretations and applications of Stafford-Beer's vsm*. Edited by Espejo R., and Harnden R., John Wiley & Sons, Chichester, 1989. pp. 39-50.

- [Bellamy 99] Bellamy L.J., Papazoglou I.A., Hale A.R., Aneziris O.N., Ale B.J.M., Morris M.I., Oh J.I.H. *I-Risk: Development of an integrated technical and management risk control and monitoring methodology for managing and quantifying on-site and off-site risks*. Contract ENVA-CT96-0243. Report to European Union. Ministry of Social Affairs and Employment. Den Haag, 1999
- [Benner 83] Benner Jr. L: *Task report no. 1: Accident Models and Investigation Methodologies employed by U.S. Government Agencies in Accident Investigation Programs*. OSHA, Washington DC, 1983
- [Benner 85] Benner Jr. L: *Rating Accident Models and Investigation methodologies*. In: *Journal of Safety Research* 16(3):105-126, Fall 1985
- [Bergmann 98] Bergmann R. *Introduction to Case-Based Reasoning*. University of Kaiserslautern, 1998 [www.cbr-web.org/CBR-Web/cbrintro/sld001.html]
- [Bird 76] Bird Jr. F.E., Loftus R.G. *Loss Control Management*. ILCI, Institute Press, Loganville GA, 1976.
- [Boeing 70] *Sneak Circuit Analysis Handbook*. Boeing Aerospace Company, Space Systems Division, Doc.no. D2-118341-1, Houston TX, 1970
- [Briscoe 91] Briscoe G.J. *MORT-based Risk Management*. EG&G, SSDC WP-28, Idaho Falls, Idaho, March 1991
- [Bullock 79] Bullock M.G. *Work Process Control Guide*. EG&G, DOE-76/45-15 SSDC-15, Idaho Falls, 1979
- [Bullock 81] Bullock M.G. *Change Control and Analysis*. EG&G, DOE-76/45-21 SSDC-21, Idaho Falls, Idaho, 1981
- [Buys 95] Buys J.R., Clark J.L.; Rev.: J. Kingston-Howlett, H.K. Nelson. *Events and Causal Factors Analysis*. SCIE-DOE-01-TRAC-14-95. Scientec, Idaho Falls, 1995
- [CDRH 96] Center for Devices and Radiological Health. *Medical Device Reporting: An Overview*. CDRH, Rockville MD, April 1996 [www.fda.gov/cdrh/mdrovrwv.pdf]
- [CDRH 00] Center for Devices and Radiological Health. *Reporting Problems with Medical Devices*. Medical Devices Reporting System On-line Data Search Page. CDRH, Rockville MD, 2000 [www.fda.gov/cdrh/mdr.html]
- [CEC 90] Committee on Energy and Commerce. *The Björk-Shiley heart valve: "Earn as You Learn". Shiley Inc.'s Breach of the honor system and FDA's failure in medical devices regulation*. U.S. House of Representatives, February 1990. US GPO 26-766. Washington DC, 1990
- [Chopra 92] Chopra V., Bovill J.G., Spierdijk J., Koornneef F. *Reported Significant Observations during Anaesthesia: a prospective analysis over an 18-month period*. *British Journal of Anaesthesia* 1992; **68**: 13-17.
- [Chopra 94] Chopra V. *Aspects of Quality Assurance in Anaesthesia - Development of an Anaesthesia Simulator*. PhD Thesis, Leiden, 1994
- [Cornelison 89] Cornelison J.D. *MORT-based Root Cause Analysis*. EG&G, SSDC WP-27, Idaho Falls, Idaho, May 1989
- [Cromheecke 00] Cromheecke M., Koornneef F., Gaalen G. van, Mol B. de. *Controlling the Risks of Mechanical Heart Valve Failure using Product Life Cycle-based Safety Management*. In: Vincent C., Mol B.A. de (eds.) *Safety in Medicine*. Pergamon, Amsterdam, 2000

- [Crosetti 82] Crosetti P.A.: *Reliability and Fault Tree Analysis Guide*. DOE-76/45-22 SSDC-22, Idaho Falls, 1982
- [Dewey 38] Dewey J. *Logic: the Theory of Inquiry*. Holt, Rinehart and Winston. New York, 1938
- [Dijen 2000] Dijen M. van. *Gecombineerd management van kwaliteit, arbeidsomstandigheden, milieu en patiëntenveiligheid*. PhD Thesis, Leiden University. Leiden, 2000
- [DIN 94] *DIN 19250 - Grundlegende Sicherheitsbetrachtungen für MSR-Schutzrichtungen*. DIN (Deutsches Institut für Normung e.V.) (1994). Draft norm DIN V 19250. May 1994
- [DMSA 93] *Memorie van Toelichting: Aangepaste Arboret 1993*. (Explanatory Memorandum: Revised Working Conditions Law 1993). Dutch Ministry of Social Affairs. Staatsuitgeverij, Den Haag, 1993
- [DOE 92] *Root Cause Analysis Guidance Document*. DOE Guideline DOE-NE-STD-1004-92. Washington DC, 1992
- [DSI 00] *REASON Root Cause Analysis*. DSI, Longview TX. <http://www.rootcause.com/>
- [EC 93] *European Directive on Medical Devices*. MDD: 93/42/EEC. June 14, 1993
- [Eicher 76] Eicher R.W. *Reported Significant Observation (RSO) Studies*. ANC, ERDA-76/45-5 SSDC-5, Idaho Falls, Idaho, March 1976
- [Eijkel 95] Eijkel G.C. van den, E. Backer E., Gerbrands J.J.. *A Proposal for Fuzzy Rule Generation using Temporal Reasoning*. In: Katwijk J. van, Gerbrands J.J., Steen M.R. van, Tonino J.F.M. (eds.). Proceedings of the first annual Conference of the Advanced School for Computing and Imaging (ASCI '95), Delft, 1995, pp. 47-55.
- [Eijkel 99] Eijkel C.G. van den. *Fuzzy probabilistic Learning and Reasoning*. Ph.D. Thesis. TU Delft, 1999
- [Elsea 83] Elsea K.J., Conger D.S.: *Management Oversight and Risk Tree*. In: the Risk Report, Vol. VI, no. 2. International Risk Management Institute, 1983
- [Espejo 89] Espejo R., Harnden R.J. (eds.) *The Viable System Model: Interpretations and Applications of Stafford-Beer's vsm*. John Wiley & Sons, Chichester, 1989
- [FDA 92] Federal Food, Drug, and Cosmetic Act (Sec. 513), amended by the Safe Medical Devices Act of 1990 (SMDA) and the Medical Device Amendments of 1992.
- [FDA 99] Task Force on Risk Management. *Managing the Risks from Medical Product Use - Creating a Risk Management Framework*. FDA, Rockville MD, 1999 [<http://www.fda.gov/oc/tfrm/riskmanagement.pdf>]
- [Ferry 88] Ferry T S: *Modern Accident Investigation and Analysis*. 2nd edition. Wiley, New York, 1988
- [Flanagan 54] Flanagan J.C. The Critical Incident Technique. *Psychological Bulletin*, 51, 327-358. 1954
- [Gaba 94] Gaba D.M., Fish K.J., Howard S.K. *Crisis Management in Anesthesiology*. Churchill Livingstone, New York, 1994
- [Glantz 98] Glantz S.A., Slade J., Bero L.A., Hanauer P., Barnes D.E. (eds.). *The Cigarette Papers*. The University of California Press, San Francisco, 1998
- [Graaf 92] Graaf Y. van der, Waard F. de, Herwerden L.A. van, Defauw J.J. *Risk of Strut Fracture of Björk-Shiley Valves*. The Lancet 1992; **339**: 257-261.
- [Graaf 98] Graaf P.M.A. *How to Design an Intelligent Anesthesia Monitor – Safety and Cognition based Design in Anesthesia Monitoring*. Ph.D. Thesis. TU Delft, 1998

- [Groeneweg 98] Groeneweg J. *Controlling the controllable: the management of safety*. 4th Edition. DSWO Press, Leiden, 1998
- [Guldenmund 97] Guldenmund F., Graaf P.M.A. de. *AMC – ISA Project – Nulmeting*. Internal report. TU Delft, Safety Science Group, 1997.
- [Guldenmund 00] Guldenmund, F.W. *The Nature of Safety Culture: a Review of Theory and Research*. *Safety Science* **34** (2000): 215-257.
- [Haddon 66] Haddon Jr. W. *The Prevention of Accidents*. In: *Preventive Medicine*, Little Brown, Boston, 1973
- [Haddon 73] Haddon Jr. W. *Energy Damage and the Ten Counter-Measure Strategies*. In: *Human Factors Journal*, August 1973
- [Hale 70] Hale A.R., Hale M. *Accidents in Perspective*. *Occupational Psychology* 1970; **44**: 115-122.
- [Hale 94] Hale, A.R., Pelt E. van, Heimplaetzer P., Swuste P., Guldenmund F., Heming B.. *Influencing Perception of Accident Causation*. Proceedings 23rd International Congress of Applied Psychology, Madrid, 1994
- [Hale 87a] Hale A.R., Glendon A.I. *Individual Behaviour in the Control of Danger*. Elsevier Science Publishers. Amsterdam, 1987
- [Hale 87b] Hale A.R., Vroege D., Otto E.H., Burdorf, A. *Accidents in Ports. Report of a Pilot Study to set up an Accident Registration System in the Port of Rotterdam*. TU Delft Delft/EU Rotterdam, October 1987
- [Hale 91] Hale A.R., Karczewski J., Koornneef F., Otto E.H. *IDA: an interactive Program for the Collection and Processing of Accident Data*. In: Schaaf, T.W. van der, Lucas, D.A., Hale, A.R. (eds.) *Near Miss Reporting as a Safety Tool*. Butterworth Heinemann. Oxford, 1991
- [Heinrich 31] Heinrich H.W. *Industrial Accident Prevention*. McGraw Hill, New York, 1931
- [Heinrich 80] Heinrich H.W., Petersen D., Roos N.: *Industrial accident prevention; a safety management approach*. 5th ed. McGraw-Hill, New York, 1980
- [Hendrick 87] Hendrick K., Benner Jr. L.: *Investigating Accidents with STEP*. Marcel Dekker, Inc., New York, 1987
- [Hirschfeld 63] Hirschfeld A.H., Behan R.C.: *The Accident Process*. In: *Journal of the American Medical Association*, 186:193-199, October 19, 1963
- [Horman 84] Horman R.L. *Glossary of SSDC Terms and Acronyms*. EG&G, DOE-76-45/28 SSDC-28, Idaho Falls, Idaho, October 1984
- [HSE 86] *Deadly Maintenance: a study of fatal accidents at work*. HSE. HSMO. ISBN 0 11 883806 7. London, 1986
- [HSE 87] *Dangerous Maintenance: a study of maintenance accidents in the chemical industry and how to prevent them*. HSE. HSMO. ISBN 0 11 883957 8. London 1987.
- [IEC 98] IEC 16508 – *Functional Safety of Electrical/Electronic/Programmable Electronic Safety-related Systems*. IEC 1998 & 2000
- [ISO 87] ISO 9004 – *Quality Management and Quality System Elements. Guidelines*. CEN, Brussels, 1987.
- [Ives 91] Ives G. *"Near Miss" Reporting Pitfalls for Nuclear Plants*. In: Schaaf, T.W. van der, Lucas, D.A., Hale, A.R. (eds.) *Near Miss Reporting as a Safety Tool*. Butterworth Heinemann. Oxford, 1991

-
- [Johnson 73] Johnson W.G. *MORT-The Management Oversight and Risk Tree*. ERDA SAN 821-2, Idaho Falls, 1973.
- [Johnson 80] Johnson W.G.: *MORT Safety Assurance Systems*. Marcel Dekker, Inc., New York and Basel, 1980.
- [Johnson 85] Johnson W.G., Nertney R.J., et al. *Accident/Incident Investigation Manual (2nd edition)*. EG&G, DOE-7645/27 SSDC-27, Idaho Falls, Idaho, November 1985
- [Jørgensen 98] Jørgensen K. *The Use of Accident Statistics for Prevention - Method of Compiling and Analysing*. In: De Rademaeker E., Pineau J.-P. *Accident Databases as a Management Tool*. Proceedings of the 15th ESReDA Seminar. Technologisch Instituut, Antwerpen 1998
- [Johnson 99] Johnson C. *Using Case-Based Reasoning to Supporting the Indexing and Retrieval of Incident Reports*. Draft. 1999 [www.dcs.gla.ac.uk/~johnson/papers/case_based_reasoning]
- [Kallewaard 97] Kallewaard M. *Björk-Shiley convexo-concave Heart Valves: the Risk of Outlet Strut Fracture and Clinical Implication*. PhD Thesis. Utrecht, 1997
- [Karczewski 93] Karczewski J.T. System Komputerowej Analizy Wypadków Przy Pracy ISA-PL - Podręcznik użytkownika. CIOP, Warsaw, 1993
- [Kepner 81] Kepner C.H., Tregoe B.B. *The New Rational Manager*. Princeton Research Press, Princeton NJ, 1981.
- [Kingston 96] Kingston-Howlett J.C. *The Evolution and Design of Safety Management Systems*. PhD Thesis. Aston University, Birmingham, UK, 1996
- [Kjellén 83a] Kjellén U. *Analysis and Development of Corporate Practices for Accident Control*. OARU, Royal Institute of Technology, Stockholm, 1983
- [Kjellén 83b] Kjellén U. *The Deviation Concept in Occupational Accident Control Theory and Method*. OARU, Royal Institute of Technology, Stockholm, 1983
- [Kletz 92] Kletz T. *HAZOP and HAZAN: Identifying and Assessing Process Industry Hazards*. Hemisphere Publishing Corporation, Bristol USA, 1992
- [Knox 92] Knox N.W., Eicher R.W. *MORT User's Manual. For Use with the Management Oversight and Risk Tree Analytical Logic Diagram*. DOE-76/45-4 SSDC-4 (Rev. 3), Idaho Falls, 1992
- [Kolodner 83] Kolodner J. *Reconstructive memory, a computer model*. Cognitive Science, Vol. 7, s.281-328. 1983.
- [Kolodner 93] Kolodner J. *Case-based reasoning*. Morgan Kaufmann, 1993
- [Kohn 99] Kohn L.T., Corrigan J.M., Donaldson M.S. (Eds.). *To Err Is Human - Building a Safer Health System*. Committee on Quality of Health Care in America. Institute Of Medicine. National Academy Press, Washington DC, 1999
- [Koornneef 90] Koornneef F., Mol B.A.J.M. de, Heimlaetzer P.V. *Risk Control of Biomedical Products in Health Care Systems*. In: Cornelis J., Peeters S. (eds.). *Proceedings North Sea Conference on Biomedical Engineering 1990*. Technologisch Instituut, Antwerpen, November 19-22, 1990
- [Koornneef 93] Koornneef F., Mol B.A.J.M. de. *Report on Safety and Quality Assurance of the Sorin Bi-leaflet Prosthetic Heart Valve*. Delft University of Technology, Safety Science Group. February 1993

- [Koorneef 96a] Koorneef F., Spijkervet A.L., Karczewski J.T. *Organisational Learning using Near-Miss and Accident Data within and outside your Organisation*. In: Redmill F. Proceedings Safety-critical Systems Symposium 1996. Safety-Critical Systems Club. Springer, 1996
- [Koorneef 96b] Koorneef F., Gaalen G.L. van, Mol B.A. de. *A Risk Assessment and Control Model for the failing BjörkShiley convexo-concave Heart Valve*. International Journal of Technology Assessment in Health Care 1996; **12** (1): 141-145.
- [Koorneef 97a] Koorneef, F., Hale, A.R. *Learning from Incidents at Work*. In: Redmill F., Rajan J. (eds.) *Human Factors in Safety Critical Systems*. Butterworth Heinemann, Oxford, 1997
- [Koorneef 97b] Koorneef F., Hale A.R. *Using MORT to generate Organisational Feedback from Single Accidents at Work*. In: Hale A.R., Wilpert B., Freitag M. (eds.) *After the Event: from Accident to Organisational Learning*. Oxford, New York, Tokyo. Pergamon, 1997
- [Koorneef 98] Koorneef F., Kingston-Howlett J. *Accident Data in Organisational Learning, or What Makes Accident Databases Useful?* In: De Rademaeker E., Pineau J.-P. *Accident Databases as a Management Tool*. Proceedings of the 15th ESReDA Seminar. Technologisch Instituut, Antwerpen, 1998
- [Koton 89] Koton P. *Using experience in learning and problem solving*. MIT, Laboratory of Computer Science (Ph.D. diss, October 1988). MIT/LCS/TR-441, 1989
- [Lepelaars 97] Lepelaars E.S.A.M., Ooijen W.D.R. van, Tijhuis A.G., Mol B.A.J.M. *Detecting Fractures in Artificial Heart Valves*. In: IEEE Antennas and Propagation Society International Symposium 1997, Digest – Volume 2, pp. 1198-1201. Montreal 1997
- [Leveson 93] Leveson N.G, Turner C.S. *An Investigation of the Therac-25 Accidents*. IEEE Computer 1993; **July**: 18-41.
- [Leveson 95] Leveson N.G. *Safeware - System Safety and Computers*. Addison-Wesley, 1995.
- [MacDonald 72] MacDonald G L: *The Involvement of Tractor Design in Accidents*. Research Report 3/72, Dept. of Mechanical Engineering, University of Queensland, St. Lucia, USA, 1972
- [Maturana 84] Maturana H.R., Varela F.J. *The Tree of Knowledge: the biological Roots of human Understanding*. Shambala, Boston, 1984
- [McFarland 67] McFarland R A: *Application of Human Factors Engineering to Safety Engineering Problems*. National Safety Congress Transactions, Chicago, 1967
- [Mintzberg 83] Mintzberg H. *Structure in fives: designing effective organizations*. Prentice-Hall, Englewood Cliffs, 1983
- [Mol 95] Mol B.A. de, Koorneef F, Gaalen G.L. van. *What can be done to improve the safety of heart valves?* International Journal of Risk & Safety in Medicine 1995; **6**: 157-168.
- [Nertney 75] Nertney R.J., Clark J.L., Eicher R.W. *Occupancy-use Readiness Manual*. ANC, ERDA-76/45-1 SSSDC-1, Idaho Falls, Idaho, September 1975
- [Nertney 78] Nertney R.J. *Management Factors in Accident and Incident Prevention*. DOE-76/45-13 SSSDC-13, Idaho Falls, 1978
- [Nertney 87] Nertney R.J. *Process Operational Readiness and Operational Readiness Follow-on*. EG&G, DOE-76-45/39 SSSDC-39, Idaho Falls, Idaho, February 1987

-
- [Nimmo 95] Nimmo I. *Abnormal Situations Management: giving your control system ability to 'cope'*. Journal for Industrial Automation and Control, Vol.2 No.2, August 1995, pp. 6-7. <http://www.iac.honeywell.com/Pub/Journal/a9507abnorm.html>
- [OSHA 98] Occupational Safety and Health Administration. *The New OSHA: Reinventing Worker Safety and Health*. December 16, 1998. Available at: <http://www.osha-slc.gov/html/Reinventing/intro.html>
- [Paradies 91] Paradies M., Unger L. *TapRoot Incident Investigation System: Executive Summary*. Volume 1. Systems Improvements, Inc., Knoxville, TN, 1991.
- [Perrow 84] Perrow C. *Normal accidents: living with high-risk technologies*. Basic Books, New York, 1984
- [Porter 90] Porter B., Bareiss R., Holte R. *Concept learning and heuristic classification in weak theory domains*. Artificial Intelligence, vol. 45, no. 1-2, September 1990. pp. 229-263.
- [Powel-Price 98] Powel-Price M. *The Accident Database: capturing corporate Memory*. In: De Rademaeker E., Pineau J.-P. Accident Databases as a Management Tool. Proceedings of the 15th ESReDA Seminar. Technologisch Instituut, Antwerpen, 1998.
- [Rankin 73] Rankin J.P. *Sneak Circuit Analysis*. Nuclear Safety, Vol. 14, No. 5, 1973
- [Rasmussen 80] Rasmussen J. *What can be learned from human Error Reports?* In: Duncan et.al. (eds.): Changes in Working Life, Wiley, Chichester, 1980
- [Reason 87] Reason J.T. *A Framework for Classifying Errors*. In: Rasmussen et.al.: New Technology and Human Errors, Wiley, Chichester, 1987
- [Reason 90] Reason J.T. *Human Error*. Cambridge University Press. Cambridge, 1990
- [Robens 72] Lord Robens: *Safety and Health at Work: Report of the Committee 1970-1972*. HMSO, London, 1972
- [Roggeveen 99] Roggeveen V., Brink H. van den, Lambers M., Groeneweg J. *Tripod Delta: Detecting organisational Vulnerability by Mapping Operational Performance*. In: Goossens L.H.J. (Ed.). *Risk Analysis: Facing the New Millenium*. Proceedings of the 9th Annual Conference . SRA Europe, TU Delft. Delft University Press, 1999
- [Rosenthal 97] Rosenthal I. *Organisational Epidemiology: a Tool for Investigating Organisational Factors related to the Occurrence and Prevention of Accidental Chemical Releases*. In: Hale A.R., Wilpert B., Freitag M. (eds.) After the Event: from Accident to Organisational Learning. Oxford, New York, Tokyo. Pergamon, 1997
- [Sale 96] Sale A.E. *The Colossus Rebuild project*. 1996. <http://www.cranfield.ac.uk/cc/bpark/colossus-update.html>
- [Schaaf 91a] Schaaf T.W. van der, Lucas D.A., Hale A.R. (eds): *Near Miss Reporting as a Safety Tool*. Butterworth Heinemann, Oxford, 1991
- [Schaaf 91b] Schaaf T.W. van der. *A Framework for Designing Near-Miss Management Systems*. In: Schaaf, T.W. van der, Lucas, D.A., Hale, A.R. (eds.). Near Miss Reporting as a Safety Tool. Butterworth Heinemann. Oxford, 1991
- [Schaaf 96] Schaaf T.W. van der. *PRISMA: A Risk management Tool based on Incident Analysis*. Proceedings of the International Conference and Workshop on Process Safety Management and Inherently Safer process, pp. 242-251. Orlando, October 8-11, 1996

