

Fair Mechanisms for Smart Grid Congestion Management

Hekkelman, B.

DOI

[10.4233/uuid:bb3f0f87-ea66-4e75-b414-21425344248e](https://doi.org/10.4233/uuid:bb3f0f87-ea66-4e75-b414-21425344248e)

Publication date

2022

Document Version

Final published version

Citation (APA)

Hekkelman, B. (2022). *Fair Mechanisms for Smart Grid Congestion Management*. [Dissertation (TU Delft), Delft University of Technology]. Delf University of Technology. <https://doi.org/10.4233/uuid:bb3f0f87-ea66-4e75-b414-21425344248e>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Fair Mechanisms for Smart Grid Congestion Management



Brinn Hekkelman

Fair Mechanisms for Smart Grid Congestion Management

Brinn Hekkelman

Fair Mechanisms for Smart Grid Congestion Management

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology

by the authority of the Rector Magnificus
Prof.dr.ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates

to be defended publicly on
Monday 12th of December 2022
at 12:30 o'clock

by

Brinn HEKKELMAN

Master of Science in Mathematical Physics,
University of Amsterdam, The Netherlands

born in Weststellingwerf, The Netherlands

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus	chairperson
Prof.dr.ir. J.A. La Poutré	Delft University of Technology Centrum Wiskunde & Informatica, promotor

Independent members:

Prof.dr. P. Palensky	Delft University of Technology
Prof.dr. F.P. Kelly	University of Cambridge
Prof.dr.ir. G. Deconinck	KU Leuven
Prof.dr. J.L. Hurink	University of Twente
Prof.dr. J.K. Kok	Eindhoven University of Technology
Prof.dr. F.M. Brazier	Delft University of Technology



The research in this dissertation was performed at the Intelligent & Autonomous Systems (IAS) group of Centrum Wiskunde & Informatica (CWI), the national research institute for mathematics and computer science in the Netherlands.



SIKS Dissertation Series No. 2022-33 - The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.



This work is part of a Perspectief programme of the former Dutch Stichting voor de Technische Wetenschappen (STW), since 2017 NWO Domain Toegepaste en Technische Wetenschappen (TTW).



This work is part of the research programme SES-BE with project number 14182 P13-21, which is partly financed by the Dutch Research Council (NWO).

Cover by Geert Mullens, printing by Ridderprint

ISBN 978-94-6366-632-9

Copyright © 2022 Brinn Hekkelman

An electronic version of this dissertation is available at

<https://repository.tudelft.nl>



Summary

We consider energy systems in the built environment. With the transition to a more sustainable, distributed, and ‘smart’ energy system, such local grids are undergoing significant changes. Among other developments, the new role of end-users as ‘prosumers’ – users that can either produce or consume power depending on the situation – is turning energy systems in the built environment into autonomous microgrids with complex internal interactions.

One of the primary challenges for these local grids is maintaining grid stability, which requires constant balancing of supply and demand. Because local grids were not designed for distributed energy generation and large loads such as electric vehicle charging, their limited capacity is now leading to congestion. Since the responsibility for resolving congestion falls increasingly on the individual prosumers and their flexibility, the concept of fairness must take a central role in congestion management.

In this dissertation we present our research on supply-demand matching mechanisms for fair congestion management. The local networks populated by users can be represented by radial multi-agent commodity flow systems. For the resource allocation problems in this setting we draw on the fields of mechanism design and fair division to design provably fair congestion management mechanisms. We evaluate the merit of different notions of fairness and present algorithmic mechanisms that align agent incentives with fair allocations.

We find that notions of fairness regarding congested commodity flow networks can either focus on local or global fairness. Agents can have differing opinions on the two, depending on how wide they draw the circle of peers that they compare themselves to. We find that the mix of producers and consumers requires slight adaptation of notions of fairness, with agents envying one group while welcoming the other. Furthermore, we find that it is possible to combine notions of fairness with welfare optimization by letting individual agents decide which of the two is more important, and protecting their fair shares.

We are able to use the radial structure prevalent in energy systems in the built environment to design algorithmic mechanisms of consistently low computational complexity. The congestion solutions of these mechanisms satisfy different local and global fairness criteria, for which we provide rigorous proofs. We prove that our mechanisms are individually rational and, for variations of egalitarian fairness, also incentive compatible. Finally, we introduce a congestion aftermarket where agents compensate their peers for flexibility.

Samenvatting

We richten ons op energiesystemen in de bebouwde omgeving. Met de transitie naar een duurzamer, gedistribueerd en ‘slimmer’ energiesysteem verandert er veel voor zulke lokale netten. Dankzij ontwikkelingen zoals de nieuwe rol van eindgebruikers als ‘prosumenten’ – gebruikers die zowel kunnen verbruiken als opwekken – veranderen energiesystemen in de bebouwde omgeving in autonome micronetten met complexe interne interacties.

Een van de belangrijkste uitdagingen voor deze lokale netten is het borgen van netstabiliteit, wat vereist dat vraag en aanbod constant gebalanceerd worden. Omdat lokale netten niet zijn ontworpen voor gedistribueerde energieopwekking en grote belasting zoals die van elektrische auto’s, leidt hun beperkte capaciteit nu tot congestie. Omdat de verantwoordelijkheid voor het oplossen van congestie steeds meer bij de individuele prosumenten en hun flexibiliteit ligt, moet eerlijkheid een centrale rol krijgen in congestiemanagement.

In dit proefschrift presenteren we ons onderzoek naar eerlijke congestie-managementmechanismen voor het koppelen van vraag en aanbod. De lokale netwerken kunnen we modelleren als radiale multiagentsystemen met interne stromen. Om bewijsbaar eerlijke mechanismen voor allocatie te ontwerpen, putten we uit de vakgebieden van mechanismeontwerp en eerlijk verdelen. We evalueren verschillende begrippen van eerlijkheid en presenteren algoritmische mechanismen die prikkels voor agenten in lijn brengen met eerlijke allocaties.

We zien dat eerlijkheid bij congestie in stroomnetwerken over lokale ofwel globale eerlijkheid kan gaan. Agenten hebben hier wisselende meningen over, afhankelijk van hoe groot ze de kring waarmee ze zich vergelijken trekken. We zien dat de mix van producenten en consumenten een lichte aanpassing in begrippen van eerlijkheid vraagt, omdat agenten jaloers zijn op de ene groep terwijl ze de andere verwelkomen. Verder zien we dat het mogelijk is om eerlijkheid te combineren met welzijns optimalisatie door individuele agenten te laten beslissen wat belangrijker is, en hun eerlijke deel te beschermen.

We gebruiken de radiale structuur die veel voorkomt bij energiesystemen in de bebouwde omgeving om algoritmische mechanismen van lage computationele complexiteit te ontwerpen. Hun congestieoplossingen voldoen aan verschillende lokale en globale eerlijkhedscriteria, wat we rigoreus bewijzen. We bewijzen dat voor onze mechanismen deelname rationeel is en, voor variaties van egalitaire eerlijkheid, ook waarheidsgetrouw is. Tot slot introduceren we een congestienamarkt waar agenten elkaar compenseren voor flexibiliteit.

Acknowledgments

Many people have helped me along my PhD journey. Closest to my work as a PhD has been my research group at the CWI. First and foremost I would like to thank my supervisor and promotor Han La Poutré for letting me explore and follow my way, and for the thorough readings, acute feedback, and great stories. I have always felt that working through your comments has really improved my work. I would like to thank Georgios Methenitis and Roland Saur for being my partners in crime during this PhD and for supporting each other in any way we could. I would like to thank Tim Baarslag and Michael Kaisers for being my academic older brothers, providing me with context and perspective on what being an academic is like. I would also like to thank all the other members of the Intelligent and Autonomous Systems group that have come and gone, and with whom I shared so many chats, laughs, lunches, and drinks. I have had a lot of fun spending time with every one of you!

During my PhD the CWI has been my home away from home. I am grateful for my comfortable office which I greatly enjoyed spending time in, and for all the other facilities that I was free to use. I would like to thank Isabella Pozzi for tremendous emotional support. I would like to thank Alvaro Piedrafita for introducing me to a fantastical world that I got to explore with colleagues from across the CWI. I would also like to thank all my fellow PhD's and postdocs for the camaraderie and all the support staff for making the CWI a great place to spend time. In addition I would like to thank Anneke Kuipers and the other staff that I encountered at TU Delft, who were always friendly and helpful.

Outside of my workplace I would like to thank Zerline Henning for being a true force of happiness in my life and for always being so amazingly kind and loving to me. I realise now that I likely could not have done it without you. I would like to thank Luuk de Gruijter and Mees de Vries for all the sneaky synergies between us. I would like to thank Ernst Schäfer and Jim de Groot for being my brothers from another mother, venturing out into the world. I would like to thank Wouter Hekkelman and Annemie Mullens for always thinking of me and for sending me as much help and love as is humanly possible. I would like to thank Coby Hekkelman, Jef Mullens, and Golly Kinable, all of whom I said goodbye to over the course of the PhD, for always taking a loving interest in my endeavours (and for keeping a secret). I would also like to thank all of my friends and family for being in my life; I am blessed to have you and have drawn much support from you during my PhD.

Contents

Summary	vii
Samenvatting	ix
Acknowledgments	xi
1 Introduction	1
1.1 Context of the Research	1
1.1.1 Topic of this dissertation	3
1.1.2 Formulating the research questions	5
1.2 Important Concepts in Power Systems	7
1.2.1 Traditional power system operation: physical aspects	7
1.2.2 Traditional power system operation: market aspects	9
1.2.3 Changes brought about by the energy transition	10
1.2.4 Congestion management	11
1.3 Important Concepts in Mechanism Design	13
1.3.1 Individual rationality	14
1.3.2 Incentive compatibility	14
1.4 Important Concepts in Fairness	15
1.4.1 Egalitarian fairness	16
1.4.2 Fairness in practice	18
1.5 Structure of the Dissertation	19
1.6 List of Publications	21
2 Measures of Fairness in Congestion Management	23
2.1 Introduction	23
2.2 Setting and Notation	24
2.3 Fairness Optimization and Measurement	25
2.4 Comparative Utility and Network Topology	27
2.5 Congestion Management Mechanism	31
2.6 Conclusions	35

3	Principal Notions of Fair Division and Local, Outer Matching	37
3.1	Introduction	37
3.2	Setting	39
3.3	Model	40
3.3.1	Congestion Management	40
3.4	Principal Notions of Fair Division	43
3.5	Local, Outer Matching and Fairness	44
3.6	Algorithm: Local, Outer Matching combined with Fair Division	49
3.6.1	Algorithmic Local, Outer Matching	49
3.6.2	Division for Different Notions of Fairness	51
3.6.3	Setting the Water Level	52
3.6.4	Result of the Algorithm	54
3.6.5	Example Congestion Solutions	54
3.7	Incentive Compatibility	56
3.8	Computational Complexity	58
3.9	Conclusion and Discussion	58
4	Parameterized Globally Max-Min Fair Solutions	61
4.1	Introduction	61
4.2	Preliminaries	63
4.2.1	Agent-Based Commodity Flow on Trees	63
4.2.2	Leximin Fairness for Commodity Flows	64
4.3	Main Theorems	66
4.4	Proof of the Main Theorems	66
4.4.1	Outline of the proof	66
4.4.2	Curtailment Allocations and Root Flow	67
4.4.3	Breaking Down in Binary Congestion Trees	70
4.4.4	Formulating the Induction Hypothesis	71
4.4.5	Verifying the Induction base	72
4.4.6	Performing Induction on Subtrees	72
4.4.7	Parameterization of Allocations	77
4.5	A Leximin Fair Mechanism	78
4.5.1	Curtailment using a Water Level Algorithm	78
4.5.2	A Mechanism for Leximin Fair Allocations	79
4.6	Properties of the Mechanism	82
4.6.1	The Computational Complexity is Low	82
4.6.2	The Mechanism is Individually Rational	82
4.6.3	The Mechanism is Incentive Compatible	82
4.6.4	The Mechanism as a Distributed Algorithm	83
4.7	Conclusions	83

5	Fairness, Welfare, and a Congestion Aftermarket	85
5.1	Introduction	85
5.1.1	Related Work	87
5.2	Preliminaries	88
5.3	Fair Share Congestion Management	90
5.4	Maximal Welfare Solutions	90
5.4.1	Maximal Welfare Congestion Algorithm	92
5.5	A Hybrid Solution With A Choice	98
5.6	An Aftermarket Pricing Scheme	101
5.7	Full Aftermarket Participation	110
5.8	Numerical Example	111
5.9	Conclusions	113
6	Conclusions	115
6.1	Answering the Research Questions	115
6.1.1	Answering our overarching research question	118
6.2	Outlook	119
A	The SES-BE Program	121
A.1	Research Program Outline	121
A.2	Individual Research Projects	122
	Bibliography	123
	SIKS Dissertation Series	133
	Curriculum Vitae	145

1 Introduction

The core of this dissertation is based on four research papers. This introductory chapter lays out the necessary groundwork for the reader to both follow the technical content of the scientific research as well as place it in its wider context.

1.1 Context of the Research

Since the beginning of the 21st century, the ‘energy transition’ of our energy system has been steadily gaining momentum. It entails a transition from traditional centralized demand-follows-supply systems that rely predominantly on large ‘grey’ power plants (e.g. coal), to distributed and ICT-connected energy systems that rely on ‘green’ power (e.g. solar) and the flexibility of their users. This transition is for a large part enabled by the increasing deployment of renewable energy sources (RES) such as wind and solar, together with new types of appliances such as electric vehicles, heat pumps, battery storage and other domestic electric usage technologies. Renewable energy sources are often intermittent and local in nature, while new electric usage technologies often allow flexibility in their consumption patterns. Such developments are transforming end-users into ‘prosumers’ that are able to switch between the role of producer and consumer depending on the local situation of generation.

The distributed nature of the future energy system has already started to become visible in society. After large windmills on the horizon and solar panels on building rooftops, in recent years we have seen a surge in solar fields for local communities and in electric vehicle ownership, sometimes coupled with sizeable local energy storage. This goes to show how many aspects of the energy transition have been picked up with enthusiasm and have made their way to integration in the built environment.

However, the limitations of our energy system are also beginning to show. Utilizing the full potential of flexible prosumers at all levels of the energy system will require infrastructure to gather and share detailed grid use information, as well as new ways to coordinate prosumers moment to moment using this information. But more tangibly, our current grid is unable to accommodate

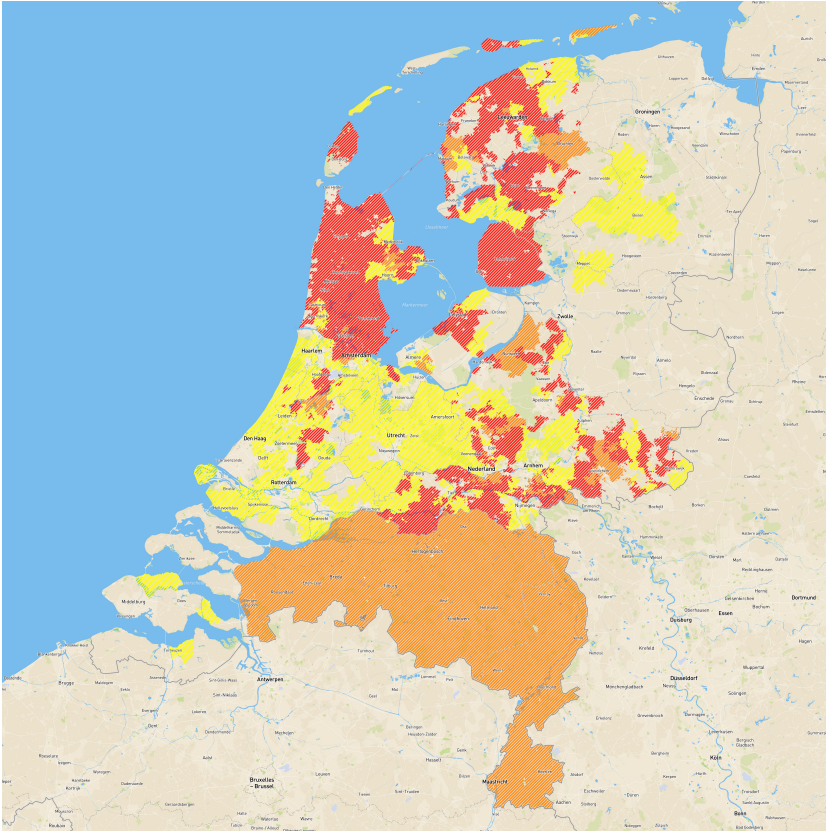


Figure 1.1 Capacity map for electricity demand in the Netherlands on August 8th, 2022 (Netbeheer Nederland, 2022). Color coding: Transparent: no congestion (yet). Yellow: congestion imminent. Orange: advance notice of structural congestion. Red: structural congestion, new applications for network use are denied.

all the new renewable energy initiatives. In the Netherlands, towards the end of 2018, new fully planned and funded local solar fields started being denied connection to the grid by the grid operators (van den Berg, 2019). The current grid capacity has proven insufficient to support new generation at the local level.