-
- [Sivro 97] Sivro S., Graaf R. de. *Protocol Cardiopulmonale Chirurgie G3 Zuid*. AMC, Amsterdam, 1997
- [Spek 97] Spek R. van der, Spijkervet A. *Knowledge Management - Dealing Intelligently with Knowledge*. Knowledge Management Network, Utrecht, 1997
- [Spijkervet 94] Spijkervet A.L. *Eind-/Effectrapportage PSO-Project ARA - Accident Registration and Analysis*. CSC Europe, February 1994
- [Surry 69] Surry J. *Industrial Accident Research: a Human Engineering Appraisal*. University of Toronto, Department of Industrial Engineering, Toronto, 1969
- [Thomas 99] Thomas E.J., Studdert D.M.; Newhouse J.P., et al. *Costs of Medical Injuries in Utah and Colorado*. Inquiry 36:255-264, 1999.
- [Thomas 95] Thomas M.T. *An Analysis of the causal Factors leading to fatal Accidents in Agriculture in Great Britain*. Journal of Health and Safety, Number 12, 1995
- [Tripod 99] Tripod Beta version 2 HSE Manual. Tripod Foundation. EQE Int. – 301/00 Issue 01. April, 1999
- [Trost 95] Trost W.A., Nertney R.J. Rev.: J. Kingston-Howlett, R.J. Nertney, H.K. Nelson. *Barrier Analysis*. SCIE-DOE-01-TRAC-29-95. Scientec, Idaho Falls, 1995
- [Vullings 99] Vullings E. *Biomedical Waveform Validation*. Ph.D. Thesis. TU Delft, 1999
- [Vuuren 98] Vuuren W. van. *Organisational Failure – An exploratory study in the steel industry and the medical domain*. Ph.D. thesis. TU Eindhoven. Eindhoven, 1998
- [Watson 94] Watson I., Marir F. *Case-Based Reasoning: A Review*. In: The Knowledge Engineering Review, Vol.9, No.4, 1994
- [Weegels 96] Weegels M.F. *Accidents involving Consumer Products*. PhD Thesis. TU Delft. Delft, 1996
- [Wells 77] Wells L.D. *Quality Assurance of Patient Care*. In: Rhodes B.A. (ed.) *Quality Control in Nuclear Medicine*. 1977
- [Wickens 92] Wickens C.D. *Engineering psychology and human performance*. 2nd edition. New York, NY, Harper Collins Publishers, 1992

MORT Chart: main Structure and Use in Accident Analysis

The Management Oversight and Risk Tree (MORT) is a powerful investigation tool. The MORT (single page) chart covers 98 generic safety problems, broken down into over 1500 basic 'causes' within the safety system of the process owner [Elsea 83]. Combining principles from the fields of management and safety and using fault tree methodology, the MORT tree aims at helping the investigator discover what happened and why. Both specific control factors and management system factors are analysed for their contributions to the accident. People, procedures, and plant-hardware are considered first separately, then together, as key system safety elements [Knox 92]. A great advantage of the MORT chart is that it is already drawn out and so is ready for use. It can be applied whenever a harmful energy flow or hazardous environmental condition can be identified. In the case of a major accident, the investigation can turn into a safety programme review by linking the whole M-branch on management system factors with the incident under review (see Figure A-1 below).

In the 1970s, William Johnson and his team conceived of MORT as a way to decrease, by an order of magnitude, the safety-related losses of an organisation with an already excellent safety record. Besides the MORT chart, the MORT system also consists of a system safety process and a means for problem solving [Johnson 80, Knox 92]. The philosophy underlying MORT can be summarised as follows [Elsea 83]:

- 1 *Management takes risks of many kinds.* These risks are in the areas of product quantity and quality, costs and schedules, and in the area of environment, health and safety.
- 2 *Risks in one area affect operations in other areas.* Management's job may be viewed as one of balancing risks. For example, to focus only on safety and environmental issues might increase the risk of losses from deficiencies, schedule delays, and related costs.
- 3 *Risks should be made explicit where practicable.* Since management must take risks, it is helpful for management to know the potential consequences of those risks.
- 4 *Risk management tools should be flexible enough to suit a variety of diverse situations.* While in-depth analytical tools are needed for complex situations, other situations require simpler, quicker approaches.

Part of the MORT system is a toolbox containing analytical techniques for accident investigation aimed at problem solving. In increasing order of complexity, these techniques are: Change Analysis [Bullock 81], Energy-Barrier-Target analysis [Trost 85], Fault Tree Analysis [Crosetti 82], Events and Causal Factors Charting [Buys 78], and the MORT chart analysis [Knox 92, Johnson 80]. Only the MORT chart has comprehensive built-in (safety) expertise; the other tools are purely procedural. Essentially, the MORT chart is a distillation of forty years of research work in safety analysis.

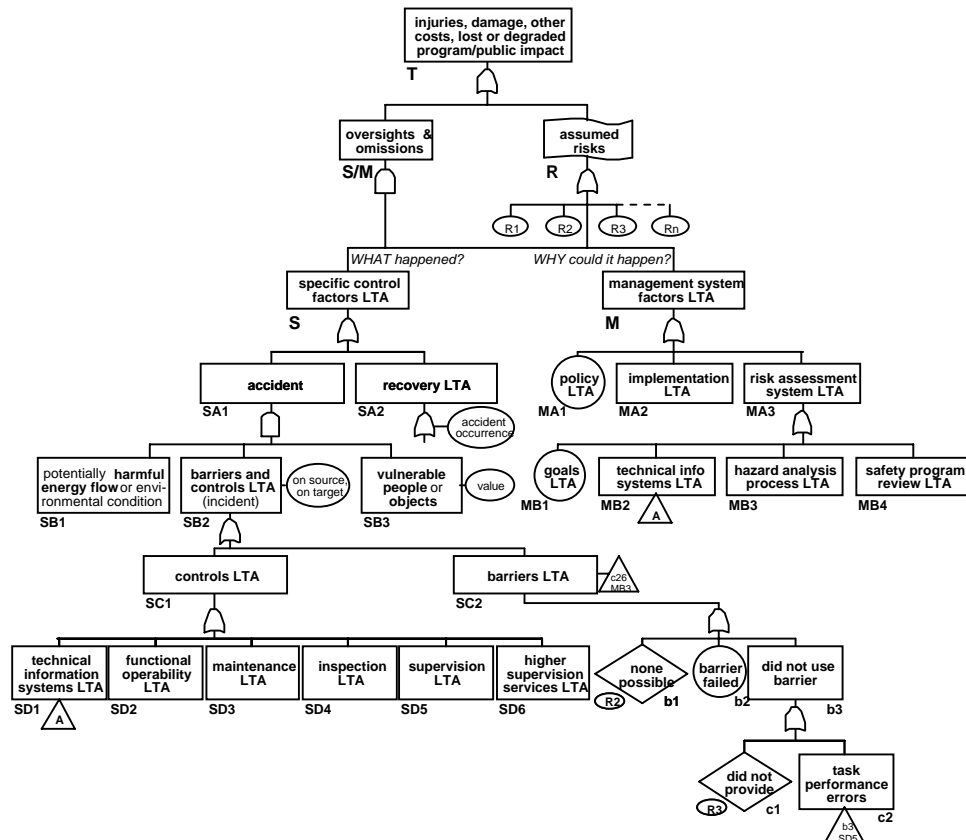


Figure A-1: the basic Management Oversight and Risk Tree (simplified)

Using the MORT Tree

The basic MORT-tree in Figure A-1 is best described as 'structured common sense'. It has three major tasks:

- it identifies oversights and omissions;
- it enables residual risks to be seen and, where possible calculated;
- it enables the analyst, the employer, or whoever is using the chart, to put the effort and the finance where it is most needed.

The MORT chart looks like a fault tree, but rules for going through the tree are modified. In addition to the fault tree logic (AND and OR gates) the user must work not only from top to bottom, as in a fault tree, but also from left to right [Conger 84]. In many branches, a tree event to the left is a prerequisite to the one under review; e.g. in the barrier sub-tree (**SC2**), it would not make sense to investigate whether a prescribed, specific barrier meets its required specifications if such a barrier could not be applied under the actual pre-incident circumstances.

In incident analysis, the MORT chart says that adverse outcomes or losses are the result of oversights or omissions concerning relevant risks which management should have known of, or the anticipated consequence of explicit acceptance of an identified, assessed risk. This 'accepted' risk must have been evaluated by a person who had authority delegated by management to assume that (residual) risk, because it was considered to be:

- tolerably low (minor) in frequency or consequence,
- high in consequence but impossible to eliminate, (e.g. a hurricane), or
- simply too expensive to correct when weighed against the risk consequences.

The overall losses form the Top-event (**T**) of the tree. Analysis of the Specific factors branch (S-branch) will show whether any specific relevant risks have been accepted by a management authority or not. If this decision has been made adequately within the context of the specific problem area, the risk is transferred to an oval with code **R#** in the R-branch just below the top of the tree. If the acceptance decision has not been correctly made, this remains an oversight or an omission and, therefore, indicates a management problem.

Moving down the tree from the top leads to the **S/M**-event (Oversights and Omissions) and thence to the first AND-gate. This gate indicates that, in any accident, Specific Control Factors as well as Management System Factors will always be present in a state which can be assessed qualitatively as 'Less Than Adequate': LTA. The important principle here is that, in MORT, specific failures alone cannot be responsible for accidents; the possibility of these specific failures should have been predicted by management and, for those risks which have not been accepted, systems should have been put in place to prevent or detect and correct them.

Moving down further and to the left leads to the 'Accident' (**SA1**) branch and, subsequently to 'Recovery' (**SA2**), i.e. rescue, fire fighting, prevention of secondary accidents, public relations and rehabilitation of victims and operations.

After assessment of the specific control factors, the management system factors may have to be evaluated by going through the M-branch. This is only cost-effective in the case of very serious incidents with a high loss or loss potential or of a series of incidents with a pattern of similar specific failures.

The structure of the MORT chart below the 'Accident' event, **SA1**, is based on the basic energy trace & barrier accident model of Figure A-1. **SB1** represents the harmful energy flow or environmental condition, **SB3** stands for the vulnerable targets and **SB2** for the barriers between the two as well as for energy flow control functions. The AND-gate here says that there will be no accident if at least one of the Specific Control Factors SB1, SB2 or SB3 functions according to adequate plans and specifications. Obviously, the safety margin is increased in relation to other accidents if all three, or at least two, are present

and functioning, although, strictly, only one needs to be brought fully under control in order to prevent future accidents with similar underlying problems.

Post-accident recommendations will focus on resolution of the problems discovered at the level below SB1, SB2 and SB3. The main problem areas will be found by analysing the tree event **SB2**: barriers and controls LTA. This applies even when a vulnerable target was not hit by the energy flow, either by chance or because there was, also by chance, no vulnerable target around. The incident then is a near-miss occurrence.

Specific Control Factor SB2

The tree event **SB2** sits above two branches, see Figure A-1. The first one, starting with **SC1**, is used to evaluate the factors that condition the actual operations. These are:

- Technical Information Systems (**SD1**) with respect to the hazard: knowledge of the work flow process and downstream/upstream communication, monitoring systems, data collection and analysis, triggers for initiating hazard analysis, audit outcomes.
- Facility Functional Operability testing (**SD2**) especially of procedures, Man-Machine interfaces and operational skills. This is relevant at initial start-up of a plant or other safety-critical operation, or at the first start-up after a major plant modification.
- Maintenance (**SD3**),
- Inspection (**SD4**),
- first line Supervision (**SD5**)
- support of the first line supervisor by Higher Supervision Services (**SD6**).

All these factors are aimed at the realisation of an undisturbed process as intended in the system design stage, whereby inherent hazards are kept under control by design as well as by real-time control and adjustments. Problem areas under **SC1** may, therefore, point back to the drawing board of the plant. More often, identified problems call for measures regarding internal communication about known risks and risk-control options, as well as about supervisory improvements in daily operations.

The barrier subtree **SC2** is relevant whenever a harmful energy flow might reach a vulnerable, valuable target, especially a human being. Barriers are things that separate a harmful energy flow or a harmful environmental condition from a vulnerable target, in time or space, either completely or by reducing the energy to a harmless level before the target can be exposed. This sub-tree should be used for each relevant barrier successively. Thus, the relevant barriers must be identified first. By checking each of the four barrier-categories, i.e. on the hazard source (a1), between hazard and vulnerable target (a2), on the target (a3) or as separation in time or space (a4), a list of relevant barriers can be prepared efficiently. These categories have been omitted for clarity in Figure A-1.

I will follow the barrier subtree further, as example, in order to discuss how a MORT-tree analysis works. The first item is tree event **b1**, 'none possible'. For the barrier under review, the question is whether this type of a barrier was possible in the (pre-incident) situation. The answer to MORT-questions is always qualitative and requires a judgement by the investigator. The assessment categories are: *adequate*, *less than adequate (LTA)*, *more information needed*, or *not relevant*. If the barrier under review, in the case in question, was not possible under the pre-incident circumstances, the answer would be:

LTA. This means that a relevant barrier could not be applied where it was needed. This leads to the following question, whether the organisation went through the proper process for accepting the risk of not having the relevant barrier in place. This is denoted in Figure A-1 by the small oval **R2** on the left below the tree event SC2-**b1**. If this risk was accepted in a proper way at the appropriate level of responsible management (**R2** = adequate), then there is an assumed risk **R2** which can be noted in the **R**-branch just below the top of the MORT chart. In that case the other subtree events **b2** and **b3** can be ruled out as being not relevant with respect to the barrier under review.

The circular, basic tree event **b2**, 'barrier failed' relates to the issue as to whether the barrier worked as intended, i.e. conform its specifications. If not then **b2** scores LTA, thus indicating a basic 'cause' of the harm process in the specific accident. Next is **b3**, 'did not use barrier'. The meaning of this can be explored by going down one tier deeper: **b3** results from either **c1**, 'did not provide', or **c2**, 'task performance error'. If the barrier was not provided then it could not be used (**c1** = LTA). Again, the issue of acceptance of an identified risk (**R3**) by appropriate management is raised for assessment at this point. If the barrier was provided (**c1** = adequate) then it might have been used wrongly or not at all (**c2** = LTA). This can be assessed in detail by entering the task performance subtree section on the chart (**SD5-b3**), denoted by the triangular transfer symbol underneath tree event **SD5**. This branch is not worked out in Figure A-1, but the topic will be discussed in the next section. Note that the analysis of a specific barrier stops when each of the subtree events **b1**, **b2** and **b3** has been assessed. The approach outlined above applies to the whole MORT-tree.

Task Performance Errors in the MORT Chart Analysis

The role of human operators is consistently reviewed in relation to the specific task or activity on which they were working either during the incident or at some earlier time relevant to the incident, e.g. a maintenance task which was not performed according to plan. If a task performance error is identified, e.g. in relation to maintenance, inspection, supervision, the use of a specific barrier or during initial recovery action, then an extensive subtree section (**SD5-b3**: not shown in Figure A-1) is provided by the MORT chart in order to evaluate underlying management factors such as the functional situation of the operator, training, the role of the supervisor, motivation (including issues like social climate), time pressure, job interest building, group norms, physical hindrances. Management system factors underlying task performance errors must be analysed in greater detail. There are seven main problem areas with respect to task performance errors, of which only one, namely 'personnel performance discrepancy' as a sub-problem of 'employee motivation', permits the user to indicate that deviating behaviour by an individual operator beyond management control was a determinant of the accident. This underlines an important principle within the MORT system, that the analyst should be discouraged from pointing the finger too hastily at the operator as a scapegoat when it comes to task performance errors in a safety-critical system: the bottom line is not to *blame* somebody but to *gain* improved control over safety-critical systems.

(this page is left blank intentionally)

Appendix B

ISA Program for Incident Data Collection and K-model based Analysis

The listing of the ISA-PL program (excluding the modules) as used in the pilot project ARA-1, see section 2.3.1. A dedicated interpreter is needed to execute the program code that allows navigational control of input and ruled-based generation of diagnostic reports.

```
REM <Graph ISA: ARA1+ MORT APPLICATION>
REM registration, classification and diagnosis of incidents
REM Author: Floor Koornneef / TU Delft
REM relates to MORT graphs & db-definition d.d. version 1d, d.d. 951114
REM Version: 1e / 951213
REM 2-way data link integrity implemented for specSD5b
REM CHECK <record> statements have been REM-ed for the ARA1+ trial

REM Modifications & Corrections made per version date
REM Version 1d has an improved version of the performance error message
  generator
REM tested with ara.exe shell version 10/11/95
REM ara.exe shell does NOT yet support MULTI-VALUE fields
REM no "NEWLINE" occurrence: replaced by changing ADD <text> into ADDLN
  <text>
REM assign & if...then statements wrt SB1.Haz_cond/SB3.Tar_cond corrected
  to YN
REM IF readref THEN tst<assumed risk> replaced by IF readref THEN
  rref<assumed risk> ELSE tst<assumed risk>
REM In database definition (TABELE.DBF) the H07 must be linked to HAZ_R4_S
  and H08 to HAZ_R4RE_S
REM IF SC1.SD3/SD4 <> mort4 corrected into IF SC1.SD3/SD4 = mort2 THEN
  event63/65 ELSE event64/66
REM sometimes after an initially wrong attempt, a corrected Date entry is
  not stored in database
REM init SC1.SD3/4err corrected from "N" to "Y" according to the phrasing
  of the questions
REM wrong label notarassnext corrected to notaroprnext
REM incorrect field names specSC2.Bar_al...4 renamed to specSC2.Bar_a_s
REM ADD(LN) empty statements REMmed
REM logical tests IF Record.Field = true/false changed to IF record.Field =
  T/F
REM no taskinfo asked for visitors, inspectors and others (empmod4/5/7)
REM AND gates SA1 = SB1+SB2+SB3 and O=S+M taken into account
REM ANAL-UIT graph still to be implemented
```

```
REM to be done before trial release: test cases harmless incident,  
disaster, material damage, victims  
REM reviewed: logic of T, O, R, S, M, SC1, SB1, SB3  
REM to be reviewed: logic of SC2, SB2; generation of messages for test  
cases (see above)  
REM IF specSB3.TarStat = tarstat2 AND morttop.SB3 = mort1 THEN morttop.SB3  
:= mort3  
REM IF (morttop.) SB1 = mort2 AND SB2 = mort2 AND SB3 = mort3 THEN near-  
miss incident <=> SA1 := mort3  
REM filtering by SIL < 2 criterion disabled (IF eventcat.SIL < 2 remmed)  
REM incorrect second label inccat1b: replaced by inccat1c:  
REM IF specSB3.TarStat = tarstat1 THEN morttop.SB3 := mort2 THEN  
morttop.SA1 := mort2  
REM IF errcat2 THEN specSD5b.tperfb3_s := mort3  
REM wrong IF NOT EXIST NEXT <record> / SET <record> NEXT combinations  
corrected  
REM wrong IF analysis.Countflag = "true" statements corrected into IF  
analysis.Countflag = T  
  
REM <Graph MAIN>  
  
CALL init  
CALL general  
CALL spectarget  
CALL spechazard  
CALL controls  
CALL specbarrier  
CALL finish  
CALL updtproots  
CALL analysis  
CALL finish  
REM CALL prot  
COMMIT  
END
```

Appendix C

RSO Data Model

Data model for case recording in the RSO Study in Medical Centre A (see 2.2.5 and also 2.3.2)

D A T A B A S E S C H E M E

1 - Record: patient

	Field name	type	table	question	help
1 -	naam	string[32]	---	1	DEF
2 -	leeftijd	integer	---	2	DEF
3 -	ziekenhuisnr	string[16]	---	3	DEF
4 -	fysieke status	code 1	1	4	DEF
5 -	sex	code 1	2	5	DEF

2 - Record: melding

	Field name	type	table	question	help
1 -	melder code	string[8]	---	6	DEF
2 -	RSO datum	date	---	7	DEF
3 -	RSO tijd	time	---	8	DEF
4 -	RSO beschrijving	text[20]	---	9	DEF
5 -	patient	boolean	---	10	DEF
6 -	OR procedure	string[32]	---	11	DEF
7 -	tijdsduur ORproc	time	---	12	DEF
8 -	anest techniek	multi-code 1	3	13	DEF
9 -	anest tekst	string[32]	---	14	DEF
10 -	plaats	code 1	4	15	DEF
11 -	plaats tekst	string[32]	---	16	DEF
12 -	anest fase	code 1	5	17	DEF
13 -	fase tekst	string[32]	---	18	DEF
14 -	uitkomst	code 1	6	19	DEF
15 -	uitkomst tekst	text[10]	---	20	DEF
16 -	major neg. uitk.	code 1	10	21	DEF
17 -	o-uitstel tekst	text[10]	---	22	DEF
18 -	correct uitstel	boolean	---	23	DEF
19 -	c-uitstel tekst	text[10]	---	22	DEF
20 -	te lang werk	boolean	---	31	DEF

```

3 - Record: factor

  Field name      type          table    question  help
1 - bijkomende fact multi-code 2    7        24        DEF
2 - chir ongev tekst text[10]      ---      25        DEF
3 - overig tekst  text[10]      ---      26        DEF
4 - categorie     code 1        8        27        DEF
5 - classificatie multi-code 2    9        28        DEF
6 - opmerkingen  text[10]      ---      29        DEF
7 - aanbevelingen text[20]      ---      30        DEF
8 - classific. tekst text[10]      ---      32        DEF
9 - potentieel risic code 1    11       33        DEF
10 - text pot.risico text[5]       ---      36        DEF
11 - prevmog incident text[10]      ---      34        DEF
12 - prevmog n-effect text[10]      ---      35        DEF
13 - inc.voorkoombaar code 1    12       37        DEF

E N D   O F   S C H E M E

```


Appendix D

Rapid & dirty Prototyping: K-model + RSO model

Main program deployed in the ISA-Dialysis trial (see 2.3.2). It combines the RSO-data model (see Appendix C) with the ISA-PL program containing the K-model (see section 2.3 and Appendix B). The program allows pausing in data entry at a predefined stage until additional information is acquired.

```
CONVERSATION SCHEME : : PILOT

1 : CALL SUBGRAPH : reset
2 : GOTO : 9
3 : IF Analysis.Menu= 3 THEN 52 ELSE NEXT
4 : IF Do you want to continue? [y/n] THEN 8 ELSE NEXT
5 : SET CONSTANT : 1 IN Analysis.Pause-mode
6 : BREAK
7 : GOTO : +2
8 : IF Analysis.Menu= 1 THEN 28 ELSE NEXT
9 : READ VALUE : Finish.Menu
10 : IF Finish.Menu= ? THEN -1 ELSE NEXT
11 : IF Finish.Menu<>Pause... THEN +4 ELSE NEXT
12 : IF Analysis.Menu= 0 THEN 9 ELSE NEXT
13 : IF Analysis.Pause-mode= 1 THEN 9 ELSE NEXT
14 : GOTO : 5
15 : IF Finish.Menu<>Finish THEN +4 ELSE NEXT
16 : IF Analysis.Menu< 2 THEN 9 ELSE NEXT
17 : SET CONSTANT : 0 IN Analysis.Pause-mode
18 : GOTO : 52
19 : IF Finish.Menu<>RSO (clinic) THEN +8 ELSE NEXT
20 : IF Analysis.Menu= 1 THEN 9 ELSE NEXT
21 : IF Analysis.Menu= 3 THEN 9 ELSE NEXT
22 : SET CONSTANT : 0 IN Analysis.Pause-mode
23 : CALL SUBGRAPH : rso-noti
24 : CALL SUBGRAPH : rso-repo
25 : CALL SUBGRAPH : rso-fact
26 : GOTO : 3
27 : IF Analysis.Menu>=2 THEN 9 ELSE NEXT
28 : SET CONSTANT : 0 IN Analysis.Pause-mode
29 : CALL SUBGRAPH : general
30 : IF Analysis.Outcome<>-1 THEN 33 ELSE NEXT
31 : CALL SUBGRAPH : observat
```

```
32 : GOTO : 46
33 : CALL SUBGRAPH : incident
34 : IF Incident.PersonsAtSpot<>yes THEN +2 ELSE NEXT
35 : CALL SUBGRAPH : person
36 : IF Incident.ObjectAtSpot<>yes THEN +2 ELSE NEXT
37 : CALL SUBGRAPH : object
38 : IF Analysis.Outcome<>1 THEN +3 ELSE NEXT
39 : MESSAGE :a-malast
40 : GOTO : 43
41 : IF Analysis.Outcome<>11 THEN +4 ELSE NEXT
42 : MESSAGE :a-prawie
43 : IF Do you want to use MORT analysis THEN +2 ELSE NEXT
44 : GOTO : +2
45 : CALL SUBGRAPH : qreview
46 : CALL SUBGRAPH : finish
47 : CALL SUBGRAPH : analiza
48 : GOTO : +3
49 : IF Omit print-out of report? [y/n] THEN 51 ELSE NEXT
50 : PRINT DOCUMENT :mort
51 : IF Analysis.Menu= 2 THEN 4 ELSE NEXT
52 : IF Analysis.Outcome= -1 THEN 54 ELSE NEXT
53 : SET CONSTANT : 1 IN SYSTEM.CurrentRecord
54 : END & SAVE OBJECT
```

ARA I Project Results

Overview of the main results of the ARA-1 technology transfer project [Spijkervet 94]

The ARA project took in 1993-1994 and turned out to be a great success in Poland. Looking at the main objectives of the project we can consider that the results of the ARA project are better than one could expected.

The first objective - development of an accident registration system for preventive purposes - has resulted in ISA-PL 1.7, which integrates the expert system elaborated in the Netherlands with the obligatory polish documents.

The second objective - to support effective co-operation in the field of work safety - has resulted in sound co-operation between:

- Ministry of Labour and Social Policy
- Trade Union Solidarnosc
- State Labour Inspectorate
- 4 research institutes and
- 11 industrial enterprises.

The group of specialists representing above mentioned institutions and enterprises has worked out a Polish terminology, a Polish pilot version of the system, course materials in polish version, and a database with the learning cases.

The third objective - transfer of the existing knowledge through the enterprises - has resulted in creating of ISA Club Poland an organisation of enterprises working with MORT and the ISA system. Moreover, the ARA project had a very big resonance in Poland what resulted in a number of activities directly related to ARA project but organised independently. A number of seminars, courses and workshops were organised in several places in Poland Universities, private consulting firms, State Labour Inspectorate. As a consequence the number of members of ISA Club Poland is significantly higher than the number of enterprises involved in the project (early 1995 more than 30 mainly big enterprises).

(this page is left blank intentionally)

Example of an Accident Analysis using ISA

In the following paragraphs, ISA is illustrated by way of an example [Koornneef 97a]. The example is of a fatal accident during maintenance of a heating installation in an apartment building. The text below provides a transcription of the investigator's typed input to the system, the system's output following analysis, and a running commentary.

The incident description, as initially available for the ISA analysis, is given in free text format in order to facilitate understanding of the subsequent incident process:

During dismantling work in a heating installation, two employees were busy cutting off a pump with a blowlamp. They got an electric shock because the body of the pump was energised. The electrician switched off the cut-off at the switch board, but the conductors were still visibly connected. At the switch board, the electrical wiring was connected wrongly: the earth line was connected to the live phase. The switch had been installed a few weeks earlier and had never been tested.

The losses due to the incident are entered first. These were:

Results: 1 fatality. Work stopped for 3 days.

The user is then required to indicate the mode of operation of the production process which is normally to be expected at the place of the incident; here it is a complete shut-down for maintenance. The assessment is used for the generation of diagnostic statements and recommendations.

Then the classification of the most relevant hazard must be entered.

Additional incident data:

Production phase: unusual conditions were applicable.

Hazard type: electricity (< 1 kV).

After the investigator has assessed and entered into ISA all the relevant issues, and before the ISA-system generates its automated analysis results, he is asked to state his professional conclusions and recommended preventive actions with respect to the incident:

Conclusions by the investigator:

1. *Leads not disconnected.*
2. *Wrongly connected electrical bridge switch.*
3. *No test if current was actually switched off.*

Preventive measures (suggested by the investigator):

1. *Check electrical installation.*
2. *Retrain employees.*
3. *Provide better supervision by infrequently performed tasks.*

In this way the professional opinion of the ISA-user can be judged against the analysis outcomes generated.

The following elements of ISA were activated when the accident assessment:

- The severity of the incident implies that a medium set of questions should be asked.
- The mode of operation of the heating system was set to a complete shut-down during the maintenance activities.
- Non-functional presence of a hazard, i.e. electricity at the heating installation, has been established.
- Maintenance was the only activity and contributed as such to the incident.
- Task Performance Errors have been scored during maintenance execution (during the previous activity on the electrical bridge switch).
- Task Performance Errors regarding the use of the cut-off switch (barrier) have been identified.
- Regarding training or re-instruction, it has been noted that both employees were competent, but did not do this task regularly.
- In such a case, the supervisor's role is to increase workers' alertness, e.g. by oral briefing, before they start the job; work site supervision was Less Than Adequate (LTA) in this respect.
- With respect to this maintenance job, the Technical Information System of the maintenance work contractor is to be assessed as LTA.
- A barrier (i.e. cut-off switch) that is not functioning conform specification, is a trap for workers.

Based on these diagnostic issues, the following analysis results were generated by the ISA-system:

REPORT OF AN ACCIDENT category: incident with serious loss potential

- With respect to the hazard, a problem has been identified with the technical information system. A review of the following issues is recommended: technical information (knowledge and/or communication), monitoring systems, data collection + analysis, Hazard Analysis Process (HAP) triggers, safety system appraisal.
- The incident appears to be linked with maintenance activities. A review of the maintenance plan & execution is called for.
- Non-functional presence has been demonstrated of a relevant hazard source.
- The presence of the identified, non-functional conditions calls for closer attention with respect to 'house keeping', permit to work/access procedures and supervision on adequate application.

- First line supervision plays a key role in work-site safety. The following supervisory problem area(s) has been found:
 - the first line supervisor did not adequately fulfil his supervisory duties. This normally indicates lack of help and training for the supervisor, lack of time, inadequate plan for responsibility transfer, LTA detection/correction of hazards, and/or performance errors. The nature of this supervision problem should be looked into.
 - Supervision on appropriate task execution during rarely occurring tasks or task circumstances is called for.
- Periodic re-instruction or re-training is recommended for tasks which are seldom performed by the employee in question.
- Frequently occurring task performance errors indicate a general management problem regarding:
 - the analysis of job/task related risks;
 - the selection, training and motivation of personnel;
 - work-site supervision of personnel.Further analysis of this problem area is called for.
- Task performance errors have been scored with respect to:
 - maintenance;
 - the use of barriers.

The categories of loss potential are based on the assessment of a combination of the human loss potential (from no injury through permanent injury or sickness to many fatalities) and the economic loss potential (from none to disastrous). This assessed loss potential is not the same as the actual losses suffered: it represents the maximum that realistically could have occurred.

ADDITIONAL DIAGNOSTIC CONCLUSIONS

- The technical information system does not meet the work process needs. Assess priorities for improvements (see MA3-MB2).
- Planned barriers did not function according to specifications. This is a serious quality assurance failure, which must be looked into immediately.

The additional diagnostic conclusions are generated as a summary of the main report to highlight safety problems that require immediate attention from higher management.

When this sort of computer-generated diagnosis was presented to those concerned in a pilot project, see section 2.3, their conclusion was that it was very relevant to the incident, and revealed a greater number of more fundamental problem areas for improvement than the recommendations from the original investigators had covered. The text above also illustrates that the report which is generated needs to be translated and made more specific before it can be presented to people, such as managers, staff or operators, who are not familiar with the concepts of the MORT tree.

(this page is left blank intentionally)

Example of an Analysis of a Clinical Accident using ISA

In the following paragraphs, ISA is illustrated by way of an example [Koorneef 97b]. The example is of an incident in a haemodialysis department of a hospital leading to the death of the patient two weeks later. The indented text below is a transcription of the text which the investigator typed into the system and/or which the system printed out.

Incident description:

During haemodialysis, apnea [a breathing arrest] occurred to a patient. Due to the subsequent asphyxia [i.e. suffocation], the patient suffered irrecoverable brain damage and died two weeks later. This patient's initial physical condition was judged as poor on the referral form for haemodialysis treatment from another hospital department.

The incident description, as initially available for the ISA analysis, is given above in free text format in order to provide understanding of the incident process afterwards. In the current hospital-version of ISA, additional questions regarding the hospital's internal assessment of the incident are asked first. In this case, the incident was considered as one with a major negative outcome, in which the occurrence of apnea during haemodialysis treatment could not have been prevented, but in which the effective avoidance of asphyxia is quite feasible. The incident was classified as a patient-related 'complication'.

The following conditions contributed to the delay in detection of the apnea:

- *standard operating procedures for haemodialysis treatment do not include the monitoring of the patient's breathing*
- *on a previous occasion this patient also felt sick, but the haemodialysis process went all right*
- *relatives present and clinical haemodialysis personnel had a difference in judgement about the actual condition of the patient: the relatives claimed that more attentive care of this patient was imperative this time*
- *the patient had been admitted to the haemodialysis department late in the afternoon during the day shift, the incident occurred during the evening shift.*

These were elicited by the question and the answer interaction in the software, based on the MORT chart logic.

The losses due to the incident were described as:

- *The patient had to be forwarded to the Intensive Care Unit instead of being returned to the sending department.*
- *One fatality: the patient died after two weeks.*
- *A formal complaint procedure initiated by relatives.*

After the local investigator had assessed and entered into ISA all the relevant issues, he was asked to state his professional conclusions and recommend preventive actions with respect to the incident, before the ISA-system generated the automated analysis results:

Conclusions by the investigator included the following:

1. *Two medical doctors were involved. The one from the sending department judged that the patient could be sent to the haemodialysis department for treatment, the other one from the haemodialysis department, accepted his colleague's judgement without questioning.*
2. *The handing over of the patient by the day shift to the evening shift was inadequate with respect to the patient's medical history and physical condition*
3. *The response by the haemodialysis department staff to alarming signals from relatives present was inadequate.*

Preventive measures as suggested by the investigator were:

1. *This class of patients should be treated only during the day shift (when more staff is present).*
2. *High-risk patients may better be treated in an ICU: a strategy for haemodialysis treatment of critical care patients must be developed.*
3. *Re-instruction/retraining of haemodialysis staff regarding communication with relatives during haemodialysis is necessary.*

This professional opinion of the ISA-user can be judged against the analysis outcomes generated by ISA.

The following elements of ISA were activated when the accident was coded:

- The severity of the incident implies that a medium set of questions should be asked.
- Planned procedures and other barriers do not provide sufficient patient safety.
- Non-functional presence of a hazard, i.e. disturbances in the respiration system of the patient, has been established.
- Supervisory problems have been pinpointed regarding task assignment and direct supervision.
- Task Performance Errors regarding the use of procedural barriers have been identified.

Based on these diagnostic issues, the following analysis results were generated by the ISA-system:

REPORT OF AN ACCIDENT category: incident with serious loss potential

- The full set of (planned) protective safety and control barriers needs to be reviewed: together, these measures may not provide adequate protection to vulnerable targets even if they are applied adequately/as intended.
- Non-functional presence has been demonstrated of:

- a relevant hazard source
- The presence of the identified, non-functional conditions calls for closer attention with respect to control of energy flows, 'house keeping', permit to access/work procedures and supervision on adequate application of presence control systems.
- Wrongly, an insufficiently qualified person has been assigned to a task. Further investigation is called for (see SD5-c14).
- First line supervision plays a key role in worksite safety. The following supervisory problem area(s) has been found in relation to at least one relevant employee:
 - the first line supervisor did not adequately take his responsibility. This normally indicates lack of help and training for the supervisor, lack of time, inadequate plan for responsibility transfer, less than adequate detection/correction of hazards, and/or performance errors. The nature of this supervision problem should be looked into (see SD5).
- Frequently occurring task performance errors indicate a general management problem regarding:
 - the analysis of job/task related risks,
 - the selection, training and motivation of personnel,
 - worksite supervision of personnel.Further analysis if this problem area is called for.
- Task performance errors have been scored with respect to:
 - the use of barriers

Additional diagnostic conclusions are generated as a summary of the main report to highlight safety problems which require immediate attention from higher management.

ADDITIONAL DIAGNOSTIC CONCLUSIONS

- Planned barriers are not aiming at sufficient protection of vulnerable targets. This indicates a systemic problem in the overall risk assessment system (see MA3-MB3-a2-b7).
- At least one barrier was not provided where possible, thus, couldn't be used. Appropriate management must assess the risk of failure to provide barriers and either assure that required barriers are in place or accept the risk (see SD5, MB4).
- At least one planned barrier didn't function conform specifications. This is a serious quality assurance failure which must be looked into immediately (see SD3, SD4, MB4).

When this sort of computer-generated diagnosis was presented to those concerned in the pilot project, their conclusion was that it was very relevant to the incident, and revealed a greater number of more fundamental problem areas for improvement than the recommendations from the original investigators had covered. A key problem in this particular incident is the lack of a systematic risk classification of patients with respect to haemodialysis treatment facilities, a procedure which is common practice regarding patients before a surgical operation. A routine haemodialysis department is not adequately suited for treatment of critical care patients.

The text above also illustrates that the report which is generated needs to be translated and made more specific before it can be presented to people, such as managers, staff or operators, who are not familiar with the concepts of the MORT tree.

(this page is left blank intentionally)

Appendix H

On Case-Based Reasoning

This appendix is an informative abstract based on [Aamodt 94]. Case-based Reasoning (CBR) is one of the methodologies from the Artificial Intelligence domain [Aamodt 94]. In fact, CBR as a generic approach of reasoning by analogy, embraces a variety of methodologies, one which itself is called 'Case-based Reasoning'. Case-Based Reasoning is both a paradigm for computer-based problem solvers and a model of human cognition. The central idea is that the problem solver reuses the solution from some past case to solve a current problem. In order to this, we must store previous experiences (cases) into a memory. When a new problem arises, we need to retrieve similar experiences about similar situations from memory and reuse solutions in the context of the new problem: the CBR cycle [Aamodt 94] starts. Four basic steps characterise the CBR cycle. Starting with a new problem, a case base is searched for *retrieval* of similar cases of which the best matching become candidates for *reuse*. The solution that was applied in the old case then is *revised* and adapted to the new situation. The new problem and the solution applied is *retained* in the case base for future reuse.

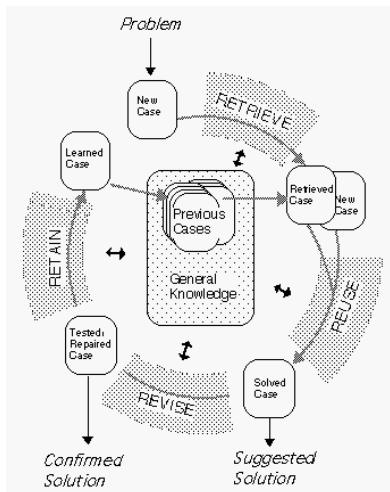


Figure H-1: CBR cycle [Aamodt 94]

An essential step in CBR is the selection of cases from the ones retrieved from the case base. The *similarity* assessment must ensure that only the most relevant cases are chosen.

As old cases usually differ at least slightly from the new problem situation, the old solution may need to be *adapted* in order to fit the new problem. The effectiveness of the adapted solution needs to be *verified*. The CBR system actually *learns*, i.e. acquires knowledge, when the new solution is combined with the problem representation to form a new case and is added to the case base [Bergmann 98]. Core problems addressed by CBR research can be grouped into five areas: knowledge representation in a case, and methods for respectively Retrieval, Reuse, Revision and Retainment. The CBR paradigm covers a range of different methods for organising, retrieving, utilising and indexing the knowledge retained in past cases. The issue of what makes a problem a case and how the case will be stored in the case base is another key issue for the case-based reasoner:

"The representation problem in CBR is primarily the problem of deciding what to store in a case, finding an appropriate structure for describing case contents, and deciding how the case memory should be organised and indexed for effective retrieval and reuse." [Aamodt 94]

Two approaches have become influential: the Dynamic Memory Model of Schank and Kolodner [Kolodner 93], and the category-exemplar model of Porter and Bareiss [Bareiss 89; Porter 90]. In the DMM, the case memory in the model is a hierarchical structure of what is called 'episodic memory organisation packets' (E-MOPs [Kolodner 83]), also referred to as generalised episodes [Koton 89].

"The basic idea is to organise specific cases which share similar properties under a more general structure (a generalised episode). A generalised episode (GE) contains three different types of objects: *Norms*, *cases* and *indices*. Norms are features common to all cases indexed under a GE. Indices are features which discriminate between a GE's cases. An index may point to a more specific generalised episode, or directly to a case. An index is composed of two terms: An index name and an index value." [Aamodt 94]

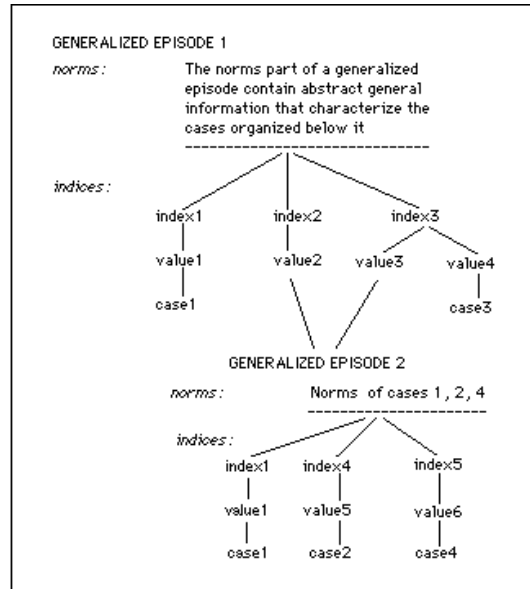


Figure H-2: Dynamic Memory Model structure of cases and generalised episodes

Figure H-2 [Aamodt 94] illustrates this structure. An index value may only point to a case or to a GE. Index-value pairs point from a GE to another GE or to a case. Indexing may be redundant since multiple paths may exist to a particular GE or case.

In the Category & Exemplar Model [Bareiss 89] a case – also referred to as *exemplar* – is linked to a category by a set of features describing a case's membership to that category. Any attempt to generalise a set of cases should - if attempted at all - be done very cautiously. This fundamental view of concept representation forms the basis for this memory model.