Figures 1.1 and 1.2 show grid capacity issues across the Netherlands in the summer of 2022. Figure 1.1 shows congestion on the demand side, which is mainly caused by electrification of industrial processes and transportation (e.g. electric vehicles). Figure 1.2 shows the severe and widespread congestion on

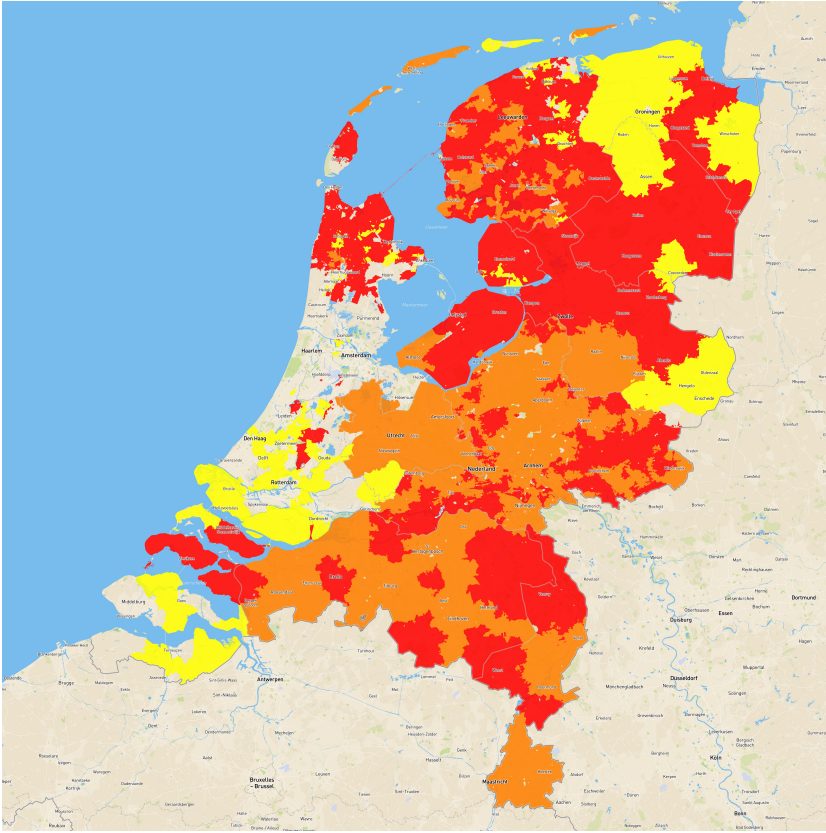


Figure 1.2 Capacity map for electricity supply in the Netherlands on August 8th, 2022 (Netbeheer Nederland, 2022). Color coding: **Transparent**: no congestion (yet). **Yellow**: congestion imminent. **Orange**: advance notice of structural congestion. **Red**: structural congestion, new applications for network use are denied.

the supply side, which is mainly caused by distributed and renewable energy resources. Both are especially prevalent at the medium voltage level.

1.1.1 Topic of this dissertation

We consider the problem of electricity network congestion management in the built environment within the context of the energy transition. Congestion management, i.e. reducing or resolving grid congestion, is an aspect of supply-demand matching which has always been fundamental to energy systems.

4 Chapter 1 *Introduction*

Unlike other goods, electricity cannot easily be stored and its production and consumption are continuously tied to each other. At each moment, supply and demand must be closely matched. Preferably, matching supply and demand should be done at the community level, to avoid transportation losses, decrease network congestion, and simplify the complex task of global balancing.

Traditionally, system operators continuously predict demand and plan appropriate supply from power plants. The generated power is sent through the transmission grid to the distribution grid in the built environment where end-users are located. With the energy transition, this situation is changing. End-users adopt the role of prosumers; users that can act as both producers and consumers, depending on the circumstances.

Moreover, the energy transition brings about a paradigm shift from ‘supply follows demand’ to ‘demand follows supply’. Where in the traditional centralized system generation followed end-user demand, in future decentralized systems the individual prosumers are tasked with so-called demand response to follow the volatile production of renewable energy sources. This places the consequences of resolving congestion increasingly on individual prosumers.

To adequately address this transition from a traditional centralized energy system towards more active participation of end-users, we will consider the future energy system as a multi-agent system: a system composed of multiple interacting intelligent agents. The agents represent prosumers which are autonomous decision-makers, and their decisions to take actions can be coordinated by a mechanism: a payoff structure for actions in an interactive setting. To replace the centralized system, we must thus design supply-demand matching mechanisms that incentivize agents to align their actions with the requirements of the grid.

More specifically, we must design such mechanisms to address the new challenges that energy systems in the built environment face. Moving from a centralized system to a prosumer-based one means that the direction and magnitude of power flow is no longer trivial. Complex power flows may now arise within the local grid, turning local grid topology into a significant and constraining factor. Where network congestion used to be a concern at the transmission level only, it is now also a concern within the local grid.

In recent years congestion management costs have soared (Nabe et al., 2017), even at the local distribution level (Haque et al., 2016). Research indicates that this congestion is not easily solved by storage (Härtel et al., 2016). Another potential solution is to increase grid capacity through physical grid reinforcement, i.e. putting larger power lines in the ground. However, this solution is expensive and above all time-consuming. Research suggests that coordinating network use to stay within existing capacity constraints is less costly than grid reinforcement (Okur et al., 2018a).

Therefore, we set out to design supply-demand matching mechanisms for multi-agent energy systems that take distribution grid topology into account and focus on local congestion management. Our multi-agent models should thus have a topological component, which is then reflected in the mechanisms. An agent's location and desired prosumption relative to other agents will determine the prosumption that they can realize within the grid constraints.

Immediately, this raises a question regarding fairness. When local congestion occurs and desired prosumptions cannot be realized, which prosumers are called upon to deviate from their desired prosumptions and by how much? Congestion, i.e. prosumption exceeding line capacities, by definition causes a conflict of interest. Traditionally, resolution of this conflict was the responsibility of grid operators, with ancillary services being called upon to accommodate end-user consumption.

With individual prosumers now being called upon to provide the flexibility necessary to resolve congestion, it becomes essential to explicitly incorporate the concept of fairness in the mechanisms governing the energy system. This requirement is highlighted by the European Commission, which states that *“energy is a critical good, absolutely essential for full participation in modern society. The clean energy transition also needs to be fair for those sectors, regions or vulnerable parts of society affected by the energy transition.”* (European Parliament, 2016).

Therefore, the topic we consider in this dissertation is that of fairness in congestion management mechanisms for the built environment. With this research we also contribute to the multidisciplinary Smart Energy Systems in the Built Environment (SES-BE) research program, see Appendix A.

1.1.2 Formulating the research questions

With this topic we set out to define our research questions. To summarize, we consider constrained systems populated by prosumers that are represented as autonomous decision-making agents. We seek to meet the fairness and congestion management requirements of energy systems in the built environment. We can use supply-demand matching mechanisms to try to align the incentives of the agents with these goals. Therefore, we formulate the following overarching research question:

In a constrained multi-agent power flow system, how can we define fairness and how can we subsequently design supply-demand matching mechanisms that manage congestion fairly?

To answer this question we must investigate which notions of fairness may apply to our setting. There are some unique aspects to consider. First, our setting has an important topological factor: an agent's role in congestion and

supply-demand matching is to a large degree affected by its predetermined location. Second, the interaction between supply and demand divides agents' dispositions towards each other into two categories: consumers compete with other consumers while they are aided by producers, and vice versa. This leads us to formulate the following research question:

Which factors determine how agents compare themselves to other agents, and how can we duly define, measure, and compute fairness among them?

The two mentioned aspects also affect how we approach congestion management. The topological factor means that congestion may occur in specific parts of the network and affect specific subsets of agents. Congestion may even occur simultaneously in multiple subnetworks across the network. The interaction between supply and demand also provides opportunities for resolving congestion. Local matching of supply and demand reduces interaction with agents elsewhere in the system and may therefore reduce strain on interconnecting lines. With these factors in mind we formulate the following research question:

How does local supply-demand matching affect congestion and how can we incorporate local balancing into notions of fairness?

The question regarding locality can also be drawn wider than the network under consideration. In electricity networks, different levels work together to create an interlocking system of energy generation, matching, and distribution. A local distribution grid, for example, acts in the transmission system with a single presumption that results from all its internal interactions. As we see, the situation of a subnetwork as a whole may depend on its interaction with a higher-level network. This then affects how the network can be fairly managed internally. Hence we formulate the following research question:

How can we adapt fair congestion solutions within a subnetwork to its collective interaction with a higher-level network?

Finally, we consider how fairness contrasts with welfare when congestion occurs. Since congestion implies a collective limitation that affects individual agents, we emphasized the importance of fairness. However, congestion also implies a market limitation that has consequences for welfare, both on a collective and individual scale. Where fairness considers the agents' relative presumptions, welfare considers the agents' total generated value. These contrasting views make it so that optimizing a congestion solution for either

one is unlikely to also optimize for the other. Consequently, we formulate the following research question:

How do fairness and welfare relate to each other in the presence of limitations imposed by congestion, and is there a way to reconcile them?

Before attempting to answer these research questions we will provide an overview of the most important concepts in our research, specifically in power networks, mechanism design, and fairness. With this basis the reader should then be well prepared to grasp the research that follows.

1.2 Important Concepts in Power Systems

Power networks have grown to become an indispensable cornerstone for our modern lifestyle. Especially in the built environment, access to the power system is never far away. Moreover, one can rely on the constant availability of power provided by this power system. In fact, nearly all systems, both large and small, among which many vital systems, are dependent on the omnipresent availability of power systems and their continuous reliability. In many places around the world power outages have become a rare phenomenon, with far-reaching consequences when they do occur.

Nevertheless, the ubiquitous availability and reliability of power systems is no small feat. In order to keep the system running, generation and consumption of electricity must constantly be balanced over the entire grid. A plethora of ancillary services and system operators constantly make sure that the balance is preserved and system constraints such as line capacities and voltage and frequency margins are satisfied. This all takes place within a volatile and uncertain context that involves open, although heavily regulated, markets and the constantly varying consumption by millions of power users.

1.2.1 Traditional power system operation: physical aspects

The electrical power system is made up of three primary components: production (e.g. generators, power plants), consumers (e.g. domestic use, factories), and the grid (e.g. electrical power lines). Production is traditionally associated with a relatively small amount of large power plants. Consumption, on the other hand, is associated with a large amount of small end-users (and some larger ones). The contrasting transportation requirements of large and small amounts of power are accommodated by a division of the grid into a transmission system and a distribution system.

8 Chapter 1 *Introduction*

The transmission system comprises a large, often national or even international, grid of high-voltage cables that transport large amounts of power. The high-voltage is to minimize energy loss across power lines due to resistance. Since power is a product of current and voltage, and line-losses scale quadratically with current, losses can be reduced by transforming the power to low-current high-voltage.

The distribution system comprises small local grids of medium- and low-voltage cables. The power from the transmission system is transformed to medium-voltage at substations, which is then transported close to end-users where it is further transformed to utilization voltage. Large power users may be connected directly to the medium- or even high-voltage networks.

Energy systems in the built environment are part of the distribution system. These local grids typically have a radial, i.e. non-cyclic, structure and are connected to the higher-level power system through a substation transformer. Due to the difference in scale, from the perspective of the local grid the transmission system can be seen as both an infinite source and an infinite sink, i.e. as a node capable of unlimited production and unlimited consumption. Interesting to note is that physically, local grids usually are not radial. There are cross-connections that can be opened in case of line failures in order to prevent parts of the grid from being completely cut off. However, during operation, whether normally or during an emergency, the active part of the grid will have a radial structure (Sallam and Malik, 2018).

The different levels of the power system are operated by different parties. The transmission and distribution systems are operated by a transmission system operator (TSO) and a distribution system operator (DSO) respectively. The TSO is often partly or wholly owned by the state or national government, and is responsible for the real-time stability of the power system and for coordinating supply and demand. The TSO is required to maintain a continuous (second-by-second) balance between electricity supply from power generators and demand from consumers. As part of this, the TSO procures ancillary services to support power system operation, e.g. generators that can quickly increase their production when needed. The TSO also has long-term responsibilities such as expanding the transmission network and planning new generation.

The DSO is responsible for the stability and maintenance of the distribution grid. Traditionally, the DSO essentially just makes sure that the incoming power from the transmission grid is properly distributed in the local grid. The DSO determines a configuration of the distribution system that minimizes power losses (Merlin and Back, 1975). Reconfiguration involves solving power flow equations that describe how power flows between nodes in the system to find an optimal power flow (OPF) (Peng and Low, 2014). Because our power systems use alternating current (AC), both active and reactive power must be

considered in an OPF. Active power actually transfers energy to the consumer, while reactive power travels back and forth over the lines and thus transfers no net energy to the consumer.

For further reading we recommend (Schavemaker and van der Sluis, 2017).

1.2.2 Traditional power system operation: market aspects

Operating on top of the physical power network we have the electricity market. This market system actually consists of several different markets that fill different roles. Closest to the end-user we have the retail market. Individual end-users purchase power from a retailer, usually through a contract with either a fixed or variable electricity price. These retailers are private parties who themselves participate in a more complex electricity market system along with large generators, large consumers, and ancillary service providers.

This electricity market system can be categorized roughly into three timescales: long-, medium-, and short-term. The long term markets include forward energy markets and capacity markets. These markets operate from about four years to one month before delivery and are used for example in the planning of new facility deployment and long-term contracts.

The medium term consists of the wholesale or spot markets, which include a day-ahead market and an intraday market. In the day-ahead market, parties purchase or sell power for time slots in the order of one hour for the 24 hours of the next day. The day-ahead market thus represents the planned production and consumption a day in advance. Its price is set by a classic supply and demand equilibrium price. After the day-ahead market is cleared, power can be traded in the intraday market up to in the order of one hour before delivery. In systems that use locational marginal pricing (LMP) the price is set per region of the network, based on the marginal cost of production in that region.

Since there is always uncertainty in the demand for, and consumption of, power, the planned consumption and production traded in the wholesale market is essentially never realised. To maintain supply demand balance and grid stability, deviations from the power traded in the wholesale market must be resolved in a short-term balancing market. The balancing market operates on intervals in the order of fifteen minutes up to real-time, and its prices can vary wildly. Designated balancing responsible parties (BRPs) use this market to negate deviations from their planned consumption or production. Each generator and consumer of electricity must have a contract with such a BRP, or be a BRP itself. The required flexibility, i.e. quickly consuming or producing on-demand for the balancing market, is provided by ancillary services.

Summarizing, three main types of participants in the wholesale and balancing markets can be discerned. Suppliers generate electricity from resources and sell it for a profit. Retailers and large consumers seek to predict their consumption for the next day as best they can, and profit from using or reselling

the purchased power. Ancillary service providers seek to take advantage of the volatile balancing market prices that arise from the constraint of constant supply demand balancing.

The TSO is tasked with maintaining grid stability across these markets, which is done through appointed BRPs and reserve power for any remaining system imbalance. The TSO also organizes the long-term markets to maintain long-term grid stability. In some countries, the TSO can use redispatch to increase or decrease output of specific generators to avoid regional grid congestion which is not taken into account in the wholesale market.

For further reading we recommend (Kirschen and Strbac, 2018).

1.2.3 Changes brought about by the energy transition

Since the beginning of the 21st century the so-called energy transition has started to gain momentum. The three primary developments that push the transition within the electricity network are the increasing penetration of renewable energy sources, the innovation of so-called smart grids, and the emergence of the so-called prosumer. Renewable energy sources such as wind and solar challenge the foundations of traditional power system management with their intermittent and distributed nature, while smart grid ICT systems elevate interconnectivity between power system users to a new level. At the center of the energy transition stands the flexible and digitally communicative end-user that can switch roles between consuming and producing: the *prosumer*.

With these developments, the tasks of the TSO are rapidly becoming significantly more complex. The intermittent nature of renewable energy sources is the cause of an increasing amount of uncertainty on the supply side. Where the TSO could traditionally plan the procurement of generation in advance with considerable accuracy, now unpredictable factors such as the weather introduce larger and more frequent inaccuracies. In addition, the distributed nature of many renewable energy sources means that the TSO can rely less on centralized control (Pepermans et al., 2005).

This movement from centralized to more decentralized control has a significant impact on DSOs. Where the DSO would traditionally just distribute the power coming in from the transmission network to its local users, it must now manage its networks as complete grids with potentially complex internal interactions between prosumers. Luckily, smart grid technologies such as smart meters and increased coordination through software communication can facilitate microgrid management (Fang et al., 2012).

On the wholesale electricity market, participants are also faced with consequences of the energy transition. Because large amounts of solar and wind power are generated under certain conditions and their operational costs are negligible, it can become difficult for traditional generators such as coal power plants to compete. Even negative prices can already sporadically be observed.

Retailers must adapt their predictive models to their customers' increasingly flexible demands that now at times even include production. Ancillary service providers must address the increasingly frequent and severe imbalances that result from uncertainty.

In general, there is an emphasis on the need for flexibility and on integration between levels of operation. The traditional centralized model cannot adequately manage the new distributed elements of the energy system. This has prompted the innovation of (market) mechanisms that coordinate between different system components such as local and national generation (Kok et al., 2005) or even combine different resources such as heat with electricity markets (Saur et al., 2019).

With the primary feature of the traditional system being centralized control of generation, we see that the energy transition is causing a paradigm shift from “supply follows demand” to “demand follows supply”. In other words: the flexibility of consumers is used to accommodate the intermittent supply of renewable energy resources. This paradigm shift manifests in the introduction of the concept of “demand-side management”. Demand-side management calls for a new approach to system control and end-user market participation (Methenitis et al., 2019a; Ramchurn et al., 2011).

1.2.4 Congestion management

With the energy transition, the nature of congestion in the electrical grid is changing. Power systems are fundamentally physical systems consisting of, among other components, metal cables called power lines. The amount of power that can be transported over such a cable is limited; if too much power flows over a cable it will heat up and eventually be damaged. In practice the most dramatic damage is avoided because the security systems will disconnect a cable before it passes its thermal threshold. The consequences of such an event may still be grave, since disconnecting the line means the power is diverted to other lines. This can lead to a cascade of disconnecting lines and ultimately a complete black-out of the power system.

Traditionally, congestion only occurs at the transmission system level. Here large amounts of power are transported over high-capacity power lines. The transmission grid spans large regions and is internally cross-connected. This grid structure allows the TSO to manage congestion by rerouting excessive power through less-congested regions, as well as procure regional ancillary services to mitigate congestion.

The ability of the TSO to manage congestion is the source of the so-called “copper plate” analogy. If the entire power system were one giant copper plate, power could flow unconstrained between producers and consumers regardless of their locations. In reality, the cost of electricity depends on the generators present in a region. In some countries the TSO's congestion management

efforts create the illusion of the copper plate. Other countries use for example locational marginal pricing (LMP) which does take inter-regional transmission system bottlenecks into account to arrive at regional electricity prices (Ma et al., 2003).

At the distribution system level, however, congestion is a relatively recent issue (Verzijlbergh et al., 2014). Renewable energy resources are being installed in increasing numbers near and even within the built environment. Distribution systems usually do not have the line capacity to support the relatively large amounts of local peak generation. Moreover, their predominantly radial structures do not provide the same rerouting opportunities found at the transmission system level.