"The case memory is embedded in a network structure of *categories*, *cases*, and *index pointers*. Each case is associated with a category. An index may point to a case or a category. The indices are of three kinds. Feature links pointing from problem descriptors (features) to cases or categories (called reminders). Case links pointing from categories to its associated cases (called exemplar links). And difference links pointing from cases to the neighbour cases that only differ in one or a small number of features. A feature is, generally, described by a name and a value. A category's exemplars are sorted according to their degree of prototypicality in the category." [Aamodt 94]

Figure H-3 [Aamodt 94] illustrates a part of this memory structure, i.e. the linking of features and cases (exemplars) to categories. The unnamed indices are reminders from features to a category. Within this memory, categories are inter-linked within a semantic network that represents a background of general domain knowledge. Bareiss distinguishes six fundamental issues of Case-based Reasoning that a CBR system must cope with: case indexing (for retrieval), similarity assessment, selection of new cases for retention, indexing for learning, acquisition of additional domain knowledge in similarity assessment, and generalisation if any.

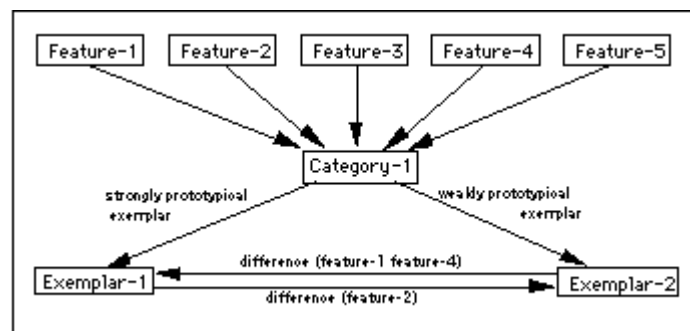


Figure H-3: Structure of Categories, Features and Exemplars

The dependence on domain knowledge in the Category & Exemplar Model makes this approach less attractive for near-miss management systems.

Nearest-Neighbour Matching

Within the DMM approach is in particular the 'nearest-neighbour matching promising for application for a lessons-learned-from-incidents' because the similarity assessment is simply done numerically. Supposing that a case C can be represented by a set of descrip-

tors ($D_1 \dots D_k$), each of these having a value V_{CD} and different set of norms N may be applicable, depending on the context of a new problem.

	<i>Norm-1</i>	<i>Norm-Z</i>	<i>Case-1</i>	$sim(C_1-C_{new})$	<i>Case-n</i>	$sim(C_n-C_{new})$	<i>New Case</i>
<i>Descriptor 1</i>	W_{11}	W_{1z}	V_{C1D1}	S_{11}	V_{CnD1}	S_{n1}	$Value_{CnewD1}$
<i>Descriptor 2</i>	W_{21}	W_{2z}	V_{C1D2}	S_{12}	V_{CnD2}	S_{n2}	$Value_{CnewD2}$
<i>Descriptor k</i>	W_{k1}	W_{kz}	V_{C1Dk}	S_{1k}	V_{CnDk}	S_{nk}	$Value_{CnewDk}$
	$\sum_{i=1}^{i=k} W_{i,1}$	$\sum_{i=1}^{i=k} W_{i,z}$		$\sum_{i=1}^{i=k} W_{i,z} \times S_{1,i}$		$\sum_{i=1}^{i=k} W_{i,z} \times S_{n,i}$	

For each descriptor and each of the candidate cases, a similarity factor S_{ij} is estimated from their respective values V_{ij} , with $0 \leq S_{ij} \leq 1$ [0 = no similarity or missing value; 1 = very similar]. Given the context of the new case, one particular set of norms N_z with importance weighting factors $w_{i,z}$ applies. A typical algorithm for calculating nearest-neighbour matching is the one reported in [Kolodner 93] where w is the importance weighting of a feature (or slot), S_{ij} is the similarity factor for descriptor i of the retrieved case j :

$$\frac{\sum_{i=1}^k w_i \times S_{i,j}}{\sum_{i=1}^k w_i} \quad [1]$$

The case with the highest (normalised) match score is the case that matches the new problem best. Pitfalls exist, however, such as the use of descriptors that are not relevant to the problem, or by giving too much or to less weight to some descriptors. The underlying DMM approach can lead to excessive retrieval effort when many indices can be applied to a large set of cases [Bareiss 89, Watson 94]. By means of templates for case retrieval and limiting descriptors by feature vocabulary [Bergmann 98] the case indexing remains manageable.

Postscript

CBR has advantages over knowledge-based AI-systems, among others because [Watson 94]:

- CBR does not require an explicit domain model and so knowledge elicitation becomes a task of gathering case histories
- CBR systems can learn by acquiring new knowledge as cases thus making maintenance easier

This CBR cycle currently rarely occurs without human intervention. For example many CBR tools act primarily as case retrieval and reuse systems. Case revision (i.e. adaptation) often being undertaken by managers of the case base. However, it should not be viewed as weakness of CBR that it encourages human collaboration in decision support.

Appendix I

Patient Throughput Phases Inventory

The table on the following pages was used to identify distinct stages in the processes of cardiac surgery, and for derivation of a detailed coding table in SINS-1 for specification of the phase of patient care in which an operational surprise was detected.

Table 1-1: Patient Throughput Phases Inventory

Stap: naar:	Eigenaar:	Actie:	Transfer naar:	Risico:	Kanttekening:
Fase 0: Klachtontwikkeling					
0.01	0.02	Patiënt	pijn op de borst	Huisarts	stoere bink: gaat niet naar arts
0.02	0.03	Huisarts	verwijst patiënt	Cardioloog	geen verwijzing (fout-negatief?)
0.03	0.04	Cardioloog	stelt diagnose ((inspannings-EKG hartspter-echo)	Huisarts? / Cardioloog	fout negatief / fout positief
0.04	0.05	Cardioloog	medicatie	Huisarts? / Cardioloog	bloedverdunner <=> trombose
0.05	0.06	Patiënt	pijn op de borst/aanval	Cardioloog	...gaat niet naar arts
0.06	1.01	Cardioloog	catheterisatie: angiogram	Cardioloog	patiënt krijgt aanval.../ fout-diagnose
Fase 1: Aanmelding					
1.01	1.02	Cardioloog	meldt patiënt aan	CPCA-chirurg	patiënt wordt niet aangemeld
1.02	1.03	CPCA-chirurg + cardioloog	intake / 'zien': patiënt operabel?	Cardioloog	patiënt-classificatie (niet, wachtlijst, voorrang, spoed)
1.03	1.04	CPCA-AMC	zet patiënt op wachtlijst	Wachtlijst	indeling niet conform classificatie
1.04	2.01, 1.05	Patiënt	wacht oproep af	Cardioloog / huisarts?	patiënt krijgt voortijdig nieuwe aanval...
1.05	1.06	Patiënt	krijgt voortijdig aanval	0611	hulpverlening wordt niet / te laat ingeschakeld
1.06	1.07	Ambulance	EHBO & transport van patiënt	EHBO-EHH	te laat / niet naar hartecentrum
1.07	3.04	Triage-arts (EHBO)	beslist tot EHH-actie voor patiënt	Cardioloog (voorwacht)	prioriteiten in eerste hulp onjuist ingeschat
Fase 2: Oproep					
2.01	2.02, 2.03	chef-chirurg	beslist tot inroosting/oproep	CPCA-AMC	afstemming schaarse middelen (mensen / bedden)
2.02	2.04	CPCA-AMC	roept patiënt op voor poli-opname	telefonisch => Patiënt	oproep komt niet goed over
2.03	2.05	CPCA-AMC	roept patiënt op voor opname	telefonisch => Patiënt	oproep komt niet goed over
2.04	2.05	Patiënt	komt opdagen bij AMC-poli	CPCA-arts	opname-keuring: beoordeling opereerbaarheid onjuist
2.05	3.01	Patiënt	bereidt zich voor op opname	Patiënt	houdt zich niet aan de instructies / complicatie
					medische pre-op amnese
					locatie manager

Stap: naar:	Eigenaar:	Actie:	Transfer naar:	Risico:	Kanttekening:
Fase 3: Opname					
3.01	4.01, 3.02	Patiënt	Verpleging G3 Zuid	opname: voorbereiding patiënt onvoldoende/vertraagd	medische en verpleegkundige pre-op amnese
3.02	4.01	Cardioloog	CCU G3 Noord	pre-operatieve NTG-behandeling... 1 dag voor operatie?	
3.03	4.01, 3.02	Cardioloog	CPCA-arts	veelal noodprocedure tijdens onderzoek van de patiënt	
3.04	3.03	Cardioloog			
3.05	3.02	Cardioloog			
3.06	4.01	Cardioloog/ CPCA-arts	CPCA-arts	via EHHO: film moet nog gemaakt worden <=> 30'	CPCA-arts: voorwacht
3.07	3.08	Cardioloog	Cardioloog		
3.08	3.09	Cardioloog	Cardioloog	EHH-patiënt is te slecht / onjuiste diagnose	
3.09	4.01	Cardioloog/ CPCA-arts	CPCA-arts		
Fase 4: Operatie					
4.01	4.02	Anesthesist	CPCA-arts/ cardioloog	anesthesie-risico's, bijv. tijdens intubatie	
4.02	4.03	CPCA-arts	Perfusionist	uitvoeringsfout / complicatie	
4.03	4.04, 4.05	CPCA-arts	CPCA-arts	andere ingreep dan gepland / fout / complicatie	
4.04	4.05	CPCA-arts	Anesthesist	inwendig leegbloeden	
4.05	4.06	CPCA-arts	OK-team	uitvoeringsfout / complicatie	
4.06	4.07	Anesthesist	Anesthesie-team	uitvoeringsfout / complicatie	
4.07	5.01	Anesthesie-team + CPCA-assistent	IC-arts (intensivist)	monitoring patiënt tijdens 15' transport afwezig	
Fase 5: Post-operatieve zorg					
5.01	5.02	Verpleging IC-G3 Noord	Verpleging IC-G3 Noord	beademing & bewaking: complicatie? <=> Interventie?!	
5.02	5.04, 5.03	Verpleging IC-G3 Noord	IC-arts	OK-patiënt is 10 uur onder narcose: complicatie?	

Stap: naar:	Eigenaar:	Actie:	Transfer naar:	Risico:	Kanttekening:
5.03	5.04, 4.01 IC-arts	diagnose van complicatie	CPCA-arts	onjuist (bijstelling medicatie of her-operatie)	
5.04	Verpleging IC- G3 Noord/ IC- arts	ontkoppelt mechanische ventilatie	IC-arts/CPCA- arts	patiënt komt niet bij uit narcose / complicatie	
5.05	IC-arts/ CPCA- arts	stuurt patiënt naar G3 Zuid - Medium Care	Verpleging G3 Zuid	meestal binnen 24 uur (gemiddeld na 48 uur)	
5.06	Verpleging IC- G3 Noord	verpleegkundige overdracht	Verpleging G3 Zuid	mondeling + schriftelijk (verpleegkundig dossier)	
5.07	IC-arts	overdracht medisch dossier	CPCA-arts	Schriftelijk	
Fase 6: Verpleging					
6.01	Verpleging G3 Zuid	neemt patiënt post-operatief op in MCU	Verpleging G3 Zuid	uitvoeringsfout / complicatie	
6.02	Verpleging G3 6.04, 6.05 Zuid	verzorgt patiënt	Verpleging G3 Zuid	uitvoeringsfout / complicatie	
6.03	CPCA-arts assistent	plaatst patiënt uit MC naar G3 Zuid	Verpleging G3 Zuid	onjuiste diagnose n.a.v. functie-onderzoek, lab-tests	
6.04	5.01, 5.02 CPCA-arts	plaatst patiënt terug naar IC	Verpleging IC- G3 Noord/IC-arts	geen IC-bed beschikbaar...	
6.05	6.06 CPCA-arts assistent	plaatst patiënt terug naar CCU	Verpleging CCU-G3 Noord	post-operatief: bij hartritestoormissen, ischemie	
6.06	5.05 Verpleging CCU-G3 Noord	bewaakt elke patiënt...	CPCA-arts		
6.07	6.08 Verpleging G3 Zuid	verzorgt patiënt	Verpleging G3 Zuid	uitvoeringsfout / complicatie	
6.08	6.09 CPCA-arts/ cardioloog	neemt ontslagbesluit	Verpleging G3 Zuid	te vroeg / ten onrechte?	
6.09	6.10 Verpleging G3 Zuid	bereidt patiënt voor op vertrek	Patiënt		
6.10	6.11 Patiënt	ontslag: vertrekt naar eigen (zieken)huis...	Patiënt		
6.11	7.01 CPCA-arts/ cardioloog	stuurt ontslagbrief naar cardioloog/huisarts	Huisarts/ cardioloog		
Fase 7: Ontslag					
7.01	7.02 Cardioloog	wikkelt nazorg/nacontrole af...	Patiënt/huisarts		
7.02	Patiënt	slikt voor altijd medicijnen...?			

Reported Significant Observation (RSO) Study CAPU

This appendix contains the description, protocol and forms (in Dutch) for the RSO study in the cardiac surgery unit (CAPU) as part of the PIERCE project (see section 6.5).

Doel:

1. Het identificeren van gevaarlijke omstandigheden en praktijken op basis van persoonlijke ervaring of historische informatie binnen het Capu-cluster, hetgeen o.a. moet leiden tot explicitering van afbreukscenario's in alle relevante trajecten die patiënten tijdens ziekenhuisopname doorlopen.
2. Het onderzoeken van de mogelijkheden voor risicomanagement binnen het cluster Capu (MCU en verpleegafdeling op G3-zuid, ICU en CCU op G3-noord).

Toepassingsbereik:

Overall binnen het cluster waar medewerkers c.q. waarnemers (historische) ervaring hebben opgebouwd.

Werkwijze:

Maximaal één keer per zes maanden wordt aan ervaren medewerkers c.q. waarnemers gevraagd mee te werken aan een interview dan wel een vragenlijst in te vullen. Zij lopen daarbij **niet** het risico over het gerapporteerde ter verantwoording te worden geroepen daar alle verkregen gegevens strikt vertrouwelijk worden behandeld en uitsluitend worden gebruikt in het kader van het PIERCE-project (PIERCE = Patient Incident Evaluation and Registration for Cardiopulmonary Environments). Hierdoor kan nauwkeurige rapportage van anderszins niet gemelde, maar wel geobserveerde "bijna-ongevallen" worden bevorderd.

Stappenplan:

1. Behoeftebepaling

Zie aanleiding PIERCE-project.

Beschrijving van de trajecten die een Capu-patiënt tijdens ziekenhuisopname doorloopt, leert het volgende:

- a) Een patiënt doorloopt verschillende trajecten die vallen binnen de "span of control" van onderscheidbare probleemeigenaren binnen het AMC.

- b) De bestaande complicatieregistratie heeft uitsluitend betrekking op patiënten die op de OK-tafel zijn beland.
- c) Deze complicatieregistratie bevat geen informatie over de omstandigheden waaronder de problemen zich voordeden.
- d) Een RSO-studie kan bijdragen aan explicitering van (afwijkings)scenario's voor de realisatie van afbreukrisico's ten koste van Capu-patiënten (direct en/of door verminderde beschikbaarheid van voorzieningen).

2. Formatie RSO-groep

Vaste kern (voorstel):

- Ruud de Graaf (coördinator AMC/CAPU)
- Floor Koornneef (veiligheidskundige TUD)
- Jaap Kloek / Bas de Mol (stafleden AMC/CAPU)
- een ervaren staflid uit elk van de cluster-afdelingen op G3-zuid en G3-noord.

Zo nodig op ad-hoc basis medewerkers van aanverwandte AMC-afdelingen (Eerste Hart-Hulp, F3-zuid, OK's, etc.).

3. Vaststelling taak RSO-groep

- a) Vaststelling van scope van de studie en mate van vrijheid bij de uitvoering.
- b) Analyse, beoordeling en validering van de verzamelde data.
- c) Implementatie van interventie maatregelen en terugkoppeling naar de betrokken medewerkers c.q. afdelingen.
- d) Formatie stuurgroep voor het beheer van de verzamelde en gevalideerde meldingen.

4. Vaststelling studiescope

1. specifiek onderwerp,
2. algemeen onderwerp,
3. vrij spel.

Algemeen onderwerp: voorvallen en/of aangetroffen omstandigheden die verstorend doorwerken en de patiënt mogelijk bedreigen, bijvoorbeeld:

- patiënt valt uit bed;
- patiënt krijgt verkeerde medicijnen;
- chirurgische mes zorgt tijdens de operatie voor problemen;
- chirurg heeft een verkeerde film bekeken.

Later in het project, scoop toespitsen op een specifieke groep CAPU-patiënten, bijvoorbeeld patiënten met hartklepprothese.

Primair de aandacht richten op problemen in de bedrijfsvoering.

5. Methode van data-acquisitie

Enquête.

Normaliter is een RSO-enquête anoniem, hetgeen in het kader van dit project ook nodig geacht wordt. (Of kunnen we een (indirecte) vorm van respondent-identification hanteren, zodat het valideren van data beter mogelijk wordt?)

Nadeel van een enquête is dat geënquêteerden de vragen foutief kunnen interpreteren zonder dat de onderzoekers dit zullen opmerken.

Nadeel van een open-eind-interview is dat het veel personele middelen vraagt en het in dit stadium te bewerkelijk is.

6. Enquêteformulier

- 1) Ontwerp een enquêteformulier.
- 2) Controleer formulier op leesbaarheid / helderheid (voorleggen aan RSO-groep en enkele medewerkers van de diverse afdelingen).
- 3) Verspreid formulier onder (ervaren) medewerkers / waarnemers. Wijd desgewenst een bespreking aan het invullen (2*2 per persoon / 30-60 minuten).
- 4) Vraag alle deelnemers om terugzending van twee formulieren die betrekking hebben op ongevallen en/of abnormale gebeurtenissen die nog net goed zijn afgelopen en twee formulieren die betrekking hebben op voorvallen die wèl tot schade hebben geleid.

7. Verzamelen resultaten

8. Analyseren resultaten

Voer een snelle analyse (*cursory hazard analysis*) uit zodra een RSO-melding binnenkomt, zodat onmiddellijk aandachtvragende gevaren c.q. praktijken geïdentificeerd kunnen worden.

Rechtstreeks voordeel van zo'n RSO-actie is de ontdekking van operationele afbreukrisico-problemen welke op een andere wijze (monitoring, auditing) niet worden onderkend.

9. Data-validatie

Normaliter moet de RSO-groep alle data valideren. Dit kan onder meer op basis van:

- persoonlijke kennis van de situatie;
- discussie met een verantwoordelijk afdelingshoofd;
- door bevestiging van lokale ((getuige-)deskundigen).

Als **snelle actie nodig lijkt**, moeten de data in de betreffende RSO-melding eerst worden gevalideerd:

1. Is de werkelijke situatie zoals beschreven?
2. Is de beschreven situatie onder welke omstandigheden denkbaar?
3. Bestaat de gevaarlijke situatie nog?
4. Kan deze terugkeren?
5. Zo ja, wat is dan het beste c.q. optimale pakket van acties en maatregelen?

10. Data zijn niet valide

Als van een melding de gegevens niet valide zijn wordt in het op veiligheid en risicomanagement gerichte PIERCE-project eerst nagegaan wat hiervan de oorzaak is.

11. Data zijn valide

Analyseer data zorgvuldig ten aanzien van

- patronen;
- systematische problemen;
- verborgen problemen;
- positieve aspecten.

Verdoe niet teveel tijd aan synthese van ongevalsscenario's door RSO-meldingen te combineren. Aan de andere kant, als een gemeld voorval een ernstig gevolg met zich meebrengt of mee kan brengen, zoek dan naar mechanismen die tot het specifieke voorval hebben geleid.

12. Wanneer is er sprake van een systematisch probleem?

Een systematisch probleem bestaat onder andere als een particuliere oplossing ("fix") wordt toegepast in de gemelde, specifieke situatie maar herhaling van het voorval in soortgelijke situaties hierdoor niet wordt voorkomen.

13. RSO-melding over een "goede" situatie

N.B. Achter een RSO-melding over een veronderstelde "goede" situatie kan een verborgen probleem schuil gaan.

Een RSO-melding die ook in werkelijkheid "goed" is, kan zinvol zijn voor de bepaling van elementen in het vigerende RM-programma die

- a) niet opgeschort of geëxtensiveerd moeten worden;
- b) ook kunnen worden toegepast in andere procesclusters en/of
- c) wellicht verwerkt moeten worden in toekomstige activiteiten en/of ontwerpen.

14. RSO-melding met betrekking tot niet-gerelateerd onderwerp

Onderneem ook actie op een RSO-melding als deze betrekking heeft op een niet-gerelateerd onderwerp, bijvoorbeeld een klacht.

15. Onbruikbare meldingen

Verwijder onbruikbare meldingen.

16. Het op de hoogte houden van de medewerkers

Houd medewerkers op de hoogte (ook wanneer het pertinente details betreft) door middel van bijeenkomsten, complicatiebesprekingen, bulletins, et cetera teneinde de betrokkenheid op peil te houden.

17. Gebruik sleutel- of trefwoorden

Voorzie gevalideerde meldingen van sleutel- of trefwoorden. Deze kunnen voor de medewerkers bruikbaar worden gemaakt door deze na de eerste ronde RSO-meldingen ter markering uit te zetten onder de lokale lijn- of stafmedewerkers.

18. RSO-bestand

Neem gevalideerde meldingen op in een toegankelijk bestand.

REPORTED SIGNIFICANT OBSERVATION (RSO) STUDIE AMC/CAPU

In het kader van het PIERCE-project (PIERCE = Patient Incident Evaluation and Registration for Cardiopulmonary Environments), zijn wij geïnteresseerd in het verkrijgen van informatie over voorvallen en/of aangetroffen omstandigheden die verstorend doorwerken en de patiënt gedurende zijn of haar ziekenhuisverblijf mogelijk bedreigen, bijvoorbeeld:

- patiënt valt uit bed;
- patiënt krijgt verkeerde medicijnen;
- chirurgische mes zorgt tijdens de operatie voor problemen;
- swan-ganz ballon is lek/stuk
- chirurg heeft een verkeerde film bekeken.

In dit verband doen wij een beroep op uw deskundigheid en ervaring wat betreft de dagelijkse praktijk met de u ter beschikking staande middelen (apparatuur, hulpmiddelen en protocollen) alsmede wat betreft communicatie en logistiek.

Eén aspect van deze studie betreft het verkrijgen van persoonlijke beschrijvingen van zelf geobserveerde gebeurtenissen uit uw dagelijkse praktijk.

Deze beschrijvingen dienen enerzijds - groene formulieren - te duiden op een **geslaagd ontwerp** van apparatuur, hulpmiddelen of protocollen, waardoor de patiëntenzorg vergemakkelijkt of beter gewaarborgd wordt, en anderzijds - gele formulieren - op een **minder geslaagd of mislukt ontwerp**.

Op de volgende bladzijden vragen wij u hiervan 2 keer 2 voorbeelden te geven.

Ieder voorbeeld moet bestaan uit een korte en chronologische beschrijving van een voorval dat aan de volgende voorwaarden voldoet:

- Het gebeurde op een bepaalde plaats en op een bepaald tijdstip en u was ervan getuige.
- Het was een uitvloeisel van het ontwerp van apparatuur, hulpmiddel of protocol dat volgens u bijzonder doeltreffend c.q. weinig doeltreffend is;
- Het droeg bij aan een positief resultaat of het voorkomen van een probleem of het droeg bij aan een abnormale probleemsituatie of een situatie die tot een probleem had kunnen leiden.

Inleveren:

U wordt verzocht de 4 formulieren binnen 2 weken bij het secretariaat G4 - CAPU in te leveren t.n.v. Ruud de Graaf / PIERCE-project.

U wordt over de bevindingen van de studie op de hoogte gebracht.

Dank voor uw medewerking.

Floor Koornneef / Ruud de Graaf (sein 59748)

[Groen Formulier]

Denk aan de meest recente situatie waarin u een activiteit observeerde waarbij het gemakkelijk bleek om op een doeltreffende en foutloze manier

- een apparaat te bedienen;
- een hulpmiddel te hanteren;
- conform een protocol te handelen

1. Waar en wanneer (bij benadering) trad deze situatie op?

- afdeling / ruimte:
- datum:
- tijd:

2. Op welk apparaat, hulpmiddel of protocol had dit betrekking?

3. Geef een korte beschrijving van de situatie ten tijde van uw waarneming.

Voorbeeld:

- Het apparaat moest worden ingesteld.
 - De disposable werd aan het apparaat aangesloten.
 - De protocollaire stappen werden afgewerkt.
- Een abnormale situatie deed zich voor, namelijk ...

4. Wat gebeurde er precies? (Gebruik zo nodig ook de ommezijde van dit formulier)

5. Waarom voert u dit voorval op als een voorbeeld van een bijzonder gemakkelijk uitgevoerde activiteit?

6. Wat had u in deze situatie qua gang van zaken verwacht van een minder goed ontworpen apparaat, hulpmiddel of protocol?

Voorbeelden:

- De activiteit is dan moeilijker uit te voeren.
- Er bestaat dan een grotere kans op foutief handelen of materiaalschade.

[Geel Formulier]

Denk aan de **meest recente situatie** waarin u een activiteit observeerde waarbij het **niet gemakkelijk** bleek om op een doeltreffende en foutloze manier

- een apparaat te bedienen;
- een hulpmiddel te hanteren;
- conform een protocol te handelen

1. Waar en wanneer (bij benadering) trad deze situatie op?

- afdeling / ruimte:
- datum:
- tijd:

2. Op welk apparaat, hulpmiddel of protocol had dit betrekking?

3. Geef een korte beschrijving van de situatie ten tijde van uw waarneming.

Voorbeeld:

- Het apparaat moest worden ingesteld.
 - De disposable werd aan het apparaat aangesloten.
 - De protocollaire stappen werden afgewerkt.
- Een abnormale situatie deed zich voor, namelijk ...

4. Wat gebeurde er precies? (Gebruik zo nodig ook de ommezijde van dit formulier)

5. Waarom voert u dit voorval op als een voorbeeld van een niet gemakkelijk uitgevoerde activiteit?

6. Wat had u in deze situatie qua gang van zaken verwacht van een beter ontworpen apparaat, hulpmiddel of protocol?

Voorbeelden:

- De activiteit is dan beter uit te voeren.
- Er bestaat dan een kleinere kans op foutief handelen of materiaalschade.

(this page is left blank intentionally)

Appendix K

Interview Protocol (Safety Culture)

This appendix contains the interview protocol (in Dutch) for the safety culture assessment study as part of the PIERCE project (see section 6.6).

INTERVIEW PROTOCOL AMC / ISA

GEÏNTERVIEWDE:

DATUM:

INTERVIEWER:

LENGTE INTERVIEW:

BIJLAGE(N)

ACHTERGROND ONDERZOEK

Zoals u waarschijnlijk nog niet weet, zal er binnenkort, op een viertal afdelingen binnen het AMC, een begin worden gemaakt met de invoer van een nieuw incidentregistratiesysteem. Met behulp van dit systeem kan men beter de oorzaak en achtergronden van ongevallen nagaan.

Als onderdeel van deze invoering willen wij het draagvlak voor en het effect van een dergelijk besluit op de organisatie in kaart brengen. Wij willen ons hiervoor een beeld vormen van de veiligheidscultuur binnen uw organisatie. Hiervoor zullen twee ronden van ongeveer twintig interviews met u en collega's van u gehouden worden. De eerste ronde vindt nu plaats, vóór de introductie van het systeem. De tweede ronde vindt vervolgens over ongeveer een jaar plaats, als men wat ervaring met het systeem heeft opgedaan. De resultaten van deze interviews stellen ons in staat het effect van deze invoering op de veiligheidscultuur van uw organisatie te bestuderen.

Ik wijs er met nadruk op, dat uw antwoorden vertrouwelijk zullen worden behandeld. Gedurende het interview zal ik uw antwoorden globaal vastleggen. Deze vastlegging dient als geheugensteun en zal nergens op deze wijze worden gerapporteerd of verspreid. Als u er prijs op stelt, kan ik u de uitgewerkte aantekeningen toesturen. Indien wij over dit onderzoek gaan publiceren, zullen hooguit een aantal voorbeelden ter illustratie worden aangehaald. Niemand zal hierin herkenbaar zijn noch zal een bepaalde gebeurtenis tot een bepaalde persoon herleid kunnen worden.

Het gehele interview duurt ongeveer één uur.

Het feitelijke interview begint nu ...

I. PERSOONLIJKE ACHTERGROND

1. Welke opleiding heeft u gevolgd (en afgerond)?
 - vooropleiding (m.n. vervolgopleiding)
 - cursussen en/of bijscholing
2. Wat is uw functie/positie in het AMC en waar bent u verantwoordelijk voor? Hoe lang werkt u in het AMC? Op deze afdeling? Welke zijn uw belangrijkste taken en activiteiten?

II. VEILIGHEIDSBEWUSTZIJN

3. Wat houdt 'Veiligheid in het ziekenhuis' voor u in? Wat is uw rol hierbij?
4. Hoe wordt er, volgens u, door uw collega's gedacht over 'Veiligheid in het ziekenhuis'?
5. Vraagt de afdeling waar u werkt nog extra/speciale aandacht voor deze veiligheid?
6. Zijn er ook speciale regels of procedures opgesteld met betrekking tot 'Veiligheid in het ziekenhuis'?
 - Zo ja, kent u deze regels? Bent erin getraind of erover voorgelicht? Kunnen deze regels ergens eenvoudig worden ingezien?
 - Worden ze ook nageleefd?
 - Zo nee, vindt u dit een gebrek?

Let op: Maak een onderscheid tussen "ARBO" en patiëntveiligheid.

III. INCIDENTEN

7. Welke was de **laatste**, onveilige situatie die u bij u op de afdeling heeft meegemaakt? Wat was uw rol hierbij?
Wat is ermee gebeurd? Wat zou er mee gedaan kunnen worden?
8. Welke was de **ergste**, onveilige situatie die u het **afgelopen jaar** bij u op de afdeling heeft meegemaakt? Wat was uw rol hiérbij?
Wat is hiermee gebeurd? Wat zou ermee gedaan kunnen worden?
9. Wat zou er, met het oog op veiligheid en het voorkomen van onveilige situaties, op het werk verbeterd kunnen worden?

IV. VERANTWOORDELIJKHEID

10. We hebben het zojuist al over uw rol bij de veiligheid in het ziekenhuis gehad. Wie is/zijn, volgens u, eigenlijk verantwoordelijk voor de veiligheid in het ziekenhuis? En bij u op de afdeling?
11. Vindt u dat zij deze verantwoordelijkheid daadwerkelijk nemen?
 - Zo ja, waar blijkt dat, onder andere, uit?
 - Zo nee, waarom denkt u dat zij hun verantwoordelijkheid niet nemen?

V. EERLIJKHEID

12. Vindt u dat er bij u op de afdeling op een 'open' manier over onveilige situaties wordt gesproken?
13. Heeft u de indruk dat iedereen die bij een onveilige situatie is betrokken, hiermee ook makkelijk naar buiten komt, **ook** als niemand anders daar bij aanwezig was?
 - Zo nee, waar ligt dat aan? Zou daar, volgens u, verandering in moeten komen?

VI. COMMUNICATIE

14. In hoeverre maakt het (mee-)denken of spreken over incidenten deel uit van uw dagelijks werk? Gaat het dan om een formele of informele 'overlegstructuur'? Is dit voldoende, volgens u?
15. Waar praat u dan over, en met wie?

De veiligheidscultuur maakt weer onderdeel uit van de zogenaamde organisatiecultuur. Deze kunnen moeilijk los van elkaar worden gezien. Ik wil u graag nog een paar vragen stellen om een beeld te krijgen van de organisatiecultuur hier op de afdeling.

VII. ARBEIDSSATISFACTIE

16. Heeft u voldoende plezier in uw werk? Schenkt het u voldoening?
17. Voelt u zich betrokken bij hetgeen wat op uw afdeling plaatsvindt? Bij de patiënten? Bij uw collega's?
18. Vindt u dat u voldoende gewaardeerd wordt voor het werk dat u hier verricht?
19. Vindt u de wijze waarop in deze organisatie beloond dan wel berispt wordt met elkaar in evenwicht?
20. Beschouwt u uw huidige functie als eindstation of als tussenstation?
21. Voorziet u in de nabije toekomst nog nieuwe ontwikkelingen voor u? Promotie? Ander(e) afdeling/werk?

VII. AANVULLENDE OPMERKINGEN, WENSEN, VERWACHTINGEN

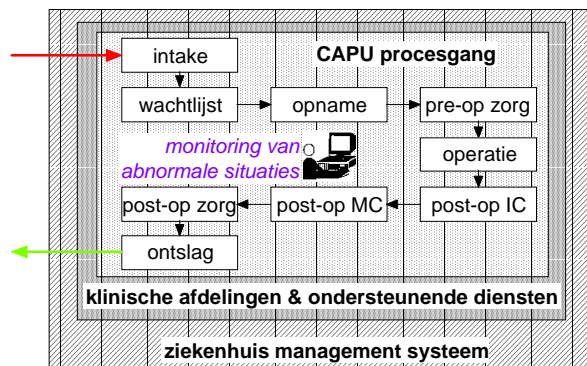
Zijn er nog vragen die u gemist heeft of heeft u nog aanvullende opmerkingen betreffende dit interview/onderzoek?

Protocol Review van Abnormale Situaties bij CAPU

This appendix contains the description of the purpose of and two alternative protocols (in Dutch) on review of abnormal situations, i.e. operational surprises, by the CAPU Review Team in PIERCE project (see section 6.7.2). At its first meeting, the Review Team decided to apply the E&CF charting protocol.

Context CAPU

In het kader van het PIERCE-project is voor de CardioPulmonale chirurgie Unit (CAPU) van het AMC een functionele procesbeschrijving opgesteld uitgaande van de fasen die iedere CAPU-patiënt doorloopt. Momenteel zijn de volgende, opeenvolgende fasen onderscheiden: intake, wachtlijst, opname, pre-operatieve zorg, operatie, post-operatieve intensieve zorg, post-operatieve medium care, post-operatieve zorg, en ontslag, zie Figuur L-1.



Figuur L-1: doorlooffasen van patiënten die worden opgenomen voor een hartchirurgische ingreep. De beslissing voor overgang naar de volgende fase berust primair op medisch-technische gronden en wordt beïnvloed door logistieke randcondities. Het CAPU-management is als aannemer van de zorg voor de CAPU-patiënt eigenaar van problemen inzake Risico-Management met betrekking tot eigen patiënten. Een RM-probleem welke optreedt in een bepaalde fase, kan veelal met relatief weinig inspanning worden ondervangen in een voorafgaande fase.

Het criterium voor onderscheid tussen fasen berust op verandering van actoren of van zorgverlenende afdeling met betrekking tot een specifieke patiënt. Bijvoorbeeld, de verpleegkundige staf van G3 Zuid en CAPU-artsen nemen de zorg over voor een patiënt van de IC-verpleegkundigen en de intensivist.

Als regel zal een patiënt 'doorstromen' naar een volgende doorloofase als aan een beperkt aantal benoembare medisch-technisch criteria wordt voldaan. Bijvoorbeeld, als de patiënt haemodynamisch en ventilatoir stabiel is zonder inotropie of beademingsmachine dan volgt overplaatsing van IC naar medium care.

Hypothese 1:

Een abnormale situatie⁴⁶ doet zich in één van de bovengenoemde doorloofasen. Uitgaande van een betrokken patiënt is dit eenvoudig in te zien. Als geen patiënt rechtstreeks betrokken is bij zulk een gebeurtenis, kan de best passende doorloofase worden afgeleid van de betrokken afdeling, dienstrooster of door een hypothetische patiënt te veronderstellen.

Als zich een abnormale situatie voordoet, kan hiervan melding worden gemaakt met behulp van SINS-1⁴⁷, tijdens de ochtendoverdracht of bij het PIERCE-team (sein 59748).

Een abnormale situatie kan optreden tijdens een activiteit van een medewerker, al dan niet betrekking hebbend op een patiënt. Zo kan op de OK de anesthesist een probleem met de beademingsmachine ondervinden, of de chirurg zich prikken met een sternumnaald.

Andere abnormale situaties betreffen rechtstreeks een patiënt. Zo kan op de OK tijdens overtillen de Swan-Ganz katheter uitgetrokken worden, en op de ICU tijdens waenen van beademing de patiënt zelf de Swan-Ganz katheter door onrust verwijderen.

Een *centrale werkhypothese* is dat een onderkende abnormale situatie in een eerder stadium, en veelal in een eerdere doorloofase, met weinig of geen (extra) inspanning kan worden voorkomen. De winst is tweeledig: tijdswinst en beheerster omgaan met procesrisico's.

Het heeft dan zin bij een geobserveerde abnormale situatie de volgende vragen te beantwoorden:

1. In welke doorloofase is de abnormale situatie geconstateerd?
2. Kan deze abnormale situatie worden voorkomen?
 - a1. zo nee, voorziet het protocol in passende scenario's voor 'Abnormale Situatie Management'?
 - a2. zo nee, zijn passende ASM-scenario's bekend en voorhanden?
 - b1. zo ja, welke acties of maatregelen in eerdere fase(n) zouden het optreden van de huidige abnormale situatie hebben voorkomen?
 - b2. welke zijn het meest optimaal, beoordeeld ten aanzien van effectiviteit en uitvoerbaarheid?
3. Wie beslist over welk ASM-scenario?

⁴⁶ Een 'Abnormale Situatie' is in dit protocol een sleutelbegrip, waarmee alle vormen van afwijkingen van normale werkprocessen kunnen worden beschreven. Dit begrip maakt een classificatie in soorten incidenten op voorhand overbodig.

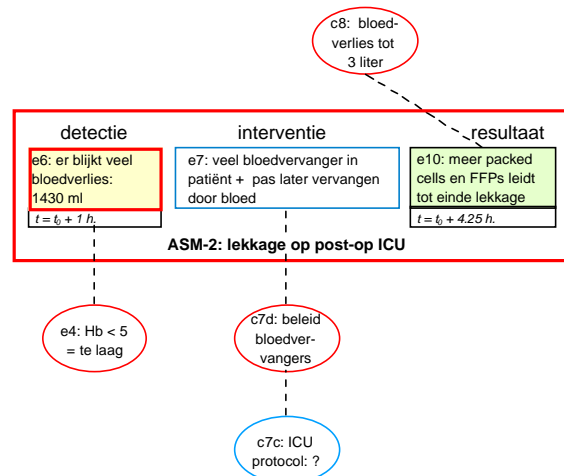
⁴⁷ SINS is een acroniem van Systematische Incident Notificatie Systeem. De eerste module, SINS-1, is per juni 1998 gereed voor dagelijks gebruik. Doel van SINS-1 is laagdrempelige melding van abnormale situaties. Toepassing van 'intelligente' formulieren moet voorkomen dat voor verwerking onnodige gegevens worden opgevraagd en data direct in elektronisch formaat beschikbaar zijn.

Methode 1: Abnormale Situatie Management: ASM-model

Het ASM⁴⁸-model beschrijft het optreden en de aanpak van een abnormale situatie⁴⁹ in termen van ASM-units. Een ASM-unit is omvat 3 functionele 'gebeurtenissen': **detectie**, **interventie** en **resultaat**. De gebeurtenis 'detectie' beschrijft de onderkenning van een abnormale situatie welke vroeger of later interventie vereist opdat te beheersen risico's onder controle blijven. De gebeurtenis 'interventie' beschrijft de maatregel (serie maatregelen) welke is (zijn) getroffen ter opheffing van een abnormale situatie. De gebeurtenis 'resultaat' beschrijft de toestandsveranderingen in het systeem waarop de abnormale situatie betrekking heeft, volgend op de interventie.

AS-voorwaardenscheppende omstandigheden en condities kunnen worden gekoppeld aan de relevante functionele gebeurtenis in een ASM-unit, zie Figuur L-2.

Bij een opgemerkte abnormale situatie hoort een interventie welke zich laat opsporen! Er bestaat een normaliter beoogd verband tussen interventie en resultaat.



Figuur L-2: Abnormale Situatie (AS) Unit. De AS-unit omvat gebeurtenissen 'detectie' 'interventie' en 'resultaat'. De ovals buiten de AS-unit beschrijven omstandigheden die van directe [e4, c7d], dan wel van systematische [c7c] invloed zijn op het AS-verloop.

Een abnormale situatie kan een gevolg zijn van een eerder opgetreden, mogelijk nog niet onderkende abnormale situatie.

⁴⁸ ASM staat voor Abnormale Situatie Management.

⁴⁹ Een optredende abnormale situatie (AS) is een zich ontwikkelende afwijking in procestoestanden van een systeem, welke vroeger of later interventie vereist teneinde niet-beoogde schadevorming te voorkomen dan wel te beperken.

Hypothese 2:

Toepassing van het ASM-model maakt de representatie van een abnormale situatie-gebeuren transparant voor de doelgroep (Review Team) zonder in essentie belangrijke RM-opties⁵⁰ te verliezen in de vereenvoudiging.

Validatie:

in Review Team vergelijkbare (of dezelfde) cases laten beoordelen m.b.v. beide beschreven methoden voor gebeurtenisrepresentatie: ASM-model en E&CF-model

Sleutelvragen ASM-model

Beantwoord de volgende vragen voor de beschouwde Abnormale Situatie.

Inzake de AS-reconstructie:

1. Wat is de eerst opgemerkte *abnormale situatie*?
In welke fase van zorgverlening is de abnormale situatie opgetreden?
Kon de abnormale situatie eerder worden gedetecteerd?
 2. Wat is de diagnose t.a.v. het onderliggende probleem?
Is urgente interventie geboden?!
Is interventie opportuun in het licht van het denkbare risico?
 3. Welke *interventie*(-reeks) is als beleid gekozen?
Is het interventiebeleid uitgevoerd?
 4. Welk *resultaat* had de interventie op het onderhavige proces?
Is dit resultaat conform plan of verwachting?
 5. Welke omstandigheid of -heden is of zijn kritisch voor de 3 respectieve functionele gebeurtenissen: **detectie**, **interventie** en **resultaat**?⁵¹
- Inzake de mogelijkheden voor verbetering van RM:*
6. Welke RM-opties zijn alleen of in een specifieke combinatie afdoende effectief om de opgetreden Abnormale Situatie op te heffen.
Welke hiervan zijn uitvoerbaar?
 7. Welke RM-maatregelen worden aanbevolen?
Wie moet(en) beslissen over keuze tussen, respectievelijk implementatie van de aanbevolen maatregelen?

⁵⁰ 'RM' staat voor RisicoManagement in de betekenis van omgaan met rest-resico's. RM-opties zijn mogelijkheden voor maatregelen voor verbetering van de beheersing van de rest-risico's. Rest-risico's zijn de risico's die nog inherent in het (werk-)proces achterblijven nadat maatregelen voor beheersing van bekende risico's zijn getroffen.

⁵¹ Het antwoord op deze vraag geeft directe ingangen naar opties voor maatregelen voor verbetering van het omgaan met (rest-) risico's in de werkprocessen (RM-opties). Een kritische omstandigheid kenmerkt zich doordat bij afwezigheid de gelieerde gebeurtenis niet optreedt, omdat dan optreden en verloop van een Abnormale Situatie anders zouden zijn gegaan.