There are three main approaches to distribution system congestion management. The first is grid expansion. However, laying down more and larger cables is a costly and above all time consuming process. At least on the short to medium term, alternative solutions are more efficient (Spiliotis et al., 2016). The second is increased control of when, where, and in what quantity power flows, for example through power storage. However, the relatively small buffer that storage provides can not easily resolve congestion (Härtel et al., 2016). The third is demand response and flexibility.

For congestion management of this last type, i.e. coordinating prosumers in the local network, the primary consideration should be active power curtailment (Bach Andersen et al., 2012; Hu et al., 2014; Rivera et al., 2015; Tonkoski et al., 2011; Verzijlbergh et al., 2014). Active power curtailment can take the shape of incentives provided by mechanisms or, more directly, the shape of quotas and direct control. In essence, active power curtailment is the allocation of a limited amount of capacity to a number of prosumers whose unaltered desired prosumptions would otherwise collectively cause congestion.

New solutions to distribution system congestion are being proposed (Philipsen et al., 2016; Xue et al., 2009), often focusing on demand-side management (Esmat and Usaola, 2016; Haque et al., 2017), electric vehicles (Li et al., 2014; Rivera et al., 2015; Shao et al., 2017), and/or decentralized multi-agent aspects of the problem (Ciavarella et al., 2019; Hu et al., 2015; Vandael et al., 2011; Vytelingum et al., 2010). Solutions often try to utilize the opportunities for local matching of supply and demand provided by the mix of prosumers found in these distribution systems. Local matching can be used to reduce strain on bottlenecks in the network.

With the paradigm shift from “supply follows demand” to “demand follows supply” and accompanying demand-side management, the consequence of resolving congestion lies increasingly with the prosumer; the individual end-user. Resolving congestion always entails a choice between which prosumers will sacrifice part of their desired prosumption. Since energy is a critical

good required for full participation in society (European Parliament, 2016), prosumer-based congestion management thus immediately raises an important question regarding fairness: who gets what?

1.3 Important Concepts in Mechanism Design

Secure and efficient operation of modern and future electricity distribution systems relies increasingly on coordinating the prosumption of individual prosumers present in the network. As opposed to the traditional centralized optimization approach, demand-side management requires consideration of the individual prosumers' motives and constraints. In smart grid energy markets, each prosumer can be represented as an autonomous decision-making agent. An agent is an independent intelligent entity that observes the state of its environment and subsequently takes actions in this environment based both on what it observes and its internal valuation of states. An environment with a number of agents acting and interacting within it is called a multi-agent system. In practice, prosumers in smart grids are normally represented by software agents: computer programs that act on behalf of a user.

Since prosumers are not directly controlled by the distribution system operator, coordination in a distribution system calls for a multi-agent approach (Vytelingum et al., 2010). Agents are at their core self-interested, and thus if certain behaviour is required from a system operation perspective, this behaviour must be incentivized. The design of a system of rules and rewards that incentivizes certain behaviour and leads to certain outcomes is called mechanism design. Mechanism design is sometimes referred to as “reverse game theory” (Jackson, 2014; Nisan et al., 2007): where game theory is about finding optimal agent strategies given the rules of a game, mechanism design is about designing the rules of a game in such a way that optimal agent strategies lead to a certain outcome for the system. A good mechanism aligns the incentives of the agents with the requirements of the system.

For our application of congestion management we will in specific consider resource allocation mechanisms. Markets are a form of resource allocation mechanisms that use prices as incentives. Locational marginal pricing (LMP) relies on such price incentives for congestion management; when congestion occurs due to excessive consumption, the price is raised until agents sufficiently reduce their consumption. This approach views the capacity at the bottleneck(s) as a scarce good and applies scarcity pricing accordingly. However, resource allocation mechanisms need not be market- or price-based.

In this work we will consider rational agents, i.e. agents that always take an optimal action given their information, as is usual in such settings. This is not necessarily the most accurate representation of reality; some agents, humans for example, may instead exhibit bounded rational behaviour (Methenitis et al., 2019b). Under the assumption of rationality, mechanisms can possess some useful properties, the two most important of which we will highlight here.

For further reading we recommend (Weiss, 2013) and (Nisan et al., 2007).

1.3.1 Individual rationality

Arguably the most important aspect of a mechanism is whether any of the agents that it was designed for actually *want* to participate in it. If, for example, an agent sets out to purchase a good but by doing so risks being forced to pay a price higher than it was prepared to, then that agent may think twice about attempting to make a purchase. In other settings agents may have no alternative to participating in the mechanism. Even then, it is important that agents are not worse off than trying to minimize their participation in the mechanism.

We say that a mechanism is individually rational when each participating agent is guaranteed to get a non-negative utility, where utility refers to the total satisfaction of, or usefulness for, the agent (Nisan et al., 2007). In other words, for each individual agent it is rational to participate in the mechanism. Individual rationality is also sometimes referred to as the participation constraint.

In our setting of energy systems and congestion management, individual rationality refers to two aspects. On the one hand it refers to the level of prosumption allocated to a prosumer. When a prosumer indicates a certain production we do not want the mechanism to ask the prosumer to produce more than they can nor do we want the mechanism to ask the prosumer to consume instead, and vice versa. On the other hand individual rationality refers to the monetary value for the prosumer. We do not want a prosumer to pay more than the received power is worth to them, nor do we want a prosumer to receive less for sold power than it cost them to generate it.

1.3.2 Incentive compatibility

Another crucial aspect of a mechanism is whether the participating agents are actually playing by the rules. Since multi-agent mechanisms can constitute strategic settings, and agents are self-interested, it may very well be that agents attempt to strategically take actions that increase their own utility at the expense of others. This is observed, for example, in traditional card games where players are often incentivized to bluff or otherwise misrepresent the information they hold in order to get an advantage over other players. Of course, in such games this strategic interaction is a feature. However, in a mechanism intended for social choice, i.e. aggregation of the preferences of

the agents toward a single joint decision like division of limited capacity, such strategic interactions can be detrimental.

We say that a mechanism is incentive compatible when each participating agent cannot strategically improve their situation by intentionally misreporting their preferences or other hidden information (Nisan et al., 2007). In other words, the incentives that the mechanism provides for the agents are compatible with the rules. An incentive compatible mechanism is also called truthful or strategyproof, referring to the fact that it is in an agent's best interest to participate truthfully as opposed to acting strategically.

Different degrees of incentive compatibility can be discerned. From the perspective of an agent, being truthful can either be a strongly dominant strategy or a weakly dominant strategy. We say that being truthful is a strongly dominant strategy if the agent will always be worse off when playing a strategy that is not truthful. We say that being truthful is a weakly dominant strategy if the agent is never better off when playing a strategy that is not truthful, i.e. being truthful is at least as good as any other optimal strategy. Even though these definitions subtly differ, they both define what we call dominant-strategy incentive compatibility, which incentive compatibility usually refers to. A weaker degree of incentive compatibility is Bayesian-Nash incentive-compatibility, which only requires being truthful to be a best strategy given that all other agents are also participating truthfully.

In our setting of energy systems and congestion management, incentive compatibility can refer to two aspects. On the one hand it can refer to the reported level of desired prosumption of a prosumer. We want prosumers to report their actual desired prosumption, and not an artificially inflated one aimed at receiving a larger share of available capacity when prosumptions are curtailed (i.e. instructed to reduce). On the other hand incentive compatibility can refer to the utility of the prosumer. We want prosumers to communicate their actual valuations of consumed or produced power.

As a final note we emphasize the importance of incentive compatibility in mechanisms for which fairness is a major factor. If such a mechanism is not incentive compatible, agents can "game the system" in order to improve their situation relative to other agents. Since the relative situations of agents are an integral part of fairness, the existence of such strategies would significantly undermine the mechanism's claim of fairness.

1.4 Important Concepts in Fairness

The allocation of limited resources to different agents raises an important question: what is a fair share for each agent? This problem of fairly dividing limited resources is an active field of research in political science, economics,

mathematics, and computer science (Procaccia, 2013). But the first question we must ask is; what is fairness?

The concept of fairness is a fundamentally subjective one. At its core, it relates to the opinion that agents have of their share relative to the share of other agents. Without comparison of shares the concept of fairness is meaningless. The study of such opinions on relative shares is known as social comparison theory, which originates from the work of Leon Festinger (Festinger, 1954). In the fields of psychology (Saxena et al., 2019) and behavioural economics (Fehr and Gächter, 2000; Fehr and Schmidt, 1999), researchers observe the opinions that humans inherently hold in different situations and attempt to distill from them so-called notions of fairness.

In the fields of mathematics and computer science we take a more rigorous approach. While it is futile to try to prove that any notion of fairness is objectively fair, what we *can* do is prove that a specific notion of fairness holds for a given allocation. Another thing we can do is, for a specific notion of fairness, measure how fair a certain allocation is by assigning a score to it, again without attempting to prove some objective fairness of the notion.

This rigorous study of fair division originates from the work of Hugo Steinhaus. Steinhaus considered how to divide a heterogeneous resource among several agents with different preferences such that every agent believes that they received a proportional share, resulting in the notion of proportional fairness also known as ‘simple fair division’ (Steinhaus, 1948). This work initiated the research field of fair division of divisible goods that became known as cake-cutting (Alon, 1987; Brams and Taylor, 1995; Dubins and Spanier, 1961; Stromquist, 1980).

The modern field of fair division comprises many different aspects, ranging from public decision making (Conitzer et al., 2017) to information and communications systems (Gutman and Nisan, 2012), and from multi-resource allocations (Parkes et al., 2012) to the interaction between global allocations and locally perceived fairness (Abebe et al., 2017). A large body of work in the fair division field has been developed on the topic of dividing indivisible goods rather than divisible ones (Caragiannis et al., 2019; Lang and Rothe, 2016). In our present work, however, we consider flow of electrical power which is a commodity and hence a divisible good.

For further reading we recommend (Brams and Taylor, 1996), (Robertson and Webb, 1998), and (Moulin, 2003).

1.4.1 Egalitarian fairness

Among all the different notions of fairness we want to highlight a particularly intuitive and ubiquitous one: egalitarian fairness. The egalitarian notion of fairness is captured by a few different definitions, the connections between which we will touch upon. The central idea behind egalitarianism is that all

individuals are equal and should be treated as such. This concept has varying levels of usefulness across different situations, most often deriving from, for better or worse, the fact that individual aspects and preferences of agents are essentially ignored (Pazner and Schmeidler, 1978).

One of the first forms of egalitarian fairness appearing in the fields of resource allocation and fair division is that of envy-freeness (Foley, 1967; Gamow and Stern, 1958). An allocation is said to be envy-free if no agent would prefer another agent's allocated share over their own. In settings where a homogeneous divisible good is divided over agents that all value the good equally, e.g. money from an inheritance, an envy-free solution allocates an equal share to each agent.

Envy-freeness can also be applied to situations where agents do not share the same valuation of goods, for example when dividing sets of indivisible goods that are valued differently by different agents. Since in these settings an envy-free allocation does not always exist, the adjacent definition of "envy-free up to one good" may be preferred there (Lipton et al., 2004).

Reaching an egalitarian (or envy-free, in homogeneous settings) allocation may be thought of as taking any allocation and then equalizing the agent shares, reducing envy. This approach takes from the Pigou-Dalton principle in welfare economics (Dalton, 1920; Pigou, 1912). The Pigou-Dalton principle is a condition on social welfare functions that is satisfied if the function prefers allocations that are more equitable. In other words, a function (strictly) satisfies the Pigou-Dalton principle if it (strictly) increases when a change of allocated goods reduces the allocation of 'rich' agents (i.e. larger allocations) while increasing the allocation of 'poor' agents (i.e. smaller allocations), as long as this does not now make the rich agents poorer (i.e. allocated less) than the poor agents.

This principle is implemented by the widely used notion of max-min fairness (Jaffe, 1981). An allocation is said to be max-min fair (or, more specifically, leximin fair) if an increase of the allocation of one agent necessarily results in a decrease of the allocation of another agent with an equal or smaller allocation. Max-min fairness can technically be taken to require only the name-sake maximization of the minimum value, but the term is popularly used synonymously with leximin fairness which is a formal extension of the same principle. Leximin fairness is a very strict criterion that implies lexicographically maximal vectors (Pióro and Medhi, 2004). In the lexicographic order, whether a vector is larger than another is determined by comparing their first two elements, and then their second two and so on if necessary. If these vectors are both sorted in ascending order, then the one vector is said to be leximin-larger than the other.

The leximin order of vectors can be used to compare different allocations by comparing leximin vectors of the values allocated to agents. This defines the leximin rule that selects the leximin-optimal vector among all possible allocations. This rule is also called the egalitarian rule and defines the leximin criterion or leximin fairness (Barbarà and Jackson, 1988). The leximin-optimal allocation is always Pareto efficient, meaning that an increase of the allocation of one agent necessarily results in a decrease of the allocation of another agent.

Finally, we have the Nash product of allocated values as a measure and optimization goal for fairness (Nash, 1950). Aside from issues when at least one value is zero, the Nash product can be shown to satisfy the Pigou-Dalton principle. The well-known issue regarding values of zero, which collapse the entire product to zero, set it apart from the leximin criterion which strictly satisfies the Pigou-Dalton principle. The Nash product, however, is extremely useful and versatile (Caragiannis et al., 2019), and a maximal nonzero Nash product implies envy-freeness and Pareto optimality (Varian, 1974).

1.4.2 Fairness in practice

The study of fairness and fair division in social choice theory is highly relevant in many domains of application such as allocation of natural resources, voting mechanisms, network engineering, and more. Even though fairness is subjective in a larger context, provably fair solutions can help groups of agents agree on solutions and provide tangible criteria on which decision-making can be based. An example of such initiatives is presented by (Kurokawa et al., 2015).

Although the scientific developments in social choice theory have been plentiful, adoption by society and the wider public seems to have lagged behind. One initiative that aims to address the limited integration into practice is the website www.spliddit.org (Goldman and Procaccia, 2015). This website collects and offers provably fair mechanisms for everyday fair division problems, i.e. practical problems faced by the general public.

The most important application for our present work is in flow networks (Kleinberg et al., 1999; Ogryczak et al., 2014; Pióro and Medhi, 2004). Classical flow problems reduce the model to a single source and sink with flow passing over a network of connections between the two. However, in order to be able to consider fairness, we must look at and compare the individual sources and sinks around the network (Megiddo, 1974). The fact that the topology of the system plays such a defining role in the constraints and asymmetric situations of agents makes flow networks a unique type of fair division problem.

Fairness applications in flow networks can be roughly divided into two categories: information flows and commodity flows. For both applications the underlying flow network model is the same and in many cases the application can be abstracted away. The primary difference is that local matching of supply and demand is possible with commodity flows. With information flows each

piece of information has a designated sender-receiver pair associated with it, preventing arbitrary re-matchings among the agents (Kleinberg et al., 1999).

For a survey on fair division in flow systems see (Fossati et al., 2018).

1.5 Structure of the Dissertation

In **Chapter 1** we provided a brief context for the research presented in this dissertation, as well as touched upon most of the important concepts. We then took a deeper look into the three areas of interest most relevant to the scientific research presented later on. Specifically, we highlighted important concepts in power networks, in mechanism design, and in fairness. With this basis the reader should be well prepared to follow the why and how of the research that follows.

In **Chapter 2** we explore notions of fairness in the context of topologically constrained multi-agent commodity flow systems. We consider different ways of adapting existing notions of fairness to the topologically constrained case. Here we investigate how prosumers may regard the prosumption of other agents when taking into account their relative locations and prosumptions. When congestion occurs locally and negatively affects prosumers in an area, it may still be considered fair by all prosumers that prosumers elsewhere in the network are not negatively affected.

Moreover, we consider models from behavioural economics that describe the inherent sense of fairness observed in humans. We translate the observation-based models into utility functions that are compatible with our setting of agents in a topologically constrained network.

Finally, we design an algorithmic mechanism for fairly resolving congestion in consumer-only energy systems modelled after the built environment. We prove desirable properties of the mechanism and show a low computational complexity.

In **Chapter 3** we dive into the congestion-resolving opportunities that local matching of supply and demand provides. When congestion occurs it is because a group of prosumers together exceeds one or more capacity constraints of the network. However, there may be prosumers present in the area whose prosumption contributes to resolving the congestion, i.e. producers amid an excess of consumption or consumers amid an excess of production.

We design a congestion management mechanism that prioritizes local matching in the radial networks that are typical for the built environment. Our mechanism locally matches supply and demand from the outer edges of the network towards its center. This outer matching approach locally resolves congestion as much as possible and allows prosumers to benefit from the counterbalance provided by those prosumers closest to them.

We combine this outer matching mechanism with three interchangeable principal notions of fair division: proportional, egalitarian, and nondiscriminatory. Moreover, the mechanism is compatible with any well-defined division rule. We prove that some properties of the mechanism depend on the chosen notion of fair division, and again show a low computational complexity.

In **Chapter 4** we investigate the impact of a local energy network's hierarchical connection to a larger network. Issues in the larger network may lead it to demand a certain power flow from the local network as a whole. Conversely, the local network may want to realise a certain power flow between it and the larger network in order to take advantage of opportunities in the larger energy market.

We design a congestion management mechanism whose congestion solutions are fully parameterized: for each possible power flow over the connection to the larger network that the local network can realize, a single parameter suffices to fully determine the solution for all prosumers or any individual prosumer. As a consequence, the mechanism can be executed once to parameterize congestion solutions for all possible power flows of the local network as a whole. Once a certain power flow is decided upon, the congestion solution can then be computed in $\mathcal{O}(1)$ time for each agent.

We prove that the unique congestion solutions provided by our mechanism are leximin fair, taking into account the topological constraints as well as supply-demand matching. Unlike greedy matching of supply and demand between agents closest to each other, a leximin fair congestion solution matches supply and demand such that no prosumer's situation can be improved without worsening the situation of a prosumer that is less well-off. In a way, such a solution is the most equitable solution that is attainable within the capacity constraints.