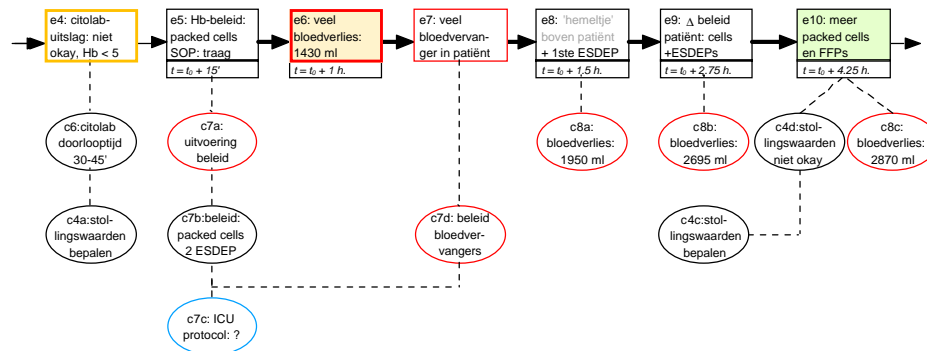
Methode 2: Gebeurtenissen en Conditie: E&CF-model

Het E&CF-model⁵² [SSDC-14] beschrijft het optreden en verloop (in de tijd) van een abnormale situatie in termen van reeks van gebeurtenissen en voorwaardenscheppende condities, zie Figuur L-3.

Een 'gebeurtenis' is een optredende actie welke leidt tot verandering van toestanden in een proces binnen het operationele systeem.

Een 'conditie' is een relevante omstandigheid waaronder de gebeurtenis optreedt. Een 'kritische conditie' is een noodzakelijke voorwaarde voor het optreden van een specifieke gebeurtenis.

Een combinatie van een gebeurtenissenreeks en kritische condities vormt een (incident-) scenario⁵³.



Figuur L-3: E&CF-representatie van het optreden en managen van een abnormale situatie. Elke rechthoek is een gebeurtenis bestaande uit actor, actie, object van actie en tijdstip. De abnormale situatie wordt opgemerkt bij gebeurtenis e6. Ter vereenvoudiging zijn in dit voorbeeld actoren weggelaten uit de omschrijving van gebeurtenissen: de focus is gericht op identificatie en beoordeling van systematische RM-opties. De ovals representeren condities waaronder de gelieerde gebeurtenissen zijn opgetreden. Naarmate condities verder van de gebeurtenissenketen afliggen, zijn ze meer systematisch van aard: in dit geval veelal afgeleid uit het ICU-protocol. Kritische condities in het voorbeeld zijn c8 en c7a, b, c en d.

De E&CF-reconstructie van een abnormale situatie en de aanpak hiervan helpt relevante data te verankeren in de specifieke context van het operationele (werk-) proces waarbinnen de abnormale situatie blijkt. De sequentie van beschreven gebeurtenissen moet het verloop van de abnormale toestanden afdoende beschrijven. Bij het ontbreken van gegevens over relevante gebeurtenissen of condities worden expliciet aannames gedaan ter toetsing. Aldus kunnen ook scenario's worden opgespoord.

⁵² E&CF staat voor 'Events & Causal Factor' charting. Deze procedurele methode is bedoeld voor ordening van gegevens inzake een te onderzoeken voorval waarbij een primaire reeks van acties en reacties (gebeurtenissen) de toedracht beschrijft en andere gegevens als directe dan wel structurele randvoorwaarden kunnen worden benoemd en gepositioneerd ten opzichte van de gebeurtenissen.

⁵³ In dit verband is een 'scenario' een generieke beschrijving van een proces van onbedoelde schadevorming, veelal aangeduid als 'ongeval'. Een foutenboomanalyse van specifieke schadevorming biedt een mogelijkheid tot verificatie van een met behulp van het E&CF-model geïdentificeerd scenario.

Hypothese 3:

Toepassing van het E&CF-model maakt de representatie van een abnormale situatie-gebeuren op basis van procesgegevens transparant voor de doelgroep (Review Team), zodanig dat effectieve RM-opties⁵⁰ kunnen worden geïdentificeerd als 'kritische condities', beschrijvende feiten kunnen worden geverifieerd en unieke scenario's kunnen worden onderkend.

Validatie:

in Review Team vergelijkbare (of dezelfde) cases laten beoordelen m.b.v. beide beschreven methoden voor gebeurtenisrepresentatie: ASM-model en E&CF-model

Sleutelvragen E&CF-model

Beantwoord de volgende vragen voor de beschouwde Abnormale Situatie.

Inzake de E&CF-reconstructie:

1. Wat is de te onderzoeken *abnormale situatie*?

In welke fase van zorgverlening is de abnormale situatie opgetreden?

2. Welke *gebeurtenissenreeks* beschrijft het optreden en managen van de abnormale situatie?

Wat zijn de onderscheiden gebeurtenissen en condities?

- a) Beschrijf een *gebeurtenis* (actor, actie, object van actie, tijdstip) of *conditie*: ga uit van feiten.

Is de gebeurtenis onderdeel van de AS-gebeurtenissenreeks?

Is de gebeurtenis als optredende conditie relevant?

- b) Beschrijf een *relevante conditie*: ga uit van feiten

Is deze conditie 'kritisch'? Zo ja, voor welke gebeurtenis?

Wat is de tijdsvolgorde van de gebeurtenissen en condities?

Vanaf welk vroeger moment kan een kritische conditie worden opgemerkt en beïnvloed?

3. Wat is de *diagnose* t.a.v. het onderliggende probleem?

Welke gebeurtenis markeert het moment van detectie van de abnormale situatie?

Is urgente interventie geboden?!

Is interventie opportuun in het licht van het denkbare risico?

Welke gebeurtenissen markeren interventieacties.

Is effect van de interventie conform plan of verwachting?

Voor welke gebeurtenissen of condities ontbreken verificatiegegevens?

Welke gegevens en waarom?

Inzake de mogelijkheden voor verbetering van RM:

4. Welke RM-opties zijn alleen of in een specifieke combinatie afdoende effectief om de opgetreden Abnormale Situatie

Welke kritische condities bieden effectieve RM-opties?

Welke RM-opties zijn uitvoerbaar?

Welke RM-opties kunnen in eigen beheer door CAPU worden ingezet?

Welke van deze kritische condities kunnen in een vroegere fase worden onderkend en beïnvloed?

5. Welke RM-maatregelen worden aanbevolen?

Wie moet(en) beslissen over keuze tussen, respectievelijk implementatie van de aanbevolen maatregelen?

Appendix M

Case of the Week Review

Cases demonstrate a great variety in severity and 'visibility'. From the point of view of risks, the most severe cases are not necessarily the most obvious ones! Much is to gain from at least more systemic review of daily operational anomalies.

Let me explain the functioning of SINS by an example taken from a cardiac surgery department a major hospital.

Casus/Example in short:

A skin needle dropped into the pleural cavity of the patient just before closure of the surgery wound. Although visible on a series of X-ray photos, this abnormal situation was not seen by several physicians in three subsequent operational phases.

During routine sampling quality check, the medical manager identified the needle 4 days later. The 'corpus alienum' was removed by re-operation without further complications.

Extended case description:

During the fixation of drain tubes to the chest of a patient who had undergone a bypass operation, one fishhook alike skin needle fell into the pleural cavity. Such a skin needle is used as a 'third hand' to keep the skin rim of the wound together in order to allow minimal scarves resulting from the closure sewing. Four skin needles are available for this purpose and arranged on a magnetic mat.

The needle remains in the pleural cavity unnoticed. Standard operating procedures includes daily X-ray photos after the operation, the first one to be made shortly after the surgical procedure. The skin needle was clearly visible on each of these photos. At the 4th day after the operation, a quality assurance inspection was applied to the patient's file, and the needle was identified as a 'corpus alienum' that had to be removed. At that time, at least 4 different physicians had seen the photos, each of them without spotting the needle. The patient had to be told the bad news, was offered a complaint submission form by the doctor, and subsequently was brought back into the OR for a brief, but nevertheless surgical procedure.

The patient recovered well, wrote a letter to the board of the hospital, praising the doctors of the surgical department, accepting that also doctors are fallible human beings. He was complaining, though, that the board has not sent its excuses and some flowers for the serious suffering... as if the board has supervisory responsibility in daily medical care and treatment.

WHAT HAPPENS WITH SUCH A CASE?

Typically, the discovery of the needle on the X-ray picture in one of the relevant operational stages will be reported through SINS-1.

This time, the needle was spotted by 2nd line the manager who then intervened in the daily operations: re-operation to be scheduled for the next day. The abnormal situation was reported to the project team for further evaluation as "Case of the Week".

The Review Team first assesses its concern with the presented situation:

- does it deserve attention now, and if not, what incidence level will be set as trigger for problem review
- is here a problem with unclear causal factors - rephrased as 'why?' could it happen - or is there a need to find potential solutions and select the best course of action.

In the case problem analysis, the evolution of the abnormal situation is reconstructed for evaluation using the Events & Causal Factor charting analysis method. The clues for risk control lessons can be found through the conditions that enable the occurrence on the abnormal situation.

About the concern, the trigger level was implicitly set to one, requiring further review. The argument is that, although such an object may remain in the body without doing any harm in future, the medical centre may be confronted with costly liability claims when during a periodical check the object is spotted after all. Such objects should not stay in the body of the patient.

E&CF-reconstruction depicted in Figure M-1 leads to the following **findings**. The **key conditions** where:

- the needle could have been dropped during the removal of surgical gauze: anyway, it happened
- the fact that the needle could leave the OR with the patient **without being noticed**
- **visibility** of the skin needle on the X-ray photos
- the fact that several physicians **did not spot** the needle visible on the X-ray photos
- a third order QA check led to spotting the needle before the dismissal of the patient

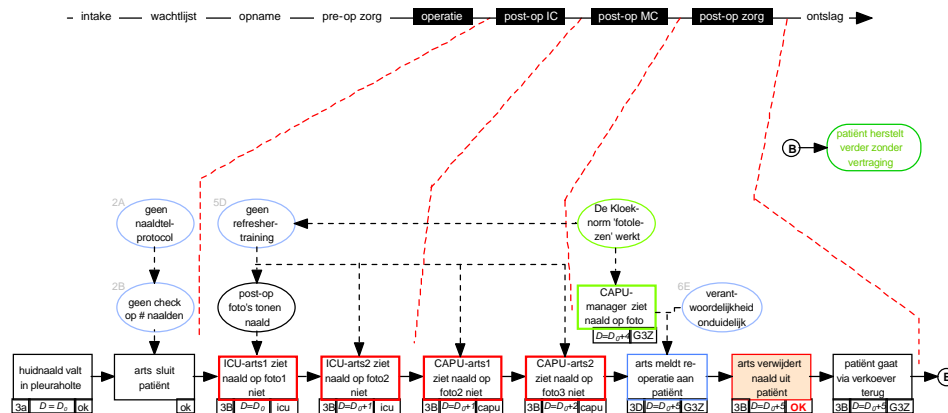


Figure M-1: E&CF charting of the fallen skin needle left behind in the patient... The colour coding is derived from MORT: red = less than adequate, green = adequate, blue = more information needed. Critical conditions have a <figure.letter> coded label at the upper left that corresponds with the causal factor classification scheme of Appendix N.

The most appropriate risk control opportunities are:

- in the ICU, by seeing of the first post-surgery X-ray photo that becomes available before the patient is wakened up
- in the OR, before closure of the patient by protocol check, i.e. counting the number of skin needles

The latter step is *preventive* in nature, where the former is *corrective*. The other subsequent X-ray pictures allow possibly also identification of the object, like in this case, but impose a second surgical procedure experience on the patient, which is patient unfriendly and thus to be avoided.

The systemic seeing of X-ray photos can be trained. This would not provide a solid solution for the following reasons:

- an object may not be visible at all due to its 3-dimensional position in the body or material characteristics resulting in lack of contrast;
- people are biased by expectations, etc. and thus don't see the obvious while looking for the standard problems;
- the implementation of individual vigilance and sensitivity regarding the seeing of photos cannot be assured at the required high level.

Fortunately, the counting of needles is feasible, can easily be added to the existing protocol for counting gauzes, clamps, etc. with the advantage that this is performed by other people, i.e. the operating room nursing staff. The ownership of this protocol lies with the OR-department, so this preventive measure must be communicated at department level if communication at between the operational units is not effective at this point.

Thus has been concluded and implemented. Also, a suggestion to reinstall the form tradition of monthly structured photo seeing sessions for residents and co-assistants has been adopted as well. The overall result is a (at least) lines of defences: the first line is in the OR, counting of the needles, AND with the second, structured X-ray scanning, and the third one remains: sample quality checks while the patient is still in house.

POINTS TO MAKE:

- incident in OR-phase
- followed by 2 obvious RM-control opportunities [count needles; first X-ray picture]
- needle spotted in 3rd line of defence
- mitigation interfering with regular rehabilitation process (ward) by corrective action in OR (in 3rd process phase after the dropping of the needle)
- RM-actions aimed at the structural management controls (protocol, training, communication policy)

Implementation followed fast - within days, because of accepted authority of review team. Summarising:

- as such, E&CF chart illustrates that depicted facts from incident reconstruction have lost almost all context variety. Understanding of what happened and why required variety compensation which was provided by the Review Team. The E&CF chart helps to focus on the more structural conditions.
- systemic case review using E&CF analysis results in evidence-based operational Risk Management !

(this page is left blank intentionally)

Appendix N

Cause Codes

This table from [DOE 92] shows causal factors grouped in categories of systemic management factors.

1. Equipment/Material Problem
 - 1A = Defective or failed part
 - 1B = Defective or failed material
 - 1C = Defective weld, braze, or soldered joint
 - 1D = Error by manufacturer in shipping or marking
 - 1E = Electrical or instrument noise
 - 1F = Contamination

2. Procedure Problem
 - 2A = Defective or inadequate procedure
 - 2B = Lack of procedure

3. Personnel Error
 - 3A = Inadequate work environment
 - 3B = Inattention to detail
 - 3C = Violation of requirement or procedure
 - 3D = Verbal communication problem
 - 3E = Other human error

4. Design Problem
 - 4A = Inadequate man-machine interface
 - 4B = Inadequate or defective design
 - 4C = Error in equipment or material selection
 - 4D = Drawing, specification, or data errors

5. Training Deficiency
 - 5A = No training provided
 - 5B = Insufficient practice or hands-on experience
 - 5C = Inadequate content
 - 5D = Insufficient refresher training
 - 5E = Inadequate presentation or materials

6. Management Problem**6A = Inadequate administrative control****6B = Work organisation/planning deficiency****6C = Inadequate supervision****6D = Improper resource allocation****6E = Policy not adequately defined, disseminated, or enforced****6F = Other management problem****7. External Phenomenon****7A = Weather or ambient condition****7B = Power failure or transient****7C = External fire or explosion****7D = Theft, tampering, sabotage, or vandalism**

Generic Cases out of Operational Surprises

This appendix is a preliminary reflection on the development of validated generic cases for the case base following case review by the Review Team.

Sources of Cases

The Review Team receives cases for validation and classification from several sources:

- from the SINS-1 notification system when a case cannot be matched with one from the SINS-2 Case Base
- from the observations by CAPU staff as reported during the morning 'overdracht' session or directly reported to the project team

The latter source generates the so-called 'Case-of-the-Week' examples.

Any case from these sources has passed the detection threshold in the operational system CAPU. But not every reported case represents a problem of which the causes must be identified before an adequate solution can be found.

Having Kepner-Tregoe in mind, any new case needs to be assessed first for appropriate action. The available options are problem analysis (PA), decision analysis (DA) and potential problem analysis (PPA).

- when a "why"-question is to be answered and relates to the past, problem analysis (PA) is the pertinent method: WHY could this happen...
- when the cause of a problem is clear or not relevant, but a choice between option for solutions need to be made, decision analysis (DA) is appropriate: WHAT course of action shall we SELECT?
- when a measure (to solve a problem) is going to be implemented, the "what can go wrong badly"-question can be answered timely using potential problem analysis (PPA):?

Generic Cases

Example 1:

New air-cushioned beds in the corridors of the ICU hinder the passage of fully equipped (patient, drains, reservoirs, ventilator tubes, etc.) beds coming from the operating theatre.

Problem-cause: known

In the past, beds came in and out with the patient; now, new beds remain at ICU and are parked on the corridor for cleaning and storage.

Related risk:

Bed with patient & devices collides with a parked bed.

Type of issue:

DA: What can be done best to reduce operational risks and improve operational conditions for personnel.

Options:

Find cleaning and storage space outside ICU; park ICU-beds in other sections of ICU-corridors; do nothing about current situation.

Example 2:

Wrong arterial graft brought into OK while patient was already opened for surgery. Correct replacement had to come from the (national) 'arterial donor bank'. Surgery had to be stalled for an hour.

Problem-cause: unknown

Anything could have gone wrong from ordering to performance errors underway and QA checks. A main issue is whether the origin lies within the AMC or elsewhere (more difficult to prevent!).

Related risks:

- duration of surgery conditions is threat to stability of patient's condition (maybe fatal!)
- delay may lead to rescheduling to a next day of another operation

Type of issue:

PA: Why could the obviously wrong arterial graft be delivered in the OK?

Example 3:

Patient arrives at ICU in state of hypothermia. Also, a lot of fluids are drained out of the body.

Type of issue:

PA: Why could the state of hypothermia remain so long?

Appendix P

Cybernetic Learning - VSM View

A theoretical viewpoint on these issues, and one that is particularly compatible with an information theoretical paradigm, is provided by cybernetics [Koornneef 98]. Within cybernetics, Beer has developed the "Viable System Model" for application to socio-technical systems (Beer, 1985). The VSM provides useful insight into the role of data processing and communication in the control of organisational systems.

The Viable System Model

The VSM considers an organisation as a system comprising five principal sub-systems. These sub-systems serve dedicated functions that together ensure the viability of the whole system. Figure P-1 is a simplified illustration of the VSM showing the relations of the sub-systems. The circles, denoting operations, each have a local management (shown as squares labelled "MU").

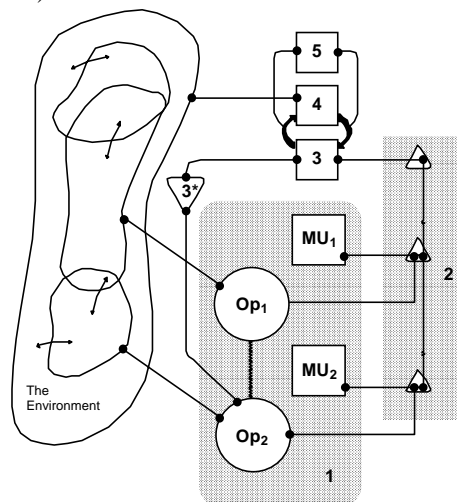


Figure P-1: *simplified illustration of the Viable System Model (after Beer, 1985)*

Figure P-1 shows just two operational elements and there might easily be more than this in practice. The circles and their associated management units are collectively referred to as "System One". The operational elements of System One will be linked through the

environments that they share, or through inter-process linkages, and – in the weakest case – competition for corporate resources. These are shown as the links from Op_1 and Op_2 to overlapping parts of the sketchily drawn environment. To maintain harmonious interaction between the operations, a means is required to monitor routine aspects of each operation and damp-down instability in System One. This routine monitoring and damping-down of instability is the function of "System Two".

Note that the perception of instability in System One is something that the individual "management units" cannot do on their own: an overseeing role is indicated. System Two is therefore provided as a service function to System One by its senior (or corporate) management. System Two is established in discussion with the local managers in System One, with particular regard to the nature and means of monitoring and the criteria for stability. However, the discussion and the establishment of senior management control, require adequate models of System One and its interaction with the environment. These functions of senior management are the purpose of the third part of the VSM: "System Three".

In order to monitor competently and achieve routine regulation via System Two, System Three must have a means to establish and maintain competent models of System One. The nature and medium of these models will vary greatly in practice. However, certain axioms are common to all models, not least among them is that models must be a simplification of the reality they represent. This simplification can be referred to as a reduction in variety. In other words, the model has a lower variety than the reality it represents. The danger of the practical imperative of simplification is that it may give rise to systematic blindness to poorly modelled aspects of System One functioning (especially over time). In the majority of cases, this "monitoring blindness" will lead to corresponding weaknesses in the services provided to, or the control asserted on, System One.

In order to establish and maintain competent models of System One, System Three must have the means of detecting (via System Two) extraordinary events that challenge the competence of its models of system one. When so triggered, System Three must have the means of interrogating System One to update or establish competent models. In the terminology of cybernetics, this process restores "requisite variety" to System Three. In the terms of this thesis, this process is a major function of an "in-depth" investigation (such as might follow from a major accident). In the language of the VSM this is referred to as "System Three-Star". The function of System Four is to assure the external and future adaptiveness of the lower sub-systems (the assembly comprising System One, System Two, System Three, and System Three-star). System Five is tasked with the generation of policy and high-level standards of performance for the whole system.

It is beyond the scope of this thesis to fully introduce the VSM and the intricacies of sub-system interconnection that provide rigour to the model. See [Beer 85] for these details. However, this sketch provides the essentials for understanding the model which can be, and has been applied to systems at all levels from a country (Chile under Allende, see [Beer 81]) to a single organisation (see examples described in [Espejo 89])

How does VSM help with organisation learning from incidents

What light does the VSM shed on issues of organisational learning and accident data? An adaptive system requires that routine monitoring (including operational ad-hoc reports of abnormal situations) provides a trigger to in-depth investigation. Routine monitoring has two roles: first, to ensure that criteria for steady state stability are maintained in real time, and; second, to ensure—if instability is detected—measures are implemented to damp down this instability. Damping down is achieved in two ways: the application of prepared plans or the acquisition of data to inform the development of new plans. In the language of the VSM, System Two provides the impetus to System Three to replenish its variety (i.e. the variety of the model of System One as maintained on System Three) via system three-star. Whilst automation can certainly help the monitoring side of the equation, the role of human expertise to re-introduce the variety necessarily lost in that monitoring process is essential to adaptation.

The contrast, between the acquisition of operational data by System Three-Star and by System Two, is worthy of note. System Two data are informative within a Single-loop situation. An important function of System Two data is to identify the current position of the operational system within the set of possible states defined for that system. On the other hand, System Three-star data are informative within a double-loop learning situation. Effectively, operational data delivered to senior management by System Three-star are the means of challenging and changing the operational model (of System One) in use. It may be that senior management (System Three), on the strength of its new understanding, may be able to specify new norms to System One either via System Two or directly via the Command channel. In either case, System One must be able to amplify these new requirements into operational reality. However, another outcome of the System Three-Star inquiry of System One may be the creation of a new System One or the termination of a System One currently in operation. The information to achieve this comes from the *environment* of the system. The direct links between System One and its environment (especially those between the operations of System One and the environment) are not sufficient here. System One's environment has less variety than System Four's. Systems Three needs input from the environment through System Four to evaluate the performance of a System One in order to come to an adequate decision.

For example, in a production unit a series of critical incidents occur in which a particular operator (System One) is involved. The unit's manager (System Three) considers throwing out this operator. Before doing so, a System Three-Star audit is invoked by System Three because direct communications between System Three and System One did not result in lasting effects. The System Three-Star audit by the personnel department (a System Four) – being aware of the large variety of drivers for individual behaviour – revealed that the operator had to work in almost complete solitude. The operator no longer could stand this and initiated the incidents to attract attention. The problem could be solved by regular coffee and lunch breaks whereby the operator joined his mates.

There is a clear need for organised research of the environment at the level of the whole system. This research is very much characteristic of double- and higher-loop learning, and is typified by the acquisition of data from outside of the system but using the purposes of the system to define what is relevant (to filter noise and transform data into information).

(this page is left blank intentionally)

Appendix Q

Unfolding the SiF in VSM

This appendix discusses the issue of levels of recursions as concept in the VSM using the case of CAPU the PIERCE project in Medical Centre C.

The System in Focus (SiF) is a cardiopulmonary unit (CAPU) in an academic medical centre (AMC). The hospital is a lower recursion of the national health care system. The CAPU is a specialised section within the surgery division B, that is one of the eight divisions that together produce the clinical functions of this hospital. The higher triple-recursions of the CAPU are depicted in Figure Q-2. The core activity of the CAPU is cardiac surgery, but members of the medical staff also initiate clinical research activities that generate new knowledge and skills as well as a second source of income.




Figure Q-1: *simplified illustration of the Viable System Model (after Beer, 1985)*

Cardiac surgery cannot be done without pre- and post surgery wards, post-op intensive care and medium care units. Typically, a patient is 'handed over' from one ward to the other, and somewhere underway temporarily taken out for surgery. In terms of the VSM, the route of the patient is through the (squiggly) lines between operational circles, see Figure Q-1. The operational systems that together produce the CAPU are the surgical

staff, the regular CAPU-ward (pre-op and post-op), post-op ICU, and post-op MCU. A Coronary Care Unit (CCU) and a clinical lab provide standby support. The surgical team cannot perform surgery without an anaesthetic team.

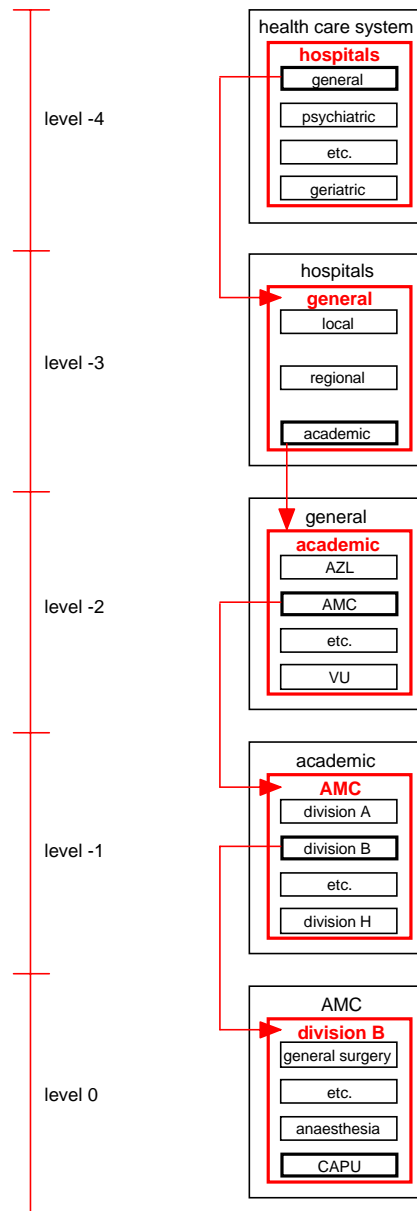


Figure Q-2: higher recursion levels of the cardiopulmonary unit (CAPU) of the AMC as embedded in the Dutch health care system. The System in Focus is at reference level 1, see Figure Q-3. The figure shows a triple-recursion at any level, so the next recursion down or up is kept in sight as well.

The CAPU ward (G3 Zuid) provides care in three different stages of the cardiac surgery process. Firstly, it receives incoming, scheduled patients and prepares them for the operation. Secondly, it runs the post-op Medium Care Unit that is used at least for one day after a patient is dismissed from the post-op ICU. And finally, it generates rehabilitation care so that the patient can be sent home as soon as possible.

Levels of recursion

A leading criterion for the identification of system producing operations is whether or not the operation is necessary for the System in Focus and has a recognisable management unit. In the CAPU organisation, the clinical processes are structured along the throughput cycle, which applies to any cardiac surgery patient. The main throughput phases are, sequentially ordered: "intake", "pre-op care", "surgery", "post-op care". The intake⁵⁴ ends with scheduling the patient, i.e. placing the patient on the waiting list until a surgery slot is available and the patient is called for admission.

Admission is the first step in pre-op care: the patient is submitted to final tests and prepared for surgery. The health condition of a patient may be very poor at this stage. The CAPU ward delivers pre-op care. The CAPU staff, including assistance from residents and when needed a perfusionist, i.e. the operator of the heart-lung machine⁵⁵, performs the surgical intervention, facilitated by an anaesthetic team and operating room staff. After surgery the patient is transferred for post-op care to the ICU and from there, possibly within 24 hours, to the MCU for another night before moving into the regular ward. In the regular ward, the patient is rehabilitated until he or she is ready for dismissal out of the cardiac centre back to the sending hospital or home.

In each stage, a designated cardiac surgeon is in charge of the medical treatment of the CAPU patients in that phase. In the wards, the care process results from the combination of the medical treatment process and nursing operations. The nursing processes are also managed separately for each stage.

Recursion 0 / Level 1: CAPU

Thus, we can say that the following mainstream operations produce the cardiopulmonary surgery system (CAPU):

- intake,
- pre-op care,
- surgery, and
- post-op care.

The CAPU system cannot do without their patients as they are the primary targets for delivery of the CAPU product, i.e. successful cardiopulmonary surgery treatment: we will find the patients at lower recursion levels.

⁵⁴ Typically, a referring cardiologist will propose the patient for cardiac surgery. A CAPU surgeon will accept the patient if the surgical intervention is feasible.

⁵⁵ A heart-lung machine is connected to the vascular system of the patient and controls the flow, the temperature and the oxygenation of the blood. Thus, many functions of the anaesthetists are transfer to the perfusionist while the perfusion pump is functioning. The changes in blood circulation states, i.e. heart-lung machine connected/disconnected, may represent serious risks to the patient.

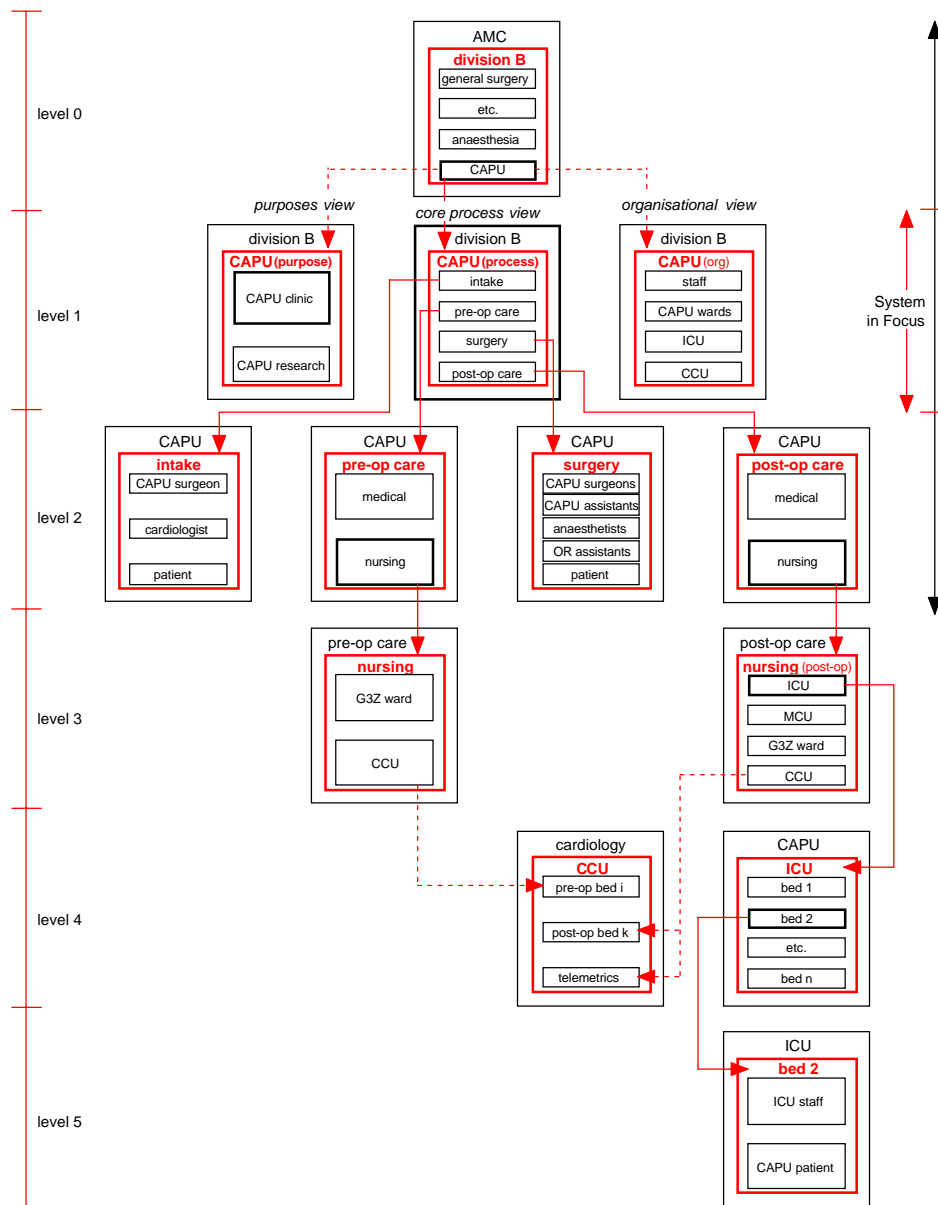


Figure Q-3: continuation of the unfolding of recursion levels for the cardiopulmonary surgery system in the AMC; see Figure Q-2 for the higher recursions. The CAPU producing operations are shown in the lower recursions (four levels deep). Recursion level 1 depicts the System in Focus, the CAPU. Abnormal situations occur at the lowest operational levels shown. In this study, a practical caesura is placed at level 2 for the "one patient at the time" batch processes, and level 3 for the continuous "multiple patients at the time" operations.

The intake stage is a particular one, because the patient is still outside the CAPU, in another hospital, another hospital department or at home. The key process activities in this phase are assessment for surgery and scheduling/prioritising patients on the waiting list⁵⁶. Only during the intake phase is a patient still not yet under direct supervision of the CAPU.

CAPU = Σ organisational entities

We can also see that the formal organisational entities, i.e. CAPU staff, CAPU ward, ICU, CCU and OR department, do not depict the core processes that produce the CAPU as a viable system. The CAPU ward is essential in two stages, pre-op and post-care, and produces three care processes: pre-op care, post-op medium care and post-op regular care. The CCU may be needed in pre-op care as well as in post-op care, but only as service 'on demand' as an extension of CAPU ward operations, whilst the medical responsibility for a patient remains with the CAPU physician in charge.

CAPU = clinic & research

Another option for unfolding the CAPU is thinkable on the basis of actual operations of a different nature, which exist in their own right. Alongside the core process of cardiopulmonary treatment and care of patients in the CAPU clinic within the AMC, individual members of medical staff have income and output producing research processes for which they are accountable. The CAPU as an organisational unit communicates its research progress to higher recursions and to the outside world, so the research is considered as an essential activity.

From these facts we may state that the CAPU is produced by two mainstream operations: the cardiopulmonary clinic and research relevant to cardiopulmonary surgery practice. This approach is not elaborated any further here, but it might be fruitful to diagnose the CAPU from this angle: it may well turn out that the research function is not well embedded yet as a viable system in terms of the VSM.

Now, the CAPU clinic will be unfolded as described above, but at the next (lower) recursion level.

Recursion 1: the CAPU producing processes (as viable systems)

At the first recursion down we see four systems that are essential for the CAPU: without any of these, the CAPU system cannot exist. The systems for intake, pre-op care, surgery, and post-op care respectively encompass their own specific operations and management, and have their own specific connections to the outside world or environment. For instance, the environment of the "intake" system includes potential CAPU patients, their family and family doctors, their cardiologists and regional hospitals, but also the emergency care/first aid unit where patients come in as victims of crime, accidents or acute major heart failure. Obviously, the "post-op care" system shares a significant part the "intake"-environment as the patient will be dismissed for final recovery care and treatment elsewhere, often at the sending institution. The environment of the "surgery"

⁵⁶ The actual maintenance of the waiting list and the call of a patient to admission is a self-organisation function of the CAPU management services: in System Three of the SiF, see Figure Q-7.

system includes also other thorax centres and professional organisations. Similar connections exist between the care subsystems and their particular environments.

If we take a closer look into the level 2 systems, we see the systems for "intake", "surgery", as well as "pre-op" and "post-op" care, see Figure Q-3.

Recursion 2

A CAPU surgeon, the proposing cardiologist, and the proposed patient, produce the "intake" system. If any of these subsystems is lacking then there is no intake⁵⁷.

In the same way, surgeons, assistants, anaesthetists and last but not least the patient produce the "surgery" system. In both systems, only one patient at the time is in focus as part of a batch process.

In the pre-op and post-op care delivery systems, individual patients are grouped according to their throughput stage, and clinical staff is divided over these groups. Care is delivered continuously, i.e. 24 hours/day. The nursing operations are separate from medical care and treatment. Thus we see that the care systems are produced by nursing as well as medical operations. Consequently, we have another recursion to go in order to find the level of respective wards with the patients at recursion level 3.

Recursions 3, 4 and 5

The need to go down this far lies in the fact that patient-related abnormal conditions become apparent at the operational level of the nursing systems, i.e. in the respective wards: regular care, Intensive Care, Medium Care, possible assisted by Coronary Care. Notifications of abnormal situations normally originate from operations in which the patients play active roles. The Intensive Care Unit system contains, see at level 4, a limited number of bed systems on a one-patient-to-one-nurse basis. So, the ICU system consists of as many operations as functioning beds. One recursion lower, at level 5, we finally see that an ICU bed system is actually produced by the nurse and his/her patient.

As said above, the Coronary Care Unit (CCU) is not a part of the CAPU, but a facilities provider 'on demand'. The CCU is produced by pre- and post-op intensive care, with exclusion of mechanical ventilation that can only be provided in the ICU, and by telemetric surveillance systems.

Discussion

We have found so far that four operations produce our System in Focus, two of which are characterised by one-patient-at-a-time batch processes and the other ones as multiple patients systems with full-continuous operations. The latter require additional recursions before we have developed adequate insight into the CAPU system.

The structure of any triple-recursion block is the same. The recursion in focus is embedded in its higher recursion and contains its lower recursions, see Figure Q-4. There is no hierarchic significance in the vertical listing of elements at one recursion level, but their coupling may be anywhere between loose and strong.

The connectivity between any pair of recursions is always the same, just as is the VSM.

⁵⁷ One can easily see that eventually the lack of intakes will result in the end of the CAPU.




Figure Q-4: depiction of a triple-recursion of a system: the next higher recursion upwards and the system producing operations at the next lower recursion are clearly shown in one block.

Although the CAPU organisation is properly depicted as belonging to *this* set of recursions shown in Figure Q-2 and Figure Q-3, it belongs to an arbitrarily large number of other sets of recursions too. For instance, cardiac surgery also breaks up by medical specialists, by thorax centre or geographically. In Figure Q-3, three dimensions for recursions of the CAPU are depicted at recursion level 1: a dimension of purposes, of core processes, and of organisational entities. Beer puts it this way:

"Whatever viable system we want to model, exists in a variety of recursive dimensions. 'What business are we in?' is the classic question for a board of directors to consider - and there may be several answers. So, the System-in-Focus may have more than one next higher and next lower recursion." [Beer 1985, p. 6]

A VSM of the CAPU: some VSM basics

Now we have identified the CAPU as System in Focus and its recursions up and down, we develop a model like the one depicted in Figure Q-1 for the CAPU as a viable system. We start with the four embedded viable systems "intake", "pre-op care", "surgery" and "post-op care" that together produce the CAPU, see Figure Q-3. The four sets of the operational circles and managerial squares, such as in Figure Q-5 below, form together *System One* of our SiF.




Figure Q-5: the black round en square objects represent a system-producing operation and its management unit of the SiF. It contains, depicted by the thin lines, a viable system at the next level of recursion.

The circle encloses the relevant operations that actually produce the SiF. The square embeds all the managerial activity needed to 'run' the operations. In Figure Q-6, the lower recursions from Figure Q-5 have been cleared, because we concentrate on the SiF. The amoebic shape has been added and represents the environment of this operation and its management unit. The arrows refer to the necessary interactions between these three basic

entities: each stand for a multiplicity of channels whereby the entities affect each other. Materials, men, machinery, money and data are manipulated and exchanged through these channels. Moreover, what is going between these entities is the management of *complexity* for which *variety* is a measure, because it counts the number of possible states of a system.





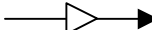
Figure Q-6: *basic entities of a viable system-producing unit as embedded System One element of the SiF.*

Usually, it is not practicable to count all possible states of a complex system, but it is possible in principle. Therefore, it is satisfactory to make comparative variety statements. Thus, we can already state that the square management box has a lower variety than the circle that contains the operations: management can never know everything that happens. Similarly, we can assert that the environment has a higher variety than the operational circle.

Consequently, in a viable system homeostatic regulation demands requisite variety⁵⁸ (RV) on the channels between the three basic entities:

1. high variety must be cut down or attenuated to the number of states that the receiving entity can actually handle;
2. low variety must be enhanced or amplified to the number of possible states that the receiving needs if it is to remain regulated.

The symbols used in the VSM to mark the attenuation of high variety input to a low variety entity, respectively the amplification of a low variety input to a higher variety entity, are borrowed from the electrician:

- attenuator: a device that reduces variety, depicted thus 
- amplifier: a device that increases variety, depicted thus 

The First Principle of Organisation states that managerial, operational and environmental varieties, diffusing through an institutional system tend to equate and should be *designed* to do so... with minimum damage to people and costs.

The elements in the Viable System Model (VSM)

The VSM consists of 5 systems, System One...Five that need each other to survive as a whole, see Figure Q-7. It is embedded in its next recursion 'higher up' and it contains the systems that produce its purposes at one recursion 'below'. Any Viable System (VS)

⁵⁸ Ashby's Law of Requisite Variety - stating that "only variety can absorb variety" - is a cornerstone in the Viable System: it has been proven that the variety of a system that can be disturbed must be met by the variety of a regulator of this system if full control of the system's behaviour is to be maintained.

interacts with its environment. The higher recursion environment encompasses the smaller environments of the embedded VS. The model is mathematically exact.




Figure Q-7: the viable system (after Beer, 1985)

System One is the set of viable systems (at one recursion level lower than that of the current 'System-in-Focus' (SiF) that produces the purposes of the SiF.

System Two is needed to damp oscillations which otherwise will occur in between the System One elements; System Two serves only this function!

System Three is the operational 'inside-and-now' management & services that enables System One to function.

System Three-Star is a means for System Three to enhance its variety by looking directly into operations of System One elements ('audits': to be used sporadically).

System Four is the 'outside-and-future' exploration of new opportunities with changing environments; this needs close and balanced channels with System Three.

System Five closes the loop of the recursive VS(M), and its main function is to maintain homeostasis between System Four and System Three. System Five may be alerted directly from System One through the algedonic channel.

The System 3-4-5 box is in itself a System One management unit of the next higher recursion, i.e. of the surgery Division B.

SiF CAPU System One

System One of CAPU is charged with conducting its operations with the purpose to perform cardiac surgery on a targeted number of qualified patients. There are four System One elements at level 1 of the SiF, CAPU: see Figure Q-3: intake, pre-op care, surgery and post-op care. Each of these elements themselves are viable systems, but at one recursion level down. The model of the CAPU system is depicted in Figure Q-8.

The Resource Bargain⁵⁹ normally comes with plans, budgets, programmes and procedures that are transmitted to the operational circles of System One. This transmission is an act of regulation and managerial variety amplification: the basic details of the Resource Bargain must be elaborated. At the same time, this regulation attenuates the operational variety, because potential operational variety must be harnessed to meet agreed objectives.

SiF CAPU System Three

SiF management and system supporting function are located in System Three. The location manager does the resource bargaining with the CAPU subsystems, holds these systems accountable and he also issues policy constraints or passes these on from higher recursions through the command channel⁶⁰. The resource bargaining with the System One elements is a two-way interaction. A homeostatic balance must be reached all the time or the System One will become disturbed.