In **Chapter 5** we consider the trade-off between fairness and welfare. Where notions of fairness compare the allocated levels of prosumption between agents, welfare looks at the cumulative value generated for the agents by the allocation. Since these are two different objectives, it is unlikely that a fair congestion solution also maximizes the total welfare of the agents. Both outcomes are desirable from different perspectives, and therefore it is not immediately clear which one should be emphasized.

We design a congestion aftermarket that aims to reconcile the two objectives: fairness and welfare. First, we compute a fair congestion solution which determines each prosumer's 'fair share'. Then, we let each prosumer choose whether to claim their fair share or participate in the congestion aftermarket. Finally, we clear the congestion aftermarket using a welfare-maximizing mechanism and determine the aftermarket prices for each agent.

Because prosumers get the option to claim their fair share at the original market price, the mechanism can arguably be said to be fair. At the same time, we prove that the aftermarket is such that it is economically optimal for each agent to participate in it. Our combination of congestion management mechanism and congestion aftermarket allows both fairness-seeking and welfare-seeking prosumers to exist in parallel without interfering in each other's goals.

In **Chapter 6** we summarize the results of the research carried out. Here we evaluate our findings in the context of the research topic and look out to future research. We return to the research questions and evaluate which answers we found for them.

1.6 List of Publications

In this section we present an overview of the research papers that comprise the dissertation. Chapters 2 to 5 are each based respectively on one of these research papers.

- Brinn Hekkelman and Han La Poutré. 2019. **Fairness in Smart Grid Congestion Management**. *IEEE PES Innovative Smart Grid Technologies Europe (ISGT-EU)*. <https://doi.org/10.1109/ISGTEurope.2019.8905496>
- Brinn Hekkelman and Han La Poutré. 2020. **Fairness in Power Flow Network Congestion Management with Outer Matching and Principal Notions of Fair Division**. *ACM International Conference on Future Energy Systems (e-Energy)*¹. <https://doi.org/10.1145/3396851.3397701>
- Brinn Hekkelman and Han La Poutré. 2023. **Fair Parameterized Congestion Management for Commodity Flow Subnetworks**. *Under review*.
- Brinn Hekkelman and Han La Poutré. 2022. **Fairness vs Welfare: a Hybrid Congestion Aftermarket**. *ACM International Conference on Future Energy Systems (e-Energy)*. <https://doi.org/10.1145/3538637.3538843>

¹ An extended abstract of this work was presented at Dutch national conference ICT.OPEN2021.

2

Measures of Fairness in Congestion Management

In the previous chapter we provided a short overview of the study of fairness. In this chapter we add to this field by proposing notions of fairness tailored specifically to network congestion problems.

2.1 Introduction

The energy system is going through a transition. This energy transition is brought about by an increasing penetration of renewable energy sources and a push towards a more decentralized system. With these developments, congestion on electrical grid lines is becoming a more widespread problem (Verzijlbergh et al., 2014); one that is not easily solved using storage (Härtel et al., 2016). The intermittent nature of renewable energy resources, the decentralized nature of consumers and producers (often *prosumers* now), and the intensive disruptive demand introduced by electric vehicles and heat pumps all contribute to grid congestion issues. According to a study concerning the German electrical grid; “over the past five years, the costs for congestion management and curtailment have increased by a factor of ten, to about one billion euro per year.” (Nabe et al., 2017).

Grid congestion management solutions appear in various forms and address different aspects of grid congestion problems (Esmat and Usaola, 2016; Haque et al., 2017; Li et al., 2014; Philipsen et al., 2016; Xue et al., 2009). However, while grid constraints raise questions concerning priority when conflicts of use arise, these studies on congestion management do not take into account an explicit notion of fairness. A recent package of measures presented by the European Commission states that “energy is a critical good, absolutely essential for full participation in modern society. The clean energy transition also needs to be fair for those sectors, regions or vulnerable parts of society affected by the energy transition.” (European Parliament, 2016). In light of this statement, the incorporation of fairness is left insufficiently covered by grid congestion management research.

The incorporation of fairness in grid congestion management is no straightforward task. Notions of fairness are fundamentally subjective, and accepted notions of fairness do not necessarily translate from one setting to another. Moreover, other goals such as efficiency may take precedence over fairness, limiting the scope of fairness that can be implemented. Once a notion of fairness has been accepted for a certain setting, it can serve one or more of the following three main uses:

- As a binary descriptor; it is expressed qualitatively and its definition either is or is not satisfied by a situation.
- As a tool for maximizing fairness; it is expressed quantitatively and may be used as an optimization goal for fairness and other qualities.
- To compare and evaluate situations; it is expressed quantitatively, preferably normalized, and measures fairness independent of other qualities.

This chapter proposes two implementations of fairness suitable for congestion management in electrical grids. Both of these implementations will be of the quantitative type and may be used as an optimization goal or a measure of fairness. The first implementation of fairness that this chapter proposes is based on the Nash product that was introduced by Nash (Nash, 1950) and closely resembles the notion of social welfare. The second implementation of fairness that this chapter proposes is based on research in behavioral economics by Fehr and Schmidt (Fehr and Schmidt, 1999), and mimics the inequity-based comparative utility (inequality between agents negatively affects their utilities) that is observed in humans.

Furthermore, this chapter presents a congestion management solution in the form of an egalitarian allocation mechanism. Based on consumer data, this mechanism allocates consumption limits to individual consumers in order to resolve congestion in acyclic networks. Finally, this chapter proves that the presented mechanism is truthful and maximizes both the social welfare and Nash product.

2.2 Setting and Notation

Consider a network \mathcal{N} , represented by a graph (V, E) . Agents $a \in A$ representing prosumers, i.e. agents that may either produce or consume at any given time, are located at the vertices $v \in V$. The edges represent electrical grid lines $l \in E$ with positive capacity constraints C_l that constrain the power flow over the line l . A connection to an external network may be represented by an edge associated with only one vertex. Prosumption of an agent a is represented by an activity y_a , with positive and negative activity corresponding to consumption and production respectively. Let A denote the set of all agents in

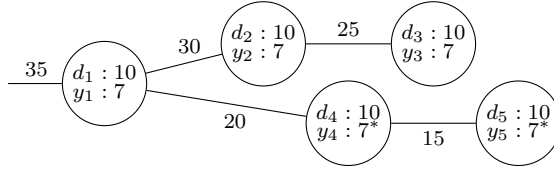


Figure 2.1 An example network of five agents named 1 through 5. Allocation 1 is as shown, while allocation 2 differs by instead setting $y_4 = 9$ and $y_5 = 5$.

the network, and let A^+ and A^- denote the subsets of consumers and producers, respectively.

When congestion occurs, i.e. at least one of the line capacity constraints C_l is exceeded by the power flow over that line, a congestion management mechanism resolves the congestion. It does so by finding an allocation Y : a set of activities y_a for the agents $a \in A$. As a result, for some agents a , there will be a difference between the agent's reported desired activity d_a , and its final activity y_a . A reported desired consumption, i.e. $d_a > 0$, must always result in a final activity $0 \leq y_a \leq d_a$. Similarly for production.

Each agent a has a utility function u_a that depends on the activity y_a , the valuation of the activity $\lambda_a(y_a)$, the price of the activity $p_a(y_a)$, and the desired presumption d_a . This chapter assumes a setting in which, for a network of limited size, the price function p_a is identical for all agents $a \in A$ and scales linearly with the activity y_a . Furthermore, the valuation function λ_a is identical for all agents $a \in A$ and scales linearly with the activity y_a . Therefore, in this setting, the utility function u_a only depends on at most the activity y_a and desire d_a .

2.3 Fairness Optimization and Measurement

A common way to optimize allocations of a divisible good to a set of agents, is to maximize the social welfare (SW). This means maximizing the sum over all agents' utilities:

$$\max_{Y \in S} \sum_{a \in A} u_a(Y). \quad (2.1)$$

Here, S denotes the solution set: the set of allocations that resolve congestion and assign consumption and production exclusively to consumers and producers respectively, bounded by their desires, as described in Section 2.2.

Since fairness is an inter-agent concept, *SW* cannot take fairness into consideration without explicitly incorporating it in the individual utility functions; Section 2.4 further considers this option.

An alternative optimization goal is the Nash product (*NP*):

$$\max_{Y \in S} \prod_{a \in A} u_a(Y). \quad (2.2)$$

This optimization problem, like the *SW* optimization in (2.1), maximizes all agents' utilities, within S . However, unlike in (2.1), the *NP* optimization in (2.2) also maximizes the minimal utility among agents, within S . The differences are illustrated with the help of a running example, shown in Figure 2.1. Taking the simplest utility function for each agent $a \in A$, i.e.

$$u_a = |y_a|, \quad (2.3)$$

the *SW* and *NP* values associated with the two allocations presented in Figure 2.1 are displayed in Table 2.1. Table 2.1 shows that the *SW* approach does not differentiate between the two allocations, while the *NP* takes a higher value when the allocated activities are closer to each other. Note that both allocations yield the same total consumption.

In order to have the *SW* and *NP* not only serve as an optimization goal but also as a measure, i.e. an indicator independent of irrelevant qualities, the average is taken. This eliminates their dependency on the size of the system. For the average Nash product (*ANP*), the optimization then takes the following form:

$$\max_{Y \in S} m \sqrt[m]{\prod_{a \in A} u_a(Y)}, \quad (2.4)$$

where m is the number of agents in A . Note that taking the average does not affect the optimization problem. The values of the averaged social welfare (*ASW*, see Table 2.1) and *ANP* on the two allocations presented in Figure 2.1 are also displayed in Table 2.1.

The *ANP*, however, still depends on the absolute level of presumption. This dependency can be removed by considering utility functions that reflect the relative activity instead of the absolute activity, that is

$$u_a = y_a/d_a. \quad (2.5)$$

This results in the normalized Nash product (*NNP*) that takes values between 0 and 1. Since the *NNP* is largely independent of qualities other than fairness, it is well suited as a measure of fairness. The values that the *NNP* takes on the two allocations presented in Figure 2.1 are displayed in Table 2.1 as well.

		Allocation 1	Allocation 2
<i>SW</i>	$\sum y_a $	35	35
<i>NP</i>	$\prod y_a $	16807	15435
<i>ASW</i>	$\frac{1}{m} \sum y_a $	7	7
<i>ANP</i>	$\sqrt[m]{\prod y_a }$	7	6.88
<i>NNP</i>	$\sqrt[m]{\prod y_a/d_a}$	0.7	0.69

Table 2.1 Social welfare and Nash product values for the two allocations presented in Figure 2.1.

The downside of the relative utility function (2.5) is that its value is influenced to a large extent by the desire of the agent. This means that an agent a intentionally reporting a large desire d_a affects the value of the *ANP* significantly. Moreover, it skews the optimization problem to allocate a potentially disproportionate amount of activity to such an agent a . Thus, while the *NNP* provides a fine measuring tool, the absolute utility function (2.3) is better suited as an optimization goal.

2.4 Comparative Utility and Network Topology

An alternative approach to incorporating fairness in congestion management, is to explicitly include a notion of fairness in the individual utility functions of the agents $a \in A$. Research in behavioral economics by Fehr and Schmidt (Fehr and Schmidt, 1999) proposes a model aimed at capturing fairness-related behaviour in humans, specifically inequity-aversion. Their findings can be used to construct utility functions for software agents that closely resemble the inherent human notions of fairness.

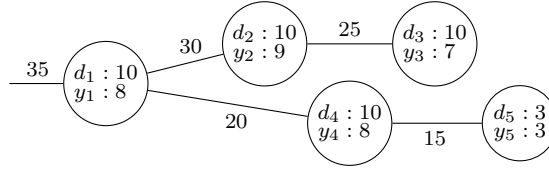


Figure 2.2 The example network with the desire of one agent lowered.

Taking $A = A^+$, the utility function presented in (Fehr and Schmidt, 1999) takes the following form:

$$\begin{aligned}
 u_a = y_a - \frac{\alpha}{m-1} \sum_{b \neq a} \Delta y_{b,a} \\
 - \frac{\beta}{m-1} \sum_{b \neq a} \Delta y_{a,b},
 \end{aligned}
 \tag{2.6}$$

where $\Delta y_{a,b} = \max(y_a - y_b, 0)$ is the positive difference between y_a and y_b . The restriction $A = A^+$, i.e. that all agents are consumers, will later be extended to the case that includes both consumers and producers. The utility function is similar when instead taking $A = A^-$.

The utility function (2.6) takes into account comparative equity; it adds two terms that compare the activity of the agent with the activity of all other agents in the network. The first term measures the utility loss from envy, i.e. others consuming more, while the second term measures the utility loss from pity, i.e. others consuming less. The parameters α and β represent the levels of envy and pity, leading to the reasonable assumptions that

$$0 \leq \alpha, \quad 0 \leq \beta \leq 1, \quad \beta \leq \alpha.
 \tag{2.7}$$

Since the comparative equity utility (2.6) explicitly considers the relation of agents to each other, it can simply be used with the SW optimization (2.1) to find a fair allocation. There are, however, a number of aspects unique to the congestion problem setting that demand adjustments to the comparative equity utility function as presented in (2.6).

Figure 2.2 presents a slightly modified version of the example network. The presented allocation includes inequalities among agents that are not clearly detrimental to its fairness. Agent 5 is allocated a significantly lower activity, but its activity y_5 equals its desire d_5 . This means that agent 5 is perfectly content, and is thus unlikely to envy other agents. Likewise, if the other agents have knowledge of agent 5's desire, it is unlikely that they will pity agent 5.

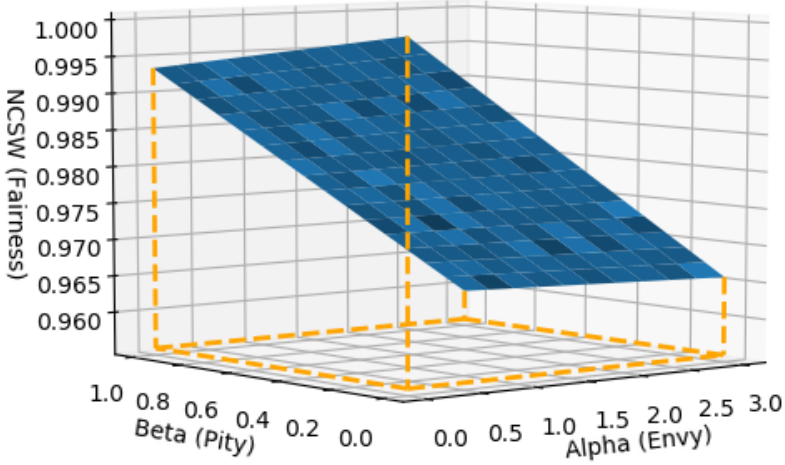


Figure 2.3 NCSW corresponding to the allocation presented in Figure 2.2.

These situations can be taken into account by adding a factor that signifies how discontent an agent a is with their allocated activity y_a relative to their desire d_a . Since it is unlikely that an agent a will pity another agent for not being allocated an activity that agent a wanted but was not allocated itself, the discontent factor applied to the pity term should take this matter of perspective into account.

The result is comparative discontent equity (*CDE*) utility:

$$\begin{aligned}
 u_a = y_a & - \frac{\alpha}{m-1} \sum_{b \neq a} \frac{d_a - y_a}{d_a} \cdot \Delta y_{b,a} \\
 & - \frac{\beta}{m-1} \sum_{b \neq a} \frac{\min(d_b, y_a) - y_s}{\min(d_b, y_a)} \cdot \Delta y_{a,b}.
 \end{aligned} \tag{2.8}$$

If, in the situation under consideration, agents do not have (full) information about other agents' desires, then the pity term should be dropped altogether.

In order to normalize the *SW* when *CDE* utility is used, instead of taking the *ASW* as suggested in Section 2.3, the *SW* is divided by the sum of activities. This results in normalized comparative social welfare (*NCSW*):

$$\frac{\sum_{a \in A} u_a}{\sum_{a \in A} y_a} \tag{2.9}$$

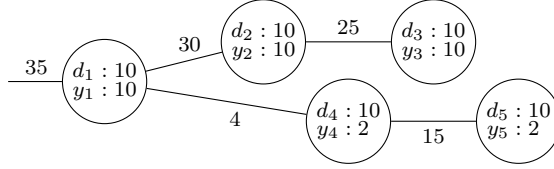


Figure 2.4 The example network with only one side congested.

Figure 2.3 depicts the *NCSW* corresponding to the allocation presented in Figure 2.2 for different values of α and β , showing how agents' characteristics determine perceived fairness. Note that when envy and pity do not play a role, i.e. both α and β are zero, the *NCSW* takes its maximum value of 1.

NCSW is a suitable fairness measurement; normalized, independent of other qualities, and customizable through the parameters α and β . However, it applies only to groups of exclusively consumers (or producers) and does not take network topology into account.

When agents in the network both consume and produce, an adjustment of *NCSW* is required. Since consumers do not compete for network capacity with producers and vice versa, neither envy nor pity between the two is expected. Thus, for each consumer, the envy and pity terms in *CDE* utility should sum over only all other consumers. Similarly, the producers only compare themselves to all other producers. For a consumer, this exclusive comparative equity (*ECDE*) utility takes the form

$$\begin{aligned}
 u_a = y_a - \frac{\alpha}{|A^+| - 1} \sum_{b \in A^+ \setminus \{a\}} \frac{d_a - y_a}{d_a} \cdot \Delta y_{b,a} \\
 - \frac{\beta}{|A^+| - 1} \sum_{b \in A^+ \setminus \{a\}} \frac{\min(d_b, y_a) - y_s}{\min(d_b, y_a)} \cdot \Delta y_{a,b}.
 \end{aligned} \tag{2.10}$$

Besides the mode of prosumption, the network topology may also play a role in determining the set of agents that any agent compares itself to. Figure 2.4 presents a version of the example network that is only congested on one side. This is an interesting situation: although all agents have the same desire and agents 2 and 3 on the non-congested side have been allocated more activity, a reduction of their activity cannot improve the situation for agents 4 and 5 on the congested side. In principle, agents 4 and 5 would not envy agents 2 and 3. This raises the question of the topological reach of comparative equity.

A possible approach to capturing this topological separation in the utility is to define regions in the network with subsets of agents associated to them. Agents from a certain subset could then have different α values depending on

whether comparing to an agent from their own subset, or an agent from another subset. For example, for the network depicted in Figure 2.4, the utility function for agent 4 could be

$$\begin{aligned}
 u_4 = & y_4 - \alpha \cdot \frac{d_4 - y_4}{d_4} \cdot \Delta y_{5,4} \\
 & - \beta \cdot \frac{\min(d_5, y_4) - y_5}{\min(d_5, y_4)} \cdot \Delta y_{4,5} \\
 & - \frac{\alpha'}{3} \sum_{b \in \{1,2,3\}} \frac{d_4 - y_4}{d_4} \cdot \Delta y_{b,4} \\
 & - \frac{\beta'}{3} \sum_{b \in \{1,2,3\}} \frac{\min(d_s, y_4) - y_s}{\min(d_s, y_4)} \cdot \Delta y_{4,b},
 \end{aligned} \tag{2.11}$$

with $\alpha > \alpha'$ and $\beta > \beta'$. This desire of specific comparative discontent equity (*SCDE*) utility function that allows different α and β values when comparing to different groups is adaptable to the distribution of population and capacity constraints of the network under consideration. For instance, to distinguish groups of agents by region or desire (e.g. hospitals).

In summary, *ECDE* utility makes a fair optimization goal with *SW* and a fairness measure with *NCSW* that both mimic notions of fairness inherent to humans. When required, *SCDE* utility may be used as a flexible way to accommodate network constraints and other attributes.

2.5 Congestion Management Mechanism

In this section, a congestion management mechanism is proposed in algorithmic form. Consider an acyclic network where all agents are consumers with simple utility (2.3). The mechanism allocates activities y_a to agents $a \in A$ based on the network topology and the agents' desires d_a . The resulting allocation Y is in the solution set S : it resolves congestion and allocated activities are bounded by agent desires. Most real-world local energy networks can be represented by such an acyclic model.

The acyclic network is interpreted as a rooted tree with its root connected to an external network. Let T_v denote the subtree of vertex v and denote the line from vertex v directed towards the root as line v . Thus, any vertex $v \in V$ has a capacity C_v associated with it that limits the flow from the subtree T_v towards the root of the network (or, in case of the root vertex, towards the external network).

The allocation mechanism, presented in Algorithm 1, takes an egalitarian approach: when congestion occurs at a line v , all consumers in the subtree T_v have the upper bound for their activities reduced to the same level. As a consequence, consumers with the lowest activity only have their activity reduced when all other consumers in the subtree T_v have their activity reduced to that same level.

Algorithm 1: Allocation Mechanism

```

1 Initialize  $y_a = d_a$  for all agents  $a \in A$ 
2 while not all vertices are marked do
3     Select unmarked vertex  $v$  with no unmarked children and mark it
4     if total consumption of the subtree  $T_v$  exceeds the capacity  $C_v$  then
5         Select value  $w$  such that  $\sum_{a \in T_v} \min(w, y_a) = C_v$ 
6         Set  $y_a = \min(w, y_a)$  for all  $a$  in the subtree  $T_v$ 
7 return  $Y = \{y_a\}_{a \in A}$ 

```

Proposition 2.1. The set of activities y_a allocated by the allocation mechanism to the agents $a \in A$ given their desires d_a maximizes the social welfare (SW) on the solution set S .

Proof. Consider an agent $a \in A$ and the final value y_a . If $y_a = d_a$, then the utility of agent a is maximal within the solution set S and cannot be changed to improve the SW.

If $y_a \neq d_a$, then, since all agents in the network are consumers, $y_a < d_a$. Let v be the last vertex where y_a was reduced. This means that the agent a is located in the subtree T_v for which, after executing lines 4 – 6, it holds that

$$\sum_{b \in T_v} y_b = C_v. \quad (2.12)$$

Since the activity y_a has not been reduced since vertex v and it was maximal among activities of agents in T_v , it follows that none of the activities of agents in T_v have changed since vertex v .

Now consider a nonempty set $A^I \subset A$ with $y_a < d_a \forall a \in A^I$ and a set $I = \{\epsilon_a \mid a \in A^I\}$ of corresponding activity increases with $0 < \epsilon_a \leq d_a - y_a \forall a \in A^I$. Let L denote the set of vertices where at least one of the agents $a \in A^I$ had their activity y_a last reduced. Since equation (2.12) holds for all vertices $v \in L$, the activity increases I cause congestion at all those vertices.

Let $A^D = A \setminus A^I$ and let $D = \{-\delta_b \mid b \in A^D\}$ be a set of corresponding activity decreases with $0 \leq \delta_b \leq y_b \forall b \in A^D$. For each vertex $v \in L$, to resolve congestion caused by I , it must hold that $\sum_{b \in A_v^D} \delta_b \geq \sum_{a \in A_v^I} \epsilon_a$, where A_v^I and A_v^D are the subsets of A^I and A^D in T_v . If such D does not exist, then the set of activities increased by I is not in the solution set S .

Since A^I is the disjoint union $\bigsqcup_{v \in K} A_v^I$ for some $K \subset L$, it follows that $\sum_{a \in A^I} \epsilon_a - \sum_{b \in A^D} \delta_b \leq 0$. Therefore, the *SW* cannot be improved by changing any number of activities. \square

As demonstrated in the proof of Proposition 2.1, for no single agent $a \in A$ can the utility y_a be improved within the solution set S . This entails the following corollary.

Corollary 2.1. The set of activities y_a allocated by the allocation mechanism to the agents $a \in A$ given their desires d_a is pareto efficient on the solution set S .

Proposition 2.2. The set of activities y_a allocated by the allocation mechanism to the agents $a \in A$ given their desires d_a maximizes the Nash product (*NP*) on the solution set S .

Proof. Using the notation and setting from the proof of Proposition 2.1, for each $v \in L$, to solve the congestion caused by I at v , it must hold that $\sum_{a \in A_v^I} \epsilon_a - \sum_{b \in A_v^D} \delta_b \leq 0$.

First, note that maximizing the *NP* is equivalent to maximizing $\log(NP) = \sum_{a \in A} \log(y_a)$. Then, consider the same expression including the increases I and decreases D , i.e.:

$$\sum_{a \in A^I} \log(y_a + \epsilon_a) + \sum_{b \in A^D} \log(y_b - \delta_b). \quad (2.13)$$

Since the derivative of $\log(x)$ is $\frac{1}{x}$, and $\log(x)$ is a strictly concave function, it holds that

$$\log(y_a + \epsilon_a) < \log(y_a) + \frac{\epsilon_a}{y_a} \quad \forall a \in A^I \quad (2.14)$$

$$\log(y_b - \delta_b) \leq \log(y_b) - \frac{\delta_b}{y_b} \quad \forall b \in A^D. \quad (2.15)$$

It follows that expression (2.13) is strictly smaller than

$$\log(NP) + \sum_{a \in A^I} \frac{\epsilon_a}{y_a} - \sum_{b \in A^D} \frac{\delta_b}{y_b}. \quad (2.16)$$

Now for all $v \in L$, let w_v denote the value selected at line 5 so that, by line 6, $y_a = w_v$ for each $a \in A_v^I$ with activity y_a last reduced at v , and $y_b \leq w_v$ for each $b \in A_v^D$. Consider the vertex v with the longest root path in L . Then the right two terms in expression (2.16) together are smaller than

$$\begin{aligned} & \sum_{a \in A^I \setminus A_v^I} \frac{\epsilon_a}{y_a} - \sum_{b \in A^D \setminus A_v^D} \frac{\delta_b}{y_b} + \sum_{a \in A_v^I} \frac{\epsilon_a}{w_v} - \sum_{b \in A_v^D} \frac{\delta_b}{w_v} \\ &= \sum_{a \in A^I \setminus A_v^I} \frac{\epsilon_a}{y_a} - \sum_{b \in A^D \setminus A_v^D} \frac{\delta_b}{y_b} + \frac{1}{w_v} \left(\sum_{a \in A_v^I} \epsilon_a - \sum_{b \in A_v^D} \delta_b \right). \end{aligned} \quad (2.17)$$

Now consider the vertex z with the longest root path in L for which v is in T_z . Since the agents $a \in A_v^I$ had their activities last reduced at vertex v , it must be that $w_z \geq w_v$. Thus, since the expression between brackets in equation (2.17) is known to be negative, it follows that expression (2.17) is smaller than

$$\sum_{a \in A^I \setminus A_v^I} \frac{\epsilon_a}{y_a} - \sum_{b \in A^D \setminus A_v^D} \frac{\delta_b}{y_b} + \frac{1}{w_z} \left(\sum_{a \in A_v^I} \epsilon_a - \sum_{b \in A_v^D} \delta_b \right) \quad (2.18)$$

which, by repeating the argument, is then smaller than

$$\sum_{a \in A^I \setminus A_z^I} \frac{\epsilon_a}{y_a} - \sum_{b \in A^D \setminus A_z^D} \frac{\delta_b}{y_b} + \frac{1}{w_z} \left(\sum_{a \in A_z^I} \epsilon_a - \sum_{b \in A_z^D} \delta_b \right).$$

These two arguments can be repeated for all vertices in L , ultimately showing that the right two terms in (2.16) together are negative. From this it follows that the logarithm including the increases I and necessary decreases D , as shown in expression (2.13), is strictly smaller than $\log(NP)$, completing the proof. \square

Important to any mechanism incorporating fairness is that the mechanism is truthful. This means that for the agents, reporting their true desire is a weakly dominant strategy; i.e. agents cannot benefit from strategizing and misreporting.

Proposition 2.3. The allocation mechanism is truthful.

Proof. Consider an agent $a \in A$ and their true desired activity d_a^* . If reporting $d_a = d_a^*$ yields a final activity $y_a < d_a^*$, then there is a last vertex v where y_a was reduced to resolve congestion. Therefore, reporting any $d_a > y_a$ would also cause congestion at vertex v and result in the same final activity y_a . Moreover, reporting any $d_a \leq y_a$ would result in a final activity d_a since y_a had already been sufficiently reduced to resolve any congestion.

Therefore, reporting $d_a = d_a^*$ is a weakly dominant strategy for maximizing y_a within S . This proves the proposition. \square

Propositions 2.1 to 2.3 provide a strong result concerning the allocation mechanism: given the specific problem setting, it provides a truthful and pareto efficient congestion management solution that optimizes egalitarian fairness within the constraints of the network topology.

The worst case computational complexity occurs when the mechanism must determine a value w at line 5 by sorting all m agents in $\mathcal{O}(m \cdot \log(m))$ time, and must do this at each of the n vertices. Hence, the worst case computational complexity is $\mathcal{O}(n \cdot m \cdot \log(m))$, where n and m are the total number of vertices and agents in the network, respectively.

2.6 Conclusions

This chapter proposed both the normalized Nash product and comparative discontent equity utilities combined with social welfare as fair optimization goals and normalized fairness measuring tools. Furthermore, this chapter presented a congestion management solution in the form of an egalitarian allocation mechanism. Finally, the allocation mechanism was proven to be truthful and maximize both social welfare and the Nash product.

Future work could provide a congestion management solution based on the human-inspired concepts of fairness presented in Section 2.4, or extend the allocation mechanism presented in Section 2.5 to more general settings.

3

Principal Notions of Fair Division and Local, Outer Matching

In the previous chapter we considered fairness expressed by comparative utility functions. In this chapter we make local fairness explicit by proposing a notion of fairness that promotes local supply-demand matching in allocations.

3.1 Introduction

With the increasing share of renewable energy sources comes an increase in electrical grid congestion (Verzijlbergh et al., 2014). The consequences of this rapid increase in congestion are already seen in soaring congestion management costs (Nabe et al., 2017), even at the low- and medium-voltage levels. Current grid congestion management is insufficiently prepared for the changes in electrical grid operation brought about by the energy transition. New solutions for congestion management come in various forms (Haque et al., 2017; Philipsen et al., 2016; Xue et al., 2009), often focusing on the introduction of electric vehicles (EVs) (Hu et al., 2015; Li et al., 2014; Rivera et al., 2015; Shao et al., 2017) or the decentralized, multi-agent aspect of distributed energy resources (DERs) (Ciavarella et al., 2019; Vandael et al., 2011; Vytelingum et al., 2010).

However, these solutions typically do not explicitly take into account a concept of fairness. Meanwhile, the energy transition is bringing about a paradigm shift from ‘supply follows demand’ to ‘demand follows supply’ that, as a consequence, places the responsibility for congestion increasingly on individual prosumers. This development makes it crucial to explicitly incorporate concepts of fairness in congestion management solutions (Hekkelman and La Poutré, 2019). The European Commission emphasizes the importance of fairness in energy, stating that “energy is a critical good, absolutely essential for full participation in modern society. The clean energy transition also needs

to be fair for those sectors, regions or vulnerable parts of society affected by the energy transition.” (European Parliament, 2016)

In this chapter we lay a theoretical foundation for fair congestion management, using a congestion model similar to those used in (Bach Andersen et al., 2012; Hu et al., 2014; Rivera et al., 2015; Verzijlbergh et al., 2014). We propose an algorithmic mechanism of low computational complexity that combines a locally oriented novel fairness concept with principal notions of fair division. We prove that this algorithmic mechanism divides the available network capacity maximally over the prosumers.

Specifically, in this chapter we propose local, outer matching as a novel concept of fairness for congestion management in low-voltage networks. This concept of fairness requires that congestion is resolved with recursive matching of supply and demand in localities outward from nodes in the network. Local, outer matching thus prioritizes matching in the peripheral of the network, reducing strain and losses on the network infrastructure.

Still, when congestion occurs, the available network capacity must be fairly divided over the affected prosumers. The fair division of goods and fairness in general are established and active fields of research in mathematics and economics (Alon, 1987; Aziz and Mackenzie, 2016; Brams and Taylor, 1995; Dubins and Spanier, 1961; Pazner and Schmeidler, 1978; Stromquist, 1980). One application domain is that of communications networks, where network capacity must be fairly divided over users (Gutman and Nisan, 2012). While congestion management in power flow networks faces similar fair division problems, power flow networks are concerned with a single-commodity flow as opposed to peer-to-peer data transmission. In the energy domain currently, fairness is considered mostly for EV charging (Danner et al., 2019) and DER related pricing (Khodabakhsh et al., 2019).

In this chapter we discuss three principal notions of fair division that may be combined with the novel fairness concept of local, outer matching to perform its division. The principal notions of fair division we discuss here are: proportional, first proposed by Steinhaus (Steinhaus, 1948) and sometimes referred to as ‘simple fair division’; egalitarian, which is closely related to the concept of envy-freeness first proposed by Gamow and Stern (Gamow and Stern, 1958); and nondiscriminatory, which is a natural counterpart to the egalitarian notion of fair division.

It is apparent that fairness in power flow networks has a combination of aspects. On the one hand, fair division of network capacity is required. On the other hand, the single-commodity flow necessitates supply-demand matching throughout the network. Matching supply and demand locally is a newly accepted paradigm for energy networks that stimulates the use of local infrastructure. This introduction of localities affects concepts of fairness (Abebe et al.,

2017). Envisioned autarkic-like local communities and neighbourhoods thus demand a local approach to congestion management and concepts of fairness.

To this end, we devise an algorithmic mechanism that computes the combination of local, outer matching with the discussed principal notions of fair division, resulting in locally oriented congestion solutions that make maximal use of the network capacity. We then prove that the egalitarian notion of fairness results in an incentive compatible mechanism, while the proportional and nondiscriminatory notions of fairness do not result in an incentive compatible mechanism. Finally, we show that the proposed algorithmic mechanism computes congestion solutions in limited computational time, which is essential for application in the energy domain.

The contributions of this chapter to the state of the art can be summarised as follows:

- We propose a novel concept of fairness for congestion management called local, outer matching.
- We discuss the principal notions of proportional, egalitarian, and nondiscriminatory fair division that we combine with the concept of local, outer matching.
- We devise an algorithmic mechanism that combines local, outer matching with notions of fair division to compute maximal congestion solutions in limited computational time.
- We prove that this mechanism is incentive compatible when using the egalitarian notion of fair division, and is not incentive compatible when using the proportional or nondiscriminatory notions of fair division.

The chapter is organized as follows. First, Sections 3.2 and 3.3 introduce the setting, model, and useful concepts. Section 3.4 then formally defines division and discusses the three principal notions of fair division. Section 3.5 defines the novel fairness concept of local, outer matching and provides congestion solutions that are proven to be local, outer matchings that make maximal use of network capacity. Finally, Section 3.6 presents the algorithmic mechanism that combines local, outer matching with the principal notions of fair division. The incentive compatibility results follow in Section 3.7 and the computational complexity results in Section 3.8. Section 3.9 concludes the chapter.

3.2 Setting

We consider an electrical power flow network that consists of prosumers connected to each other by electrical grid lines. These grid lines have a maximum reliable capacity (which is usually less than their physical or thermal limit).

Grid congestion occurs when electrical power flow caused by the prosumers exceeds some line capacities. This means that it is not always possible to realise the desired prosumption of all prosumers within the network constraints. Congestion management is the practice of reducing, resolving, or preventing grid congestion by deviating from the desired prosumptions in order to accommodate the network constraints.

We follow a modelling approach similar to those in (Bach Andersen et al., 2012; Hu et al., 2014; Rivera et al., 2015; Verzijlbergh et al., 2014) that lay the theoretical foundations for congestion management, e.g. by focusing on active power curtailment. As such, we model an electrical power flow network as a tree (representing almost all real-world low-voltage networks), the line capacities as edge weights, the prosumers as agents that are either consumers or producers located at the vertices, and the desired prosumptions as agent desires. Low-voltage networks are usually connected to larger electrical grids through a substation. This connection to an external grid may be modelled as a virtual edge, which will also have a line capacity modelled as an edge weight.

3.3 Model

Let $T = (V, E)$ be a rooted weighted tree. Let a virtual edge at the root r represent the connection to a virtual parent that represents an external network. Let the edge weights be positive, representing flow capacities. Denote the weight of an edge between vertex $v \in V$ and its parent by C_v . In addition, consider a set of agents A distributed over the vertices V . Finally, for each agent $a \in A$, consider its desire d_a . A positive desire indicates a consumer while a negative desire indicates a producer. Let $A^+, A^- \subset A$ be the sets of consumers and producers, respectively.