SiF CAPU System Two

System Two stands for the operational regulatory control centres - the upwards pointing triangles - at the horizontal level in each System One operation, as well as between all System One elements and System Three. Its function is to co-ordinate operations in terms

⁵⁹ The Resource Bargain between System Three and System One is a bi-directional interaction that must be balanced (homeostasis). It is referenced as vertical connection number (ii), see Figure Q-8.

⁶⁰ The vertical command or Corporate Intervention channel is marked as (i) in Figure Q-8

of current needs. It enables various System One elements to sort out their own problems to a large extent and, as far as possible, solve conflicts between various System One units.

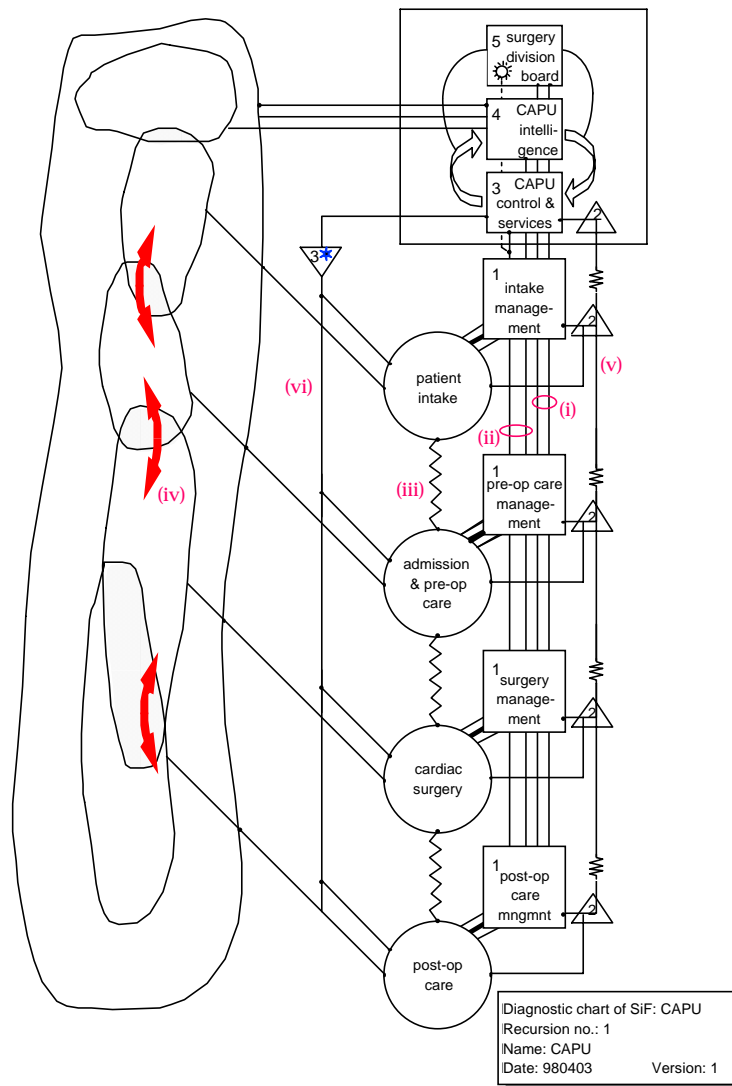


Figure Q-8: the Viable System Model of the cardiopulmonary surgery unit (CAPU) as the System in Focus (SiF), i.e. at recursion level 1. Note that the CAPU is defined by what it does: perform cardiac surgery.

Therefore, System Two must be designed to reduce high operational variety to the number of states that the receiving management can handle. At the horizontal level of a System One operation, System Two reveals operational state conditions to the local management that, in response, may have to intervene in order to maintain regulatory control. The management at System Three must cope with the variety of all System One elements.

Thus, System Two as vertical connection between System One and System Three must reduce variety from all System One operations in order to balance with the variety of System Three. In regulatory terms, the sole purpose of System Two is to damp oscillatory behaviour of the System One elements. System Two activities are based on relatively invariant rules for resource allocation, scheduling of System One operations, planning cycles, formats and coding schemes for upstream reporting, protocols, formal as well as informal procedures and directives, agendas and formal meetings and workplace bulletins.

The CAPU system has several System Two devices in place, which will be analysed separately below. Examples are the protocol⁶¹ pre- and post-op care at the regular ward and the MCU, scheduling schemes for surgery, the daily 7.30 h. shift transfer meeting of medical staff and residents covering all in-house CAPU patients, ethics and cultural mores, and - in the background - liability rules. Tracking of medical complication, abnormal situations and other indicators of operational performance disturbances are System Two activities⁶².

SiF CAPU System Three-Star

As stated before, the attenuation of operational variety by System Two must match the variety that the manager can handle⁶³. It should also trigger the manager to action in case of emerging operational instabilities that cannot be resolved autonomously by operational management. Usually, System One instabilities are being damped effectively by System Two devices that do not call the manager's attention all the time. The variety attenuation by System Two may, however, be too strong. When a System One oscillation reaches System Three management, then the manager may not have requisite variety (RV) to handle the high variety conditions effectively: he has insufficient actual insight in operational states of System One to understand the instability problem. The normal vertical channels filter out operational System One variety. System Three-Star provides the means for System Three management to overcome this low variety handicap. The 'audit channel'⁶⁴ enables System Three to penetrate straight into the operations themselves and, thus generate enormous variety. System One management should agree this. The power manifested in this way by System Three to System One management should only be employed sporadically⁶⁵.

⁶¹ This protocol has been issued in June 1997 for nursing staff and residents at the G3Z wards. It addresses diagnostic problems and medical treatment options for unscheduled interventions.

⁶² In this sense, most of today's so-called audit methods in the domain of safety and risk management, e.g. TRIPOD, ISRS, are System Two devices, because they do not peek into operational practice as is done through System Three-Star and, therefore, do not amplify the manager's perception of actual operational variety.

⁶³ Ashby's Law of Requisite Variety is a cornerstone in the Viable System: it has been proven that the variety of a system that can be disturbed must be met by the variety of a regulator of this system if full control of the system's behaviour is to be maintained.

⁶⁴ The System Three-Star audit channel is marked as (vi) in Figure Q-8.

⁶⁵ If System Three management uses System Three-Star frequently, the System Three practically takes over the operational management functions from System One. Due to the audits, the variety that the System Three management can employ regarding the System One operation meets that of the System One manager who has to accept the poking around in the System One operation by System Three management.

Samenvatting

Dit proefschrift beschrijft het zoeken naar en de resultaten inzake methoden en mechanismen die organisaties helpen om kosteneffectief te leren van (kleinschalige) ongevallen en andere ongewenste operationele verstoringen. Zoals blijkt is georganiseerd leren minder eenvoudig dan men zou kunnen denken. Kernpunten zoals *wie* zou moeten leren en *wat* er geleerd kan worden, zijn even relevant als topics als hoe effectief leren te organiseren en welke gegevens te verzamelen.

In elke bedoelde activiteit gaat een resterend (veiligheids-) risico of restrisico schuil, hoe goed deze activiteit ook is ontworpen en is gerealiseerd. In alle gevallen zijn ongevallen ongewenst en duiden op een afwijking in een anderszins gewenst, doelgericht proces. Zolang zulke ongevallen binnen de grenzen vallen van het restrisico, zou het optreden hiervan kunnen worden geaccepteerd, omdat op grond van nutsoverwegingen middelen die nodig zijn voor het voorkomen van zulke ongevallen, niet beschikbaar zijn gesteld. Het monitoren van incidenten is een manier om uit te vinden of restrisico's al dan niet groter blijken te zijn dan werk verondersteld. Ongevallen en andere *operationele verrassingen* worden gezien als ongewenste operationele omstandigheden die kunnen leiden dan wel feitelijk hebben geleid tot schade of andersoortige verliezen. De activiteiten veranderen onder invloed van zulke omstandigheden. Daardoor kunnen de werkprocessen gaan afwijken van de beoogde processen. Hoe eerder zo'n *afwijkingproces* kan worden onderkend, des te effectiever kunnen de beoogde processen worden hersteld. Operationele omstandigheden worden aangebracht en aangestuurd door een organisatorisch *stysteem*. Dit managementsysteem omvat de systematische organisatiefactoren die voorwaarden scheppen voor het optreden van operationele verrassingen. Deze aan de verrassing ten grondslag liggende factoren zijn degene die de moeite waard zijn om te identificeren in onderzoek van kleinschalige incidenten, omdat deze factoren de beste mogelijkheden bieden voor oplossing van problemen van kritische procesafwijkingen.

In dit proefschrift bespreek ik de ontwikkeling in het denken over het ontwerp van leersystemen gebaseerd op een epidemiologisch model voor het hanteren van ongevallendata tot systemen die ingebed zijn in expliciet beleid voor beheersing van risico's binnen de organisatie. Het is een evolutie van verzameling van ongevallendata voor data analyse achteraf, geleid door generieke schadescenario's in de data definitie, naar dataverzameling gebaseerd op een kennismodel inzake oorzakelijke organisatiefactoren zodat de principes van organisatorisch leren zich onmiskenbaar aandienen. Dit proefschrift volgt het proces van leren door middel van een serie projecten in uiteenlopende organisaties. Elk van deze projecten had ten doel systemen te introduceren om te leren van ongevallen en incidenten. In de loop van de tijd evolueerde en veranderde de wijze waarop dit plaatsvond, alsmede het hele ontwerp van deze systemen en de onderliggende filosofie. Het resulterende SINS-concept (SINS = Systematische Incident

Notificatie Systeem) werd geïmplementeerd in een hartchirurgieafdeling in haar meest ontwikkelde vorm. De incrementele inzichten in de vereisten voor effectief leren van zulke kleinschalige incidenten zijn op de harde manier verkregen. Principes van leren door organisaties meer dan door individuen binnen deze organisaties waren nodig om vooruitgang te boeken naar doeltreffende incident monitoring.

In de eerste projecten was de aandacht gericht op het verzamelen van data op grote schaal van vele ongevallen op zoek naar patronen in omstandigheden van ongevallen en scenario's van schadeprocessen. In de loop van de tijd en tijdens de projecten ontwikkelde deze benadering zich in naar beoordeling van elke gebeurtenis afzonderlijk. De introductie van gestructureerde kennismodellen inzake systematische factoren in het proces van incidentregistratie opende mogelijkheden tot leren van individuele gebeurtenissen. Met andere woorden, kosteneffectief organisatorisch leren van kleinschalige incidenten kwam binnen bereik. Verschillende projecten toonden de mogelijkheden in de praktijk aan van deze aanpak. Pas aan het einde van het laatste project werd de primaire vraag hoe leren te organiseren beantwoord op vanuit een theoretischer invalshoek. Op zoek naar meer fundamentele modellen en theorieën om de bevindingen te verklaren en te verankeren, legde ik het verband met Argyris' principes van organisatorisch leren [Argyris 92, 96] and met Beer's model voor levensvatbare systemen, bekend staand als 'Viable System Model' (VSM) [Beer 79, 81, 85; Kingston 96]. Dit maakte het mogelijk om de correcte organisatorische inbedding van op incidentdata gebaseerde regelsystemen vast te stellen.

De theorieën leveren veel inzicht in het succes en falen van de projecten en leiden tot een aantal heldere conclusies. Gezien het feit dat beschrijvende incidentmeldingen niet meer of minder zijn dan vereenvoudigde modellen van de werkelijkheid, moet de context van het incidentbericht bekend zijn om het bericht goed te verstaan en afdoende te interpreteren voor een doeltreffende oplossing van het in het bericht vervatte probleem. Uitgebreide incidentmeldingen zijn niet op voorhand nuttiger dan een bericht als "er is iets onverwachts gebeurd...!". Tot de sleutelbegrippen afkomstig uit de principes van Organisatorische Leren (OL) behoren het "leeragentschap" dat is benoemd om te leren ten behoeve van de organisatie, het "organisatorisch geheugen" om geleerde lessen op te slaan en te ontsluiten, het concept van de "Theorie van Actie" alsmede de onderscheidbare leerlussen.

The principes van Organisatorisch Leren zijn gebruikt om te komen tot een synthese en een beoordeling van een Systematische Incident Notificatie Systeem (SINS). Het leeragentschap werd een expliciet onderdeel in de configuratie van het systeem voor bewaking en terugkoppeling van incidenten, en vormde een sleutel tot de oplossing van het principiële probleem van verlies van contextinformatie in incidentmeldingen. Het ander winstpunt was de positionering van een organisatorisch geheugen dat de lessen bevat welke zijn getrokken uit eerdere ervaringen. De interactie tussen het leeragentschap en het organisatorisch geheugen is cruciaal voor de realisatie van een kosteneffectief systeem: hoe beter de geheugenfunctie is, des te minder middelen zijn nodig voor een leeragentschap om te functioneren. SINS werd gedeeltelijk geïmplementeerd in een chirurgieafdeling van een opleidingsziekenhuis. Het ziekenhuisproject leverde door feiten ondersteunde lessen op omtrent de mogelijkheden en valkuilen voor implementatie. Het toonde ook de zwakheden in het SINS-concept aan die nog een grondige analyse moeten worden onderworpen. Ondanks de incomplete implementatie, vormde dit PIERCE project een kritische beproeving van het SINS-concept.

Het 'Viable System Model' werd te laat gevonden om invloed uit te oefenen op de meeste van de beschreven projecten, maar het biedt een generiek raamwerk voor het begrijpen van het succes en falen van deze projecten en het generaliseren van de resultaten. Het levert een deugdelijk gereedschap om het SINS-concept alsmede lessen getrokken uit de andere projecten grondig te analyseren. In VSM perspectief wordt de positionering van de functies van incidentmonitoring- en terugkoppeling veel duidelijker. Zoals het is uitgewerkt, bied het VSM niet alleen de middelen om problemen in welk bestaand Systematische Incident NotificatieSysteem dan ook vast te stellen, maar kan het ook worden gebruikt om nieuwe systemen te specificeren, inclusief het type, de routing en de reductie van (context) variëteit van data in berichten. Het biedt een krachtig meta-instrument voor toekomstige werk.

Toekomstig werk zou zich moeten toespitsen op het ontwerp van een organisatorisch geheugen voor het opslaan van operationele verrassingen (of incidenten) en beheersingsmaatregelen welke zijn toegepast, en op de interactie met de gebruikers die verrassingen melden. Zo'n geheugen zou kunnen worden gebouwd met behulp van Case-Based Reasoning technologie. De voorgestelde en op VSM-gebaseerde methode voor het specificeren en diagnosticeren van organisatorische leersystemen moet worden gevalideerd. Met behulp van VSM, is een adequate specificatie, ontwerp en ontwikkeling van een systeem voor organisatorisch leren haalbaar en vraagt om een volledig uitgewerkte implementatie.

(this page is left blank intentionally)

Acknowledgements

This thesis has a long history going back a little less than twenty-five years. The then dean of the Electrical Engineering Department, Leo Krul, who took me on, has reminded me about my ambition to do a PhD whenever we have met since, keeping me keen in a special way. Jan de Kroes, I am proud that we could work and laugh together for so long, especially when building up the Safety Science Group, and for your shining vigour and example concerning safety in social - transportation - systems. Rogier van Wijk van Brievingh, my thanks for involving me in the field of hospital safety so early, and so persistently. My thanks also to other members of the Electrical Engineering & Safety Group within the EE Dept. To Anton Tjihuis, for his interest in risk management in the broad sense and particularly fundamentally in the case of heart valve fractures. To Philip de Graaf, Erik Vullings and Gerard van den Eijkel, for the short, but inspirational sessions during the Intelligent Anaesthesia Monitor project. I would also like to thank my Safety Science colleagues for their support over the years, especially Louis Goossens, Jurek Karczewski, Tom Heijer and Ben Quist for the discussions on basic issues regarding learning from small-scale incidents.

I am grateful to the members of the ARA task force in Poland who provided me unforgettable memories and down-to-earth insights in learning from occupational accidents in the early days of the transformation of Polish industry into modern enterprises. I also thank my temporary colleagues in the cardiopulmonary surgery unit (CAPU) of the AMC for their hospitality and their involvement in the PIERCE project, and especially Ruud de Graaf, dear "Watson". My appreciation also goes to my friends and colleagues in the on-site trade union group ABVAKABO FNV and others in the European Workshop on Industrial Computer Systems for their understanding of my priority settings in the last two years and their support to let me hang in and finish this thesis.

John Kingston, André Spijkervet and my old-time, still going-strong mentor Rien Buter: a special "thank you" for the sometimes very explorative discussions and the work we have done together: interesting to note that you independently developed your interests in Beer's viable systems modelling. I could not wish better support at the finale of my PhD work than from, Gerda Vermeer and Niek van Hussen, as my paranimfs and long-time close friends.

This thesis work would not have been realised without the lasting support and sustained faith of Andrew Hale, my promoter, during the ups-and-down of the series of projects that is discussed in the thesis. His constructive criticism and genuine interest in the matters raised have been a great help to me indeed.

The apparently unlimited energy of Bas de Mol, my co-promoter, to search for new allies and opportunities, despite his busy practice as cardiac surgeon, has been unbelievable and

stimulating. I appreciate that he opened up the opportunity of the final project in his hospital unit, as a result of which the bits and pieces left from previous projects came together in one piece.

The long-lasting support that I received from my family and friends and their understanding of work pressure during projects and certainly during the write up of the thesis is much appreciated. I count myself fortunate that my parents are able to witness my second graduation at Delft University. Pa en moe, I am grateful that you have always supported me and stimulated me to find my own way in life. I hope that you enjoy the degree-day and the graduation party.

Last, but not least, dear Maria, I am most grateful to you for all the support and understanding that you have given me over the years. Without your patience and endurance, this thesis would probably not have been finished. We will now have more time and opportunity to go out cycling in and beyond the Zoetermeerse polder.

Curriculum Vitae

Florus Koornneef was born on June 14, 1950, in Rotterdam, the Netherlands. After receiving his HBS-b diploma from the HBS Waldeck in The Hague in 1967, he studied electrical engineering at the Delft University of Technology and educational technology at the University of Amsterdam. He obtained his MSc degree in Electrical Engineering in 1975. From November 1975 to December 1980 he worked on the educational policy staff of the Electrical Engineering Department of the Delft University during the second wave of major university education reforms. He became involved in the initiation of the Safety Science Group that eventually started in 1980. He was also involved in the development of a series of course materials and courses on safety in health care. In 1981, he joined the Safety Science Group and helped to establish the interdepartmental Electrical Engineering & Safety Group. After the appointment of Andrew Hale as professor of Safety Science in 1984 they jointly developed the post-graduate accredited course on management of occupational safety, health and environment (MoSHE) that was run for the first time in 1988. In 1994, he ceased to be MoSHE course co-ordinator in order to focus on research projects on accident data registration and analysis that had been on-going since 1988. Inspired by Bas de Mol who joined the team as professor of Safety Science in Health Care Systems, he became in 1997 a collaborating staff member of the cardiopulmonary surgery unit of the Amsterdam Medical Centre in order to realise the PIERCE project. He is now a senior staff member of the Safety Science Group specialising in incident analysis.

Propositions
as a supplement to the thesis
Organised Learning from Small-scale Incidents
by F. Koornneef

1. In accident analysis, the term 'root cause' is a contradictio in terminis.
2. Reporters only report willingly when they think that the reported conditions can be changed.
3. Heinrich's Iceberg model only works when read from top to bottom.
4. The statements "Liability kills viability" and "Liability creates viability" are both valid. (*this thesis*)
5. However badly systems are designed, humans still make them work. (*Maturana et al. 1984*)
6. Organisational learning must be organised... (*this thesis*)
7. All accidents are unique; however, some are more unique than others. (*this thesis*)
8. The reaction of the Dutch Health Inspectorate to the case of the software flaw in the blood bank computer systems that allowed two blood types to be entered for one patient, has made clear that it is unaware of its System Three role in a viable Dutch health care system.
9. A bell ringer is a whistle-blower.
10. The FNV-information line, opened in spring 2000, shows that whistle-blowers are still often considered as the bad news instead of as an opportunity to strengthen the viability of the organisation.

Stellingen

behorende bij het proefschrift

Organised learning from Small-scale Incidents

door F. Koornneef

1. In ongevalanalyse is de term 'root cause' of 'basisoorzaak' een contradictio in terminis.
2. Rapporteurs rapporteren alleen dan als zij denken dat de gerapporteerde omstandigheden kunnen worden veranderd.
3. Heinrich's ijsbergmodel werkt alleen van top tot teen (top-down).
4. De uitspraken "aansprakelijkheid nekt levensvatbaarheid" en "aansprakelijkheid schept levensvatbaarheid" gaan beide op.
5. Hoe slecht systemen ook zijn ontworpen, mensen laten ze toch nog functioneren. (*Maturana et al. 1984*)
6. Organisatorisch leren moet worden georganiseerd... (*dit proefschrift*)
7. Alle ongevallen zijn uniek, echter, sommige zijn meer uniek dan andere. (*dit proefschrift*)
8. De reactie van de Inspectie voor de Volksgezondheid op de casus van de programmafout in de computer systemen van bloedbanken, waardoor twee bloedgroepen konden worden ingevoerd voor een patiënt, heeft duidelijk gemaakt dat zij zich niet bewust is van haar Systeem Drie-rol in een levensvatbaar Nederlands gezondheidszorgsysteem.
9. Een klokkenluider is een fluitist.
10. De in het voorjaar 2000 opengestelde FNV-meldlijn leert ons dat klokkenluiders nog vaak worden aangezien voor het slechte nieuws in plaats van een kans voor versterking van de levensvatbaarheid van de organisatie.