Definition 3.1 (Congestion Tree). Define a *congestion tree* $T = (V, E, A)$ as a tree $T = (V, E)$ with root r , edge weights C_v for $v \in V$, and agents $a \in A$ with desires d_a located at the vertices $v \in V$.

The subtree of a congestion tree $T = (V, E, A)$ is again a congestion tree, and is denoted by $T_v = (V_v, E_v, A_v)$ where v is its root. A subtree $T_v = (V_v, E_v, A_v)$ inherits the edge weights of $T = (V, E, A)$, with its virtual edge inheriting the weight of the edge between v and its parent in $T = (V, E, A)$. See Figure 3.1 for a representation.

3.3.1 Congestion Management

Congestion management is the practice of reducing, resolving, or preventing congestion in a network, and can take various forms. In graph theory, flow networks consider flow resulting from a single source and a single sink in a graph with flow capacities on the edges. Power networks usually deal with a

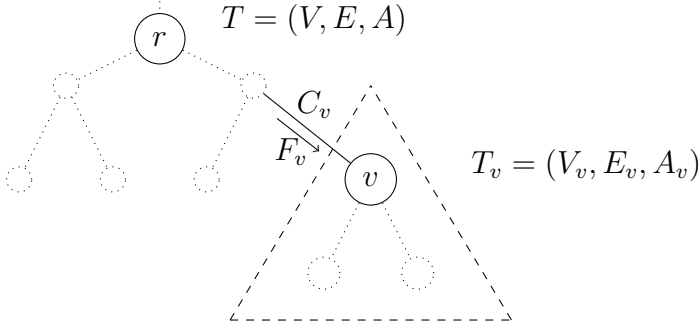


Figure 3.1 A representation of a congestion tree $T = (V, E, A)$, highlighting the situation around a vertex v .

more complex situation where many distributed prosumers are participating in a market.

In the context of congestion trees, congestion management is performed by allocating network access to agents based on their desires and the network constraints. Such allocations could be strictly enforced or used as reference for penalties or incentives.

Definition 3.2. An *allocation* Y on a congestion tree $T = (V, E, A)$ is a map $Y : A \rightarrow \mathbb{R}$.

Notation 3.1. For $B \subseteq A$, abbreviate $\sum_{a \in B} Y(a)$ as $Y(B)$.

When agents that represent prosumers are subject to such an allocation, the resulting prosumptions lead to (electrical power) flows in the network.

Definition 3.3 (Incoming and Locally Balanced Flows). Given an allocation Y on a congestion tree $T = (V, E, A)$, define the *incoming flow* $F_v(Y)$ and the *locally balanced flow* $LBF_v(Y)$ of a subtree $T_v = (V_v, E_v, A_v)$ as:

$$F_v(Y) = Y(A_v) = \sum_{a \in A_v} Y(a), \quad (3.1)$$

$$LBF_v(Y) = \frac{1}{2} \left[\sum_{a \in A_v} |Y(a)| - \left| \sum_{a \in A_v} Y(a) \right| \right]. \quad (3.2)$$

The incoming flow $F_r(Y)$ of an allocation Y on a congestion tree $T = (V, E, A)$ with root r thus represents the total amount of electrical power demanded of, or supplied to, the external grid. The incoming flow also gives

the flow over each edge, with $F_v(Y)$ being the flow to vertex v from its parent. See Figure 3.1 for a representation.

The locally balanced flow $LBF_r(Y)$ represents the total amount of electrical power that flows between the agents within the congestion tree $T = (V, E, A)$ with root r , including between agents that share a vertex. The locally balanced flow $LBF_r(Y)$ is a measure for the matching of consumer and producer desires within the congestion tree $T = (V, E, A)$.

Definition 3.4 (Desire Compatible). An allocation Y on a congestion tree $T = (V, E, A)$ is *desire compatible* if

$$0 \leq Y(a) \leq d_a \quad \text{or} \quad 0 \geq Y(a) \geq d_a \quad \forall a \in A. \quad (3.3)$$

Definition 3.5 (Congestion Free). An allocation Y on a congestion tree $T = (V, E, A)$ is *congestion free* if, for each vertex $v \in V$, the incoming flow $F_v(Y)$ of the subtree $T_v = (V_v, E_v, A_v)$ does not exceed the flow capacity C_v of its virtual edge:

$$|F_v(Y)| \leq C_v \quad \forall v \in V. \quad (3.4)$$

Definition 3.6 (Feasible). An allocation Y on a congestion tree is *feasible* if it is both desire compatible and congestion free.

The set of feasible allocations forms the solution space for the problem of congestion management. Within this solution space, some allocations are more desirable than others because they make better use of the available network capacity.

Definition 3.7 (Base Allocation). A feasible allocation Y on a congestion tree $T = (V, E, A)$ with root r is a *base allocation* if it maximizes the locally balanced flow $LBF_r(Y)$ and has incoming flow $F_r(Y) = 0$.

A base allocation maximally matches consumer and producer desires in a congestion tree, without interacting with the external grid it is connected to.

Definition 3.8 (Max Allocation). A feasible allocation Y on a congestion tree $T = (V, E, A)$ with root r is a *max allocation* if it maximizes $|F_r(Y)|$ under the condition that the locally balanced flow $LBF_r(Y)$ is maximal.

A max allocation maximizes the use of the available network capacity by making maximal use of the connection to the external grid after maximally matching consumer and producer desires internally. A max allocation can be viewed as a base allocation plus an allocation of the remaining unmatched desires.

3.4 Principal Notions of Fair Division

Congestion management leads directly to the question of fair division. The limited available network capacity does not belong to any one agent, and thus it must be fairly divided. The field of fair division of goods considers such division problems. In this case a number of agents lay claim to a portion of a divisible good, the network capacity, but the sum of their claims exceeds the availability of the good. Since energy is a critical good, it is of great importance that the division of capacity, which dictates network access, be fair to all agents. However, the notion of which choice of division constitutes a fair division is subjective.

Consider two agents a and b that have respective claims d_a and d_b to a quantity k of a divisible good. This situation may be visualized as in Figure 3.2. The point of the claims (d_a, d_b) represents the outcome desired by the agents, while the line intersecting the axes represents the available quantity k of the good.

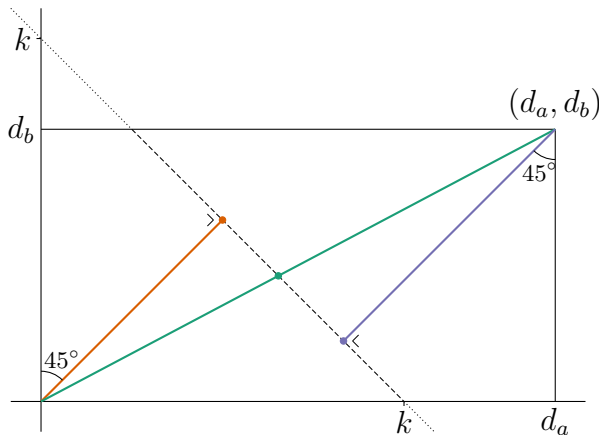


Figure 3.2 Three principal notions of fair division illustrated: proportional (green, middle), egalitarian (red, left), and nondiscriminatory (blue, right).

The set of divisions is the set of points that exactly divide the quantity k and do not allocate to agents more than their claims. In Figure 3.2 this is the set of points on the dashed line segment inside the rectangle. If this set is empty because the quantity k is larger than the sum of the claims, then there is no real division problem and the point of claims (d_a, d_b) is taken as the division.

This representation of the division of a good over claims can be extended to any number m of agents by an m -dimensional space with an $m-1$ dimensional surface representing the quantity k . This leads to the following definition.

Definition 3.9 (Division). The division $Div(k, Y, A)$ of a positive value k over an allocation Y (of claims) with positive values on a set of agents A , is an allocation D on A with

$$D(A) = \min(k, Y(A)) \quad (3.5)$$

and

$$D(a) \in [0, Y(a)] \quad \forall a \in A. \quad (3.6)$$

If $k \geq Y(A)$, then D is simply identical to the allocation Y on A .

A notion of fair division is the choice of a specific division from the set of divisions. Visually, this means that a notion of fair division is the choice of a specific point on the $m - 1$ dimensional plane that represents the quantity k , inside the m -dimensional hyperrectangle drawn by the origin and the claims. Figure 3.2 shows three principal notions of fair division:

- The **proportional** notion of fairness finds the point where the line from the origin to the point of claims intersects the set of divisions. See the green line in Figure 3.2. Each agent is allocated a portion of the good that is proportional to the ratio of its claim to the sum of all claims. This division treats all agents equally, preserving the relations between the claims.
- The **egalitarian** notion of fairness finds the point in the set of divisions that is closest to the origin. See the red line in Figure 3.2. Each agent is allocated the same portion, unless its claim is smaller than that portion. This division treats all agents equally, reducing all claims to the same amount.
- The **nondiscriminatory** notion of fairness finds the point in the set of divisions that is closest to the point of claims. See the blue line in Figure 3.2. The portion of each agent is reduced by the same amount regardless of its claim, to a minimum of zero. This division treats all agents equally, reducing all claims by the same amount.

3.5 Local, Outer Matching and Fairness

The matching of consumer and producer desires is an important aspect of congestion management in congestion trees. On the one hand, the distributed nature of prosumers in electrical power flow networks is the cause of much congestion. On the other hand, the presence of both consumers and producers provides opportunity for mitigating or resolving congestion by locally balancing excessive production or consumption.

Local matching of supply and demand is a newly accepted paradigm for future energy networks. It stimulates the use of local infrastructure as a push towards envisioned autarkic-like local communities and neighbourhoods, reducing strain and losses on the energy network in the process. This local matching of supply and demand also, again, emphasizes the importance of fairness. Factors such as the relative locality of the prosumers will play an important role in concepts of fairness regarding the energy domain.

We address this important problem by presenting local, outer matching as an efficient and fair concept for matching consumer and producer desires. A local, outer matching solution prioritizes local matching in the peripheral where prosumers are furthest away from the substation.

Definition 3.10 (Local, Outer Matching). A feasible allocation Y on a congestion tree $T = (V, E, A)$ is a *local, outer matching* if the locally balanced flow $LB F_v(Y)$ is maximal for each subtree $T_v = (V_v, E_v, A_v)$.

To make maximal use of the available network capacity, a local, outer matching solution is sought that is also a max allocation or a base allocation, depending on the envisioned interaction with the external grid. A max allocation makes maximal use of the external grid as well as the network capacity, while a base allocation is self-balanced and makes maximal use of the network

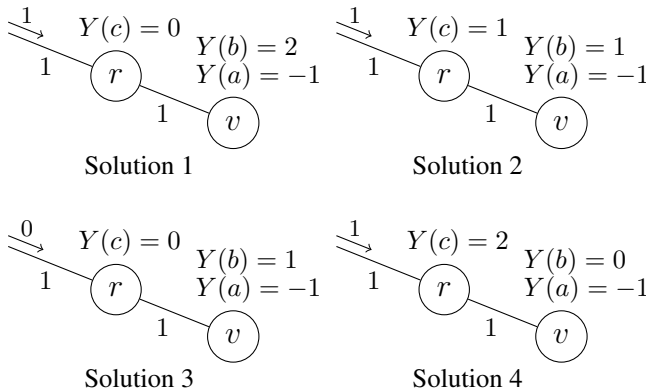


Figure 3.3 Example congestion solutions for a simple congestion tree $T = (V, E, A)$ consisting of a root r with $C_r = 1$ containing agent c with $d_c = 2$, and a second vertex v with $C_v = 1$ containing both agents b and a with $d_b = 2$ and $d_a = -1$. Solutions 1, 2, and 3 are local, outer matchings. Solutions 1, 2, and 4 are max allocations. Solution 3 is a base allocation.

capacity for internally balanced flows. See Figure 3.3 for some examples. The *FairMax* and *FairBase* allocations from Definition 3.11 present such local, outer matching solutions. Their maximal capacity use and local, outer matching are proven by Theorem 3.1.

Let $T = (V, E, A)$ be a congestion tree. Because of the priority given to outer matching from the leaves towards the root, the definitions of local, outer matching allocations are recursive. For a vertex $v \in V$, the consumer and producer desires of all consumers and producers at vertices outward from the vertex v are maximally matched. This results in a chain of matchings where consumer and producer desires are partially already satisfied by local, outer matching and partially remain unsatisfied. The remaining desires also resolve congestion by dividing the available capacity.

Definition 3.11 (*FairMax* and *FairBase* Allocations). Consider a congestion tree $T = (V, E, A)$. Let $c(r)$ be the set of child vertices of the root r . For each child $v \in c(r)$, let $FairMax_v$ and $FairBase_v$ be the *FairMax* and *FairBase* allocations on $T_v = (V_v, E_v, A_v)$. Let the remaining desire Rem for agents $a \in A$ be defined as

$$Rem(a) = \begin{cases} FairMax_v(a) - FairBase_v(a) & \text{if } a \in A_v, v \in c(r) \\ d_a & \text{if } a \in A \text{ is at } r \end{cases} \quad (3.7)$$

In the case that $Rem(A^+) \geq |Rem(A^-)|$, define the *FairMax* allocation on the congestion tree $T = (V, E, A)$ as follows:

For a producer $a \in A_v^-$ with $v \in c(r)$,

$$FairMax(a) = Rem(a) + FairBase_v(a) \quad (= FairMax_v(a)), \quad (3.8)$$

for a producer $a \in A^-$ at r ,

$$FairMax(a) = Rem(a) \quad (= d_a), \quad (3.9)$$

for a consumer $a \in A_v^+$ with $v \in c(r)$,

$$FairMax(a) = Div(Allowance, Rem, A^+)(a) + FairBase_v(a), \quad (3.10)$$

for a consumer $a \in A^+$ at r ,

$$FairMax(a) = Div(Allowance, Rem, A^+)(a), \quad (3.11)$$

where $Allowance = |Rem(A^-)| + C_r$.

The definition of the *FairBase* allocation on the congestion tree $T = (V, E, A)$ is similar, taking instead $Allowance = |Rem(A^-)|$.

In the case that $Rem(A^+) \leq |Rem(A^-)|$ the definitions are analogous, only switching consumers and producers.

Informally, the definition of the *FairMax* allocation is as follows. The remaining desire Rem in Equation (3.7) represents the fact that some parts of the desires of consumers and producers are already matched with each other in the subtrees of the congestion tree $T = (V, E, A)$. See also Figure 3.4 for a visual aide. If the root r has no children, i.e. it is a leaf, then Equations (3.8) and (3.10) do not apply; this is the base case of the recursion. When the consumers outweigh the producers, the producer desires can be fully matched with consumer desires as seen in Equations (3.8) and (3.9). The consumers, however, have to divide the sum of producer desires: the *Allowance*. The *FairMax* allocation also includes interaction with the external network up to the capacity C_r in the *Allowance*, on top of matching with the producer desires, as seen in Equations (3.10) and (3.11). Finally, of course, all agents in subtrees also get the desires that were already satisfied in those subtrees as seen in Equations (3.8) and (3.10).

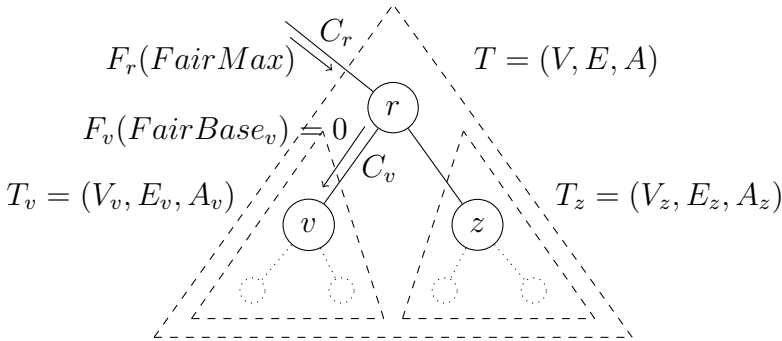


Figure 3.4 A representation of the recursive definition.

Proposition 3.1. The *FairMax* and *FairBase* allocations are, respectively, a max allocation and a base allocation.

Proof. The proof is by induction. Let $c(r)$, $FairMax_v$, $FairBase_v$, Rem , and *Allowance* be as in Definition 3.11.

For the induction basis, let $T = (V, E, A)$ be a congestion tree with only one vertex. Without loss of generality, assume that $Rem(A^+) \geq |Rem(A^-)|$. Note that $Rem(a) = d_a$ for all agents $a \in A$. Since $D(a) \in [0, Rem(a)]$ by Definition 3.9, it follows from Equations (3.9) and (3.11) and their analogs for *FairBase* that both *FairMax* and *FairBase* are desire compatible allocations. Moreover, it follows from Equation (3.11) and its analog for

FairBase, respectively, that

$$FairMax(S) \leq -FairMax(A^-) + C_r \Rightarrow FairMax(A) \leq C_r \quad (3.12)$$

$$FairBase(S) = -FairBase(A^-) \Rightarrow FairBase(A) = 0 \quad (3.13)$$

and thus that both *FairMax* and *FairBase* are congestion free allocations. Consequently, both *FairMax* and *FairBase* are feasible allocations.

Since by Equation (3.9) and its analog for *FairBase*,

$$FairMax(A^-) = FairBase(A^-) = \sum_{a \in A^-} d_a, \quad (3.14)$$

it follows from Equation (3.13) that *FairBase* is a base allocation, and from the properties of the division in Equation (3.11) that *FairMax* is a max allocation.

For the induction step, let $T = (V, E, A)$ be a more general congestion tree, and assume that the proposition holds for subtrees $T_v = (V_v, E_v, A_v)$ with v children of r . The proof of the induction step is similar to that of the induction basis. Again, without loss of generality, assume that $Rem(A^+) \geq |Rem(A^-)|$. Since *FairMax_v* is a desire compatible allocation for all children v of r , it follows in the same way as before from Definition 3.9, Equations (3.8) to (3.11) and their analogs for *FairBase* that both *FairMax* and *FairBase* are desire compatible allocations. Moreover, since $F_v(FairBase_v) = 0$ for all children v of r , it follows in the same way as before from Equations (3.10) and (3.11) and their analogs for *FairBase*, respectively, that

$$FairMax(S) \leq -FairMax(A^-) + C_r \Rightarrow FairMax(A) \leq C_r \quad (3.15)$$

$$FairBase(S) = -FairBase(A^-) \Rightarrow FairBase(A) = 0 \quad (3.16)$$

and thus, since *FairMax_v* and *FairBase_v* are congestion free allocations for all children v of r , that both *FairMax* and *FairBase* are congestion free allocations. Consequently, both *FairMax* and *FairBase* are again feasible allocations.

Since by Equations (3.8) and (3.9) and their analogs for *FairBase*,

$$\begin{aligned} FairMax(A^-) &= FairBase(A^-) \\ &= \sum_{v \text{ child of } r} \sum_{a \in A_v^-} FairMax_v(a) + \sum_{a \in A^- \text{ at } r} d_a, \end{aligned} \quad (3.17)$$

and *FairMax_v* is a max allocation for all children v of r , it follows from Equation (3.16) that *FairBase* is a base allocation, and from the properties of the division in Equations (3.10) and (3.11) that *FairMax* is a max allocation. \square

Theorem 3.1. *The FairMax and FairBase allocations are, respectively, a max allocation and a base allocation that are local, outer matchings.*

Proof. Let $c(r)$ and $FairBase_v$ be as in Definition 3.11.

The theorem follows from Proposition 3.1 and the fact, as seen from Equations (3.8) and (3.10) and their analogs for *FairBase*, that the restrictions of *FairMax* and *FairBase* to a subtree $T_v = (V_v, E_v, A_v)$ for a child $v \in c(r)$ fully contain the allocation $FairBase_v$. Since $FairBase_v$ is a base allocation and thus maximizes the locally balanced flow on the subtree $T_v = (V_v, E_v, A_v)$, so do *FairMax* and *FairBase*.

Through induction on subtrees it then follows that the *FairMax* and *FairBase* allocations maximize the locally balanced flow on all subtrees of the congestion network $T = (V, E, A)$. \square

3.6 Algorithm: Local, Outer Matching combined with Fair Division

This section presents an algorithmic mechanism that combines local, outer matching with notions of fair division to compute the *FairMax* and *FairBase* allocations on a congestion tree $T = (V, E, A)$. A sketch of the approach is as follows. First compute these allocations on subtrees $T_v = (V_v, E_v, A_v)$ for vertices $v \in V$ from the leaves towards the root r . Each step considers one vertex v and the subtree $T_v = (V_v, E_v, A_v)$. In this way, at each step, the *FairMax* and *FairBase* allocations for all subtrees with roots that are children of the current vertex will have been computed already.

Algorithm 2 uses the divide function from Algorithm 3 to fairly divide certain values over sets of agents. This is used for local, outer matching and for resolving any potential congestion. The divide function depends on the notion of fairness.

When the notion of fairness is egalitarian or nondiscriminatory, the divide function in Algorithm 3 uses the water level function from Algorithm 4. The water level function computes the level w that an allocation on a set of agents must be reduced to, per agent, in order to divide a certain value equally over the set of agents.

3.6.1 Algorithmic Local, Outer Matching

Algorithm 2 visits all vertices in V exactly once, moving from the leaves towards the root r . At each step corresponding to a vertex v , for each agent $a \in A_v$, two variables are subject to change. The first variable, $Fb(a)$, is from

the *FairBase* allocation and can be thought of as the satisfied desire that has already been allocated to agent a when considering only $T_v = (V_v, E_v, A_v)$. The second variable, $Fm(a)$, can be thought of as the remaining desire that can still be allocated to agent a in addition to its already satisfied desire, when considering only $T_v = (V_v, E_v, A_v)$.

Algorithm 2: Mechanism

Input: A congestion tree $T = (V, E, A)$ and a fairness notion f
Output: The *FairMax* (Fm) and *FairBase* (Fb) allocations

```

1  $Fb(a) \leftarrow 0$   $\forall a \in A$ 
2  $Fm(a) \leftarrow d_a$   $\forall a \in A$ 
3 while not all vertices in  $V$  are marked do
   | Select an unmarked vertex  $v \in V$ 
4   | with no unmarked children and mark
   | it
5   | if  $Fm(A_v^+) \geq |Fm(A_v^-)|$  then
6   |   |  $D \leftarrow \text{Divide}(|Fm(A_v^-)|, Fm, A_v^+, f)$ 
7   |   |  $E \leftarrow \text{Divide}(|Fm(A_v^-)| + C_v, Fm, A_v^+, f)$ 
8   |   |  $Fb(a) \leftarrow Fb(a) + D(a)$   $\forall a \in A_v^+$ 
9   |   |  $Fm(a) \leftarrow E(a) - D(a)$   $\forall a \in A_v^+$ 
10  |   |  $Fb(a) \leftarrow Fb(a) + Fm(a)$   $\forall a \in A_v^-$ 
11  |   |  $Fm(a) \leftarrow 0$   $\forall a \in A_v^-$ 
12  |   | else
13  |   | | Similarly
14  |  $Fm(a) \leftarrow Fm(a) + Fb(a)$   $\forall a \in A$ 
15 return  $Fm, Fb$ 

```

At each step, the remaining desires Fm of the consumers and producers in A_v are maximally matched with each other. This maximal matching is performed, when the consumers outweigh the producers, by dividing the sum of the remaining producer desires over the remaining consumer desires. The matched amounts are moved from the remaining desires Fm to the satisfied desires Fb . This leaves either only consumers or only producers in terms of remaining desires Fm since either all consumers or all producers have their entire remaining desires Fm moved to their satisfied desires Fb .

Simultaneously, at each step, any potential congestion is resolved by dividing the available capacity over the agents. This is done, when the consumers outweigh the producers, by dividing an amount equal to the sum of producer desires plus the capacity C_v over the consumer desires. This amount can be thought of as the *allowance* of the consumers.

As noted before, the satisfied desires Fb will constitute the *FairBase* allocation. However, the remaining desires Fm do not yet constitute the *FairMax* allocation. In order to obtain the *FairMax* allocation, after all vertices have been visited by the algorithm, the remaining desires Fm and satisfied desires Fb are added together.

At this point it has become easy to see that Algorithm 2 indeed computes the *FairMax* and *FairBase* allocations. This result will be formalized in Section 3.6.4.

3.6.2 Division for Different Notions of Fairness

The division function in Algorithm 3 implements the division from Definition 3.9. If the value k to divide is not larger than $Y(A)$, the function returns a division D on A that exactly divides the value k over the agents $a \in A$ while not exceeding the claims $Y(a)$ for agents $a \in A$.

Algorithm 3: Fair Division

```

1 Function Divide( $k, Y, A, f$ )
   Input:  $A$  positive value  $k$ , an allocation  $Y$  with positive values on
           a set of agents  $A$ , and a fairness notion  $f$ 
   Output: A division  $D$  (i.e.  $D(A) = \min(k, Y(A))$ ) and
            $D(a) \in [0, Y(a)]$  for all  $a \in A$ 
2   if  $k \geq Y(A)$  then
3      $D(a) \leftarrow Y(a)$   $\forall a \in A$ 
4   else if  $f = \text{proportional}$  then
5      $D(a) \leftarrow Y(a)/Y(A) \cdot k$   $\forall a \in A$ 
6   else if  $f = \text{egalitarian}$  then
7      $w \leftarrow \text{WaterLevel}(k, Y, A)$ 
8      $D(a) \leftarrow \min(Y(a), w)$   $\forall a \in A$ 
9   else if  $f = \text{nondiscriminatory}$  then
10     $w \leftarrow \text{WaterLevel}(Y(A) - k, Y, A)$ 
11     $D(a) \leftarrow Y(a) - \min(Y(a), w)$   $\forall a \in A$ 
12  return  $D$ 

```

The division function presented in Algorithm 3 supports the three principal notions of fair division discussed in Section 3.4. However, it is of course possible to add any other possible division.

The division for the proportional notion of fairness is computed with a straightforward ratio multiplication, allocating to each agent $a \in A$ a portion of the value k that is proportional to the ratio of the claim $Y(a)$ to the sum of claims $Y(A)$.

To compute the divisions for the egalitarian and nondiscriminatory notions of fairness, the concept of the water level is used by calling the water level function from Algorithm 4. These two divisions are computed by respectively reducing the claims $Y(a)$ to a water level w and by reducing the claims $Y(a)$ by a water level w , to a minimum of zero.

For the proportional notion of fairness, the computed division D trivially satisfies the output conditions. For the egalitarian notion of fairness, the computed division D can directly be seen to satisfy the output conditions by considering the output condition of the water level function from Algorithm 4. To see that the computed division D also satisfies the output conditions for the nondiscriminatory notion of fairness, consider that

$$\sum_{a \in A} \min(Y(a), w) = Y(A) - k \quad (3.18)$$

and thus that

$$D(A) = Y(A) - \sum_{a \in A} \min(Y(a), w) = k. \quad (3.19)$$

Informally, the part $Y(A) - k$ that will not be allocated is divided evenly over the agents A and subtracted from their claims $Y(A)$.

3.6.3 Setting the Water Level

Setting the water level refers to uniformly dividing a good over claims by computing a single value referred to as the water level. This water level value is used by egalitarian and nondiscriminatory notions of fair division.

The water level function in Algorithm 4 takes a value k and an allocation Y on a set of agents A . It then computes the unique level w that the values $Y(a)$ for agents $a \in A$ must be reduced to in order to exactly divide the value k . Simply setting a single value (k divided by the number of agents in A) for all agents $a \in A$ does not reach the intended goal as some agents may have claims lower than that value. If that is the case, the unclaimed difference can be divided over the other agents.

To do this, the water level function in Algorithm 4 starts by sorting the claims $Y(a)$ from lowest to highest. It then checks if it can allocate the lowest claim to all agents $a \in A$. If yes, the lowest claim is removed. It then checks if it can also allocate the next lowest claim to all remaining agents. Once the next lowest claim cannot be allocated to all remaining agents, the rest of the unallocated value k is evenly divided over the remaining agents.

The name of the function comes from this repeated raising of the allocation from claim to claim that resembles the rising of a water level. Like water poured into a series of connected containers, it divides the quantity equally over the recipients. See also Figure 3.5.

Algorithm 4: Setting the Water Level

```

1 Function WaterLevel( $k, Y, A$ )
   Input:  $A$  positive value  $k$  and an allocation  $Y$  with positive values
           on a set of agents  $A$  with  $k \leq Y(A)$ 
   Output: A value  $w$  such that  $\sum_{a \in A} \min(Y(a), w) = k$ 
2    $i \leftarrow 0$ 
3    $\text{list} \leftarrow \text{Sort}(Y(a), a \in A)$ 
4    $\text{total} \leftarrow k$ 
5    $\text{size} \leftarrow |A|$ 
6    $\text{level} \leftarrow 0$ 
7    $\text{rise} \leftarrow \text{list}[i]$ 
8   while  $\text{total} - \text{size} \cdot \text{rise} > 0$  do
9      $\text{total} \leftarrow \text{total} - \text{size} \cdot \text{rise}$ 
10     $\text{size} \leftarrow \text{size} - 1$ 
11     $\text{level} \leftarrow \text{list}[i]$ 
12     $\text{rise} \leftarrow \text{list}[i + 1] - \text{level}$ 
13     $i \leftarrow i + 1$ 
14   $\text{rest} \leftarrow \text{total}/\text{size}$ 
15   $w \leftarrow \text{level} + \text{rest}$ 
16  return  $w$ 

```

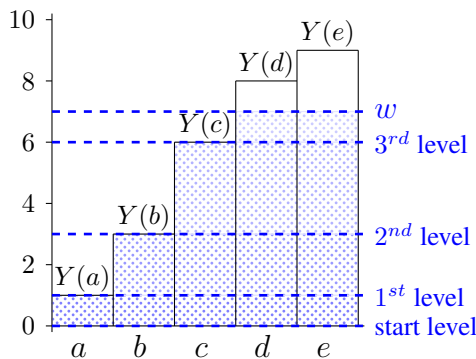


Figure 3.5 Steps of the water level function from Algorithm 4, dividing the value $k = 24$ (hatched surface area) over five agents a through e with respective claims 1, 3, 6, 8, and 9.

3.6.4 Result of the Algorithm

Theorem 3.2. *Algorithm 2 computes the combination of local, outer matching with principal notions of fair division, resulting in the FairMax and FairBase allocations that correspond to maximal or no interaction with the external grid, respectively.*

Proof. At each step of Algorithm 2, the remaining desire Fm is equal to the remaining desire allocation Rem from Definition 3.11. Instead of adding the Fb values and subtracting them again in the next step, they are saved and added to the Fm values only at the end of the algorithm. The variable updates in the algorithm are identical to those in Definition 3.11.

The division function used in Algorithm 2 is provided by Algorithm 3, and allows for combination with any of the three discussed principal notions of fair division. \square

3.6.5 Example Congestion Solutions

Figure 3.6 revisits the simple example congestion tree from Figure 3.3.

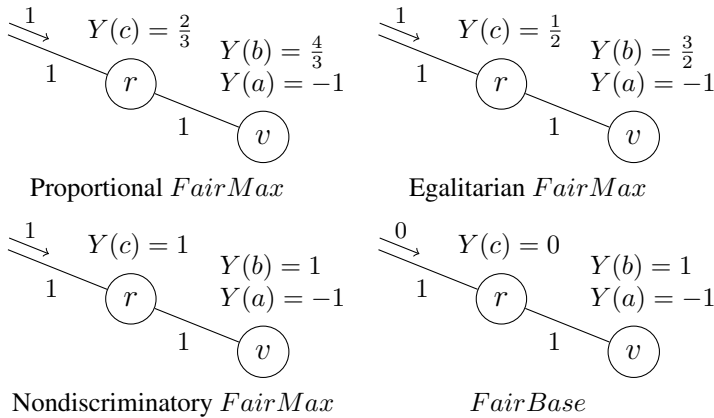


Figure 3.6 *FairMax* allocations with different notions of fair division for the simple congestion network $T = (V, E, A)$ from Figure 3.3. In this example, the *FairBase* allocation is the same for all three principal notions of fair division.

Figure 3.7 shows a more complex congestion tree for which the *FairMax* allocation is computed with the egalitarian notion of fair division. Figure 3.7 also shows the steps that Algorithm 2 takes during this computation. Once the root r is reached, the *FairMax* allocation is computed by adding the Fb values to the Fm values.

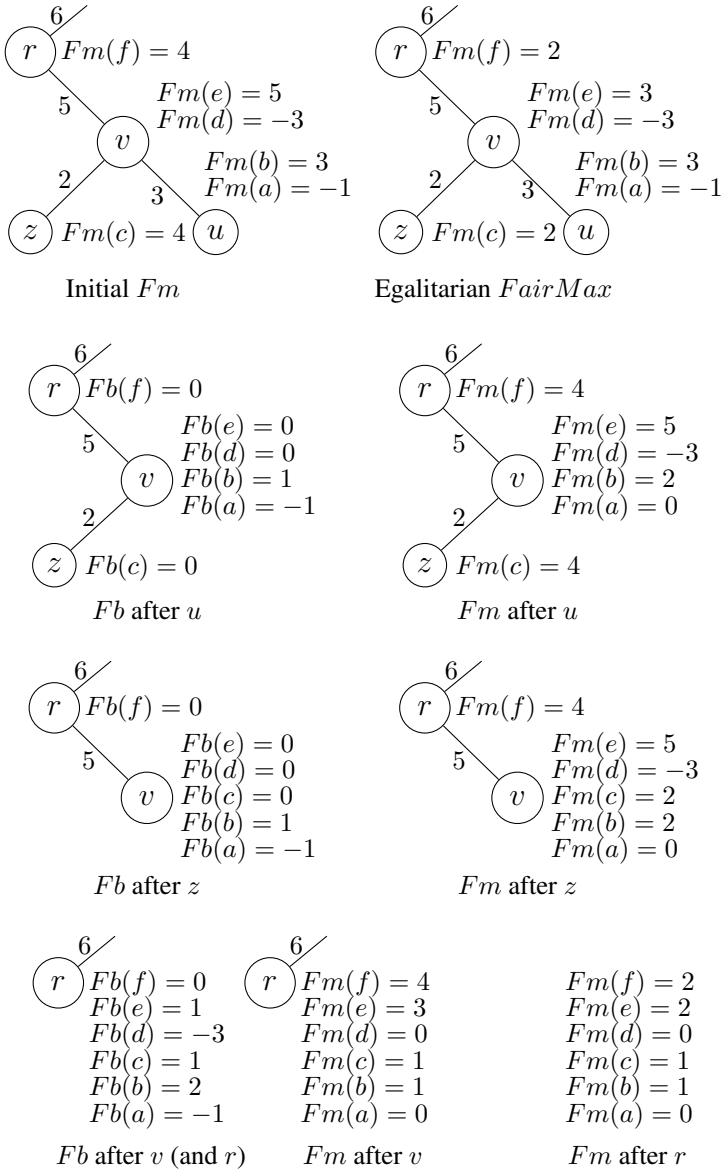


Figure 3.7 Algorithm input (top left), output (top right), and steps from vertices u to r .

3.7 Incentive Compatibility

One of the central concepts in the field of mechanism design is that of incentive compatibility. A mechanism is designed with a certain outcome in mind, for example a congestion free power flow network. The designer of a mechanism lays out the rules in such a way that agents playing the game follow strategies that together reach the intended outcome. This assumes that the agents want to play the game, and that they will play it honestly. An incentive compatible mechanism ensures that agents will participate truthfully.

The importance of truthful participation is emphasized when dealing with fairness. Fairness often heavily depends on the outcomes for the agents relative to each other. If agents can ‘game the system’, the agents that do will obtain an unfair advantage over the other agents. A mechanism that allows such strategies will require strong assumptions about the participating agents or other means of enforcing fairness.

Proposition 3.2. The mechanism presented in Algorithm 2 is incentive compatible when the notion of fairness is egalitarian.

Proof. Consider the Fm output of the mechanism, the proof for the Fb output is analogous. Let $T = (V, E, A)$ be a congestion tree containing an agent a with true desire d_a^* . Without loss of generality, assume that agent a is a consumer, i.e. $d_a^* > 0$.

If agent a reports $d_a = d_a^*$ and the mechanism returns $Fm(a) = d_a^*$, then the agent cannot improve its situation by reporting another desire d_a because it is already in its preferred situation.

If agent a reports $d_a = d_a^*$ and the mechanism returns $Fm(a) < d_a^*$, then the agent may try reporting another desire $d_a \neq d_a^*$ to improve its situation. Since the mechanism returned $Fm(a) < d_a$ when reporting $d_a = d_a^*$, it must be that the value of $Fm(a) + Fb(a)$ was reduced at at least one step of the algorithm. Consider the first step at which this happened, and the vertex v corresponding to that step. It must be that, initially, congestion occurred at this vertex, i.e. $|Fm(A_v^-)| + C_v < Fm(A_v^+)$ at Algorithm 2 Line 7. Since $Fm(a) + Fb(a)$ was reduced at this vertex, it must have been true that $Fb(a) > w$ at Algorithm 3 Line 8. Note that the output w of the water level function in Algorithm 4 does not depend on values $Y(a) > w$.

It now follows that reporting $d_a > d_a^*$ causes the same congestion at vertex v , i.e. $|Fm(A_v^-)| + C_v < Fm(A_v^+)$ at Algorithm 2 Line 7, and produces the same value w at Algorithm 3 Line 7, resulting in the same reduction of $Fm(a) + Fb(a)$ at this step as compared to reporting $d_a = d_a^*$. Thus after this step, there is no difference between reporting $d_a = d_a^*$ and reporting $d_a > d_a^*$, leading the mechanism to return the same value $Fm(a) < d_a^*$ in both cases.

It is possible that reporting $d_a > d_a^*$ causes congestion at an earlier step corresponding to a vertex z . This does not change the argument, since reporting $d_a = d_a^*$ not causing congestion at vertex z implies that when reporting $d_a > d_a^*$, $Fm(a) + Fb(a)$ can at most be reduced to d_a^* at vertex z .

Alternatively, reporting $d_a < d_a^*$ may avoid the congestion at vertex v . However, this is only the case when $Fm(a)$ at Algorithm 2 Line 7 is equal to or lower than the value w at Algorithm 3 Line 7 when reporting $d_a = d_a^*$. In other words, the reduction of $Fm(a) + Fb(a)$ to $w + Fb(a)$ at vertex v is only avoided if it is already equal to or lower than $w + Fb(a)$. The same argument then applies to each following step corresponding to a vertex z with $a \in A_z$ where congestion occurs when reporting $d_a = d_a^*$. It follows that the mechanism output $Fm(a)$ when reporting $d_a < d_a^*$ is equal to or lower than the output $Fm(a)$ when reporting $d_a = d_a^*$.

This shows that agent a cannot improve its situation by reporting anything other than its true desire d_a^* , i.e., being truthful is a weakly dominant strategy. \square

Proposition 3.3. The mechanism presented in Algorithm 2 is not incentive compatible when the notion of fairness is proportional or nondiscriminatory.

Proof. The proof is by counterexample. Let $T = (V, E, A)$ be a simple congestion tree with only one vertex r , a capacity of $C_r = 8$ on its virtual edge, and two agents a and b with true desires $d_a^* = 6$ and $d_b^* = 6$. If both agents report their true desires, the mechanism would return $Fm(a) = 4$ and $Fm(b) = 4$ for both the proportional and the nondiscriminatory notions of fairness. However, if agent a instead reports $d_a = 10$ while agent b still reports its true desire $d_b = 6$, then the mechanism would return $Fm(a) = 5$ and $Fm(b) = 3$ for the proportional notion of fairness, or $Fm(a) = 6$ and $Fm(b) = 2$ for the nondiscriminatory notion of fairness. This shows that agent a can improve its situation by reporting a desire designed for participation in the mechanism, rather than its true desire.

The same counterexample can be used for the Fb output of the mechanism by adding a third agent c with $d_c^* = -8$. \square

Propositions 3.2 and 3.3 can be intuitively understood to hold true by examining Figure 3.2. The three divisions corresponding to the three principal notions of fair division are found at the intersection of their respective lines with the line representing the quantity of the good to be divided.

Two of the three principal notions of fair division correspond to lines that depend on the point of the claims, and thus the point of intersection depends on the point of the claims as well. Evidently, for these two notions, the division correlates directly with the claims reported by the agents. Indeed, for the

proportional and nondiscriminatory notions of fairness, an agent can directly influence the division.

Conversely, for the egalitarian notion of fairness, the point of intersection only depends on claims to a limited extent. Only an agent that is allocated its entire claim would potentially be allocated a larger amount by reporting a higher claim. Agents with claims above a certain threshold cannot increase the amount allocated to them by reporting a higher claim. The egalitarian division therefore only correlates with the claims reported by the agents to the point that each agent can claim their fair equal share and receive it.

3.8 Computational Complexity

The water level function in Algorithm 4 sorts the $Y(a)$ values in $\mathcal{O}(m \cdot \log(m))$ time, where m is the number of agents in the input set A . It then enters a while loop which takes at most $m - 1$ iterations. Thus the computational complexity of the water level function in Algorithm 4 is $\mathcal{O}(m \cdot \log(m))$.

Algorithm 2 visits each vertex in V exactly once, and at each vertex calls the divide function in Algorithm 3 two times. For the proportional notion of fairness, the divide function assigns m values, where m is the number of agents in the input set A . For the egalitarian and nondiscriminatory notions of fairness, the divide function calls the water level function at most once.

Therefore, the worst case computational complexity of Algorithm 2 with the proportional notion of fairness is $\mathcal{O}(n \cdot m)$, while with the egalitarian and nondiscriminatory notions of fairness it is $\mathcal{O}(n \cdot m \cdot \log(m))$. Here n is the number of vertices in V and m is the total number of agents in A .

3.9 Conclusion and Discussion

We presented local, outer matching, Definition 3.10, as a novel concept of fairness for congestion management in low-voltage networks. Local, outer matching addresses the important problem of fairness in matching by prioritizing local matching in the peripheral where prosumers are furthest away from the substation. We then presented congestion solutions in Definition 3.11 that were proven by Theorem 3.1 to be local, outer matchings that make maximal use of the available network capacity. These congestion solutions interchangeably employ established notions of fair division for dividing quantities such as capacity. In Section 3.4 we discussed three distinct principal notions of fair division: proportional, egalitarian, and nondiscriminatory division.

Subsequently, in Section 3.6, we presented an algorithmic mechanism that combines local, outer matching with notions of fair division and computes congestion solutions which fairly resolve congestion and make maximal use

of available network capacity as proven by Theorem 3.2. In Section 3.6.2 we showed that the mechanism is able to employ different notions of fair division, and we then went on to prove that the egalitarian notion of fairness results in an incentive compatible mechanism, Proposition 3.2, while the proportional and nondiscriminatory notions of fairness do not, Proposition 3.3.

Finally, in Section 3.8, we showed that the presented congestion solutions can be computed by an algorithm with low computational complexity. This makes the notion of fairness and the algorithm suitable for sizeable and time sensitive congestion problems such as those encountered in electrical grids.

The egalitarian notion of fair division resulting in an incentive compatible mechanism is an obvious advantage over other notions of fair division, but does not render other notions of fair division obsolete. Consensus on the accepted notion of fair division should be a priority since fairness is fundamentally subjective and dependent on setting. Additional penalties or incentives could be implemented to make other notions of fair division feasible for use in this setting if they are strongly preferred.

The algorithmic mechanism we presented in this chapter is limited to the acyclic networks that are found in real-world low-voltage networks. It is likely that a similar algorithmic mechanism for more general network structures would have a higher computational complexity. Running in limited computational time is, however, essential for the application in this domain.

The theoretical foundation that this chapter lays may be extended to more detailed models, for example including line losses by discounting flows per line that is traversed. This raises the interesting question of whether fairness lies with the sent quantity or the received quantity.

Another potential avenue of research would be to look at other fair ways of matching consumer and producer desires, for example by changing the hierarchical structure of matching or by introducing time-shiftable consumers and producers.

4 Parameterized Globally Max-Min Fair Solutions

In the previous chapter we considered a local notion of fairness. In this chapter we instead consider a global notion of fairness: max-min fairness.

4.1 Introduction

Network flow congestion is often considered in the context of information systems (Bertsekas et al., 1992; Buchbinder and Naor, 2006). In these settings, a connection may have a limited capacity. When a number of users try to send too much data over the connection simultaneously, exceeding the capacity, then congestion occurs. Congestion management solutions seek to mitigate or resolve this congestion, while taking different aspects into account. One such aspect is that of fairness or fair division of network capacity (Brandt et al., 2012; Lang and Rothe, 2016; Moulin, 2003). Fairness is often considered in flow congestion problems in general (Fossati et al., 2018; Ghodsi et al., 2011; Kleinberg et al., 1999) (for a survey, see (Ogryczak et al., 2014)).

Similar network flow congestion issues arise in electricity networks (Bach Andersen et al., 2012). There too does congestion occur when agents attempt to exceed line capacity constraints. With the ongoing energy transition, moving towards more decentralized and renewable energy sources, congestion problems have become more frequent and more severe (Nabe et al., 2017; Verzijlbergh et al., 2014). Since access to electricity networks is considered essential for full participation in modern society (European Parliament, 2016), fairness plays a vital role in electricity network congestion management.

Electricity networks, however, significantly differ from information systems on a fundamental level. The flows in an electricity network are commodity flows, and the agents that populate the network act as consumers or producers of the commodity. The implication of this for congestion management is that consumption and production can be locally balanced, avoiding strain on the network elsewhere. The implication for fairness is that assigning each agent

an equal fair share may be highly suboptimal due to unused local matching opportunities. Previous work (Hekkelman and La Poutré, 2020) offers a greedy method for local matching, resulting in an efficient allocation that satisfies a specific form of local fairness. A solution for the nonlocal case is provided in the seminal work by Megiddo (Megiddo, 1977).

Within the context of electricity networks we focus on the many hierarchical interactions that play a role there. We mainly consider subnetworks, e.g. low-voltage networks which are local networks that are connected to a larger medium-voltage network through a substation. Within the subnetwork the aforementioned issues of local matching and fairness play a primary role. However, the resultant flow of the subnetwork into the higher-level network can play a significant role at that higher level. In situations such as microgrids, aggregation or virtual power plants, the local subnetwork acts in a higher-level market with its flow. This may mean that the subnetwork has to adjust itself internally to yield a certain resultant flow. In this chapter we investigate how such higher-level interactions affect fair congestion solutions within the subnetwork.

In this chapter we consider commodity flows between producers and consumers in congested subtree networks with external connections at the root. This models electricity distribution networks which in many cases have an active radial structure (their graphs are trees (Sallam and Malik, 2018)) with their root connected to a higher-level transmission network. The notion of fairness that we consider is the widely used egalitarian notion of leximin fairness (Jaffe, 1981), which is closely related to envy-freeness and the Nash product (Berliant et al., 1992; Varian, 1974). As pointed out by (Megiddo, 1974), applying fairness to flow networks requires us to consider individual agents as opposed to the usual single-source/sink models.

We extend leximin fairness to multi-agent commodity flow networks in order to capture the unique local matching opportunities that arise between producers and consumers. This natural extension implements leximin fairness, but does so among producers and consumers separately while encouraging local matching of supply and demand where possible. Resulting leximin fair congestion solutions thus possibly feature relatively high consumption and production locally, but only when a more even distribution to other parts of the network is not feasible due to capacity constraints.

In this chapter we prove that for commodity flow in a subtree network with external connection at the root, for every possible outward flow at the root, there exists a unique leximin fair congestion solution, and when no outward flow at the root is specified then there exists a unique global leximin fair congestion solution. For the proof, and later algorithm, we use an insightful method where we compute agent-specific intervals that capture local matching

and prevent congestion. We then apply a water-filling algorithm to these intervals in order to find leximin fair congestion solutions.

We then go on to fully parameterize the unique leximin fair congestion solutions for all possible root flows. We establish a correspondence between possible root flows and a parameter that, given the agent-specific intervals, allows us to find the unique leximin fair allocation with the corresponding root flow for any single agent in $\mathcal{O}(1)$ time. Increasing this parameter yields unique leximin fair allocations with progressively higher root flow values.

We subsequently devise an algorithmic mechanism to compute the leximin fair congestion solutions. Our algorithm closely follows the earlier proof, recursively computing the agent-specific intervals in the same manner. After just one run of the algorithm, we can then use the parameterization to compute the unique leximin fair congestion solution for any possible root flow in $\mathcal{O}(\#agents)$ time. We show a low computational complexity of the algorithm, which is essential for application in e.g. the energy domain.

Finally, we prove that our algorithmic mechanism is individually rational, incentive compatible, and can operate as a distributed algorithm. Especially this second property is of importance when considering fairness, since the possibility of ‘gaming’ the mechanism would undermine any fairness as well.

Our contributions can be summarized as follows:

- We prove that for commodity flow subtrees with an external connection at the root, for each possible root flow there exists a corresponding unique leximin fair congestion solution, as well as a unique global leximin fair congestion solution.
- We fully parameterize the unique leximin fair congestion solutions, finding for any agent and any possible root flow its unique leximin fair allocation in $\mathcal{O}(1)$ time.
- We present an algorithmic mechanism of low computational complexity that computes this parameterization of all unique leximin fair congestion solutions.
- We prove that our mechanism is individually rational, incentive compatible, and can operate as a distributed algorithm.

4.2 Preliminaries

4.2.1 Agent-Based Commodity Flow on Trees

We start by modelling commodity flow in congested tree networks populated by agents that act as prosumers of the commodity, meaning they can either produce or consume. We concern ourselves with allocating prosumption values

to agents so as to accommodate agent desires while ensuring fairness and resolving congestion.

Let a **congestion tree** $T = (V, E, A)$ be a rooted weighted tree (V, E) with a set of agents A located at the vertices. Let a virtual edge at the root r represent the connection to a virtual parent that represents an external network. Let the edge weights be positive, representing flow capacities, and denote the weight of an edge between vertex $v \in V$ and its parent as the **capacity** C_v . Internal flow capacities are strictly positive, while the flow capacity C_r at the root can be zero (representing an isolated network).

Let each agent $a \in A$ report a **desired prosumption** d_a . A prosumption induces a flow between the external network and the agent. An agent's desire also indicates whether the agent acts as a consumer ($d_a > 0$) or a producer ($d_a < 0$). Let $A^+ \subseteq A$ and $A^- \subseteq A$ denote the sets of consumers and producers, respectively.

Congestion occurs when commodity flows induced by the agents' prosumptions result in the tree's edge capacities being exceeded. Congestion solutions come in the form of **allocations** $Y : A \rightarrow \mathbb{R}$ that allocate a prosumption to each agent. An allocation Y completely defines the flow in a congestion tree, with the **root flow** $F(Y)$ given by the sum of allocated prosumptions, i.e. $F(Y) = \sum_{a \in A} Y(a)$.

An allocation Y on a congestion tree $T = (V, E, A)$ is **congestion free** if for each vertex v the root flow $F_v(Y)$ of the subtree $T_v = (V_v, E_v, A_v)$ with root v does not exceed the capacity C_v , i.e. $|F_v(Y)| \leq C_v$. An allocation Y is **desire compatible** if each agent a is allocated a prosumption between zero and its desire d_a . An allocation Y is **feasible** if it is both congestion free and desire compatible. A root flow f for a congestion tree T is feasible if there exists a feasible allocation Y on T with $F(Y) = f$.

4.2.2 Leximin Fairness for Commodity Flows

A widely applied notion of fairness, especially in network flow congestion settings, is that of leximin fairness. Leximin fairness is a very strong criterion that corresponds to leximin-optimal allocations and implies a maximal Nash product. With leximin fairness we are able to maximize the minimum amounts allocated to agents, i.e. an egalitarian rule, while achieving Pareto optimality within the constraints.

Our setting of prosumer commodity flows poses a unique challenge for finding leximin fair allocations because the presence of both consumers and producers allows for local matching to reduce outward flow and avoid congestion. This local matching gives rise to situations where one agent is necessarily allocated more than another, as demonstrated by the situation in Figure 4.1.

To reflect the possibility and desirability of local matching in our setting, we simply extend the classic notion of leximin fairness to absolute values.

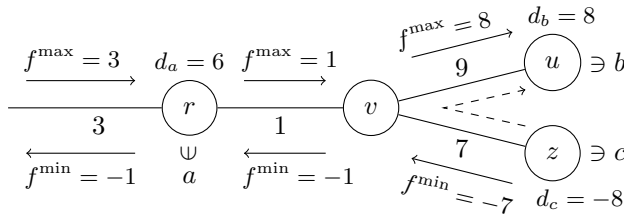


Figure 4.1 An example congestion tree T with agents a, b, c , capacities and desires as indicated. Also indicated are maximal and minimal feasible root flows for subtrees of T . The dashed flow between z and u indicates local matching of at least 6 units. The global leximin fair solution allocates 4, 6, -7 to a, b, c respectively, rather than 4, 4, -5 which enforces strict consumer equality but is not Pareto optimal.

This principle is similar to the two distinct lexicographic orders for producers and consumers used in (Megiddo, 1974).

Definition 4.1. A feasible allocation on a congestion tree is **leximin fair** if and only if increasing, in absolute value, the allocation of any agent necessarily results in the decrease, in absolute value, of the allocation of another agent with an equal or smaller, in absolute value, allocation.

Note that for commodity flows, increasing the allocation of a consumer never results in a necessary decrease in absolute value of the allocation of a producer since both these changes reduce flow. Therefore, Definition 4.1 essentially applies classic leximin fairness separately to consumers and to producers. However, it also allows for matching between the two: increasing, in absolute value, the allocations of both a consumer and a producer. Consequently, leximin fair allocations by this definition are Pareto optimal.

Since we consider divisible goods, on nonzero allocations our extended leximin fairness can be expressed by an optimal Nash product of the absolute values of the allocated presumptions. This relation, given by Lemma 4.1, is analogous to that between classic leximin fairness and the (non-absolute) Nash product (Segal-Halevi and Sziklai, 2019). The well-known association of the Nash product with fair allocations makes it a fitting tool to use to find leximin fair solutions later in certain cases.

Lemma 4.1. Given a congestion tree $T = (V, E, A)$ and a root flow f , if the maximal value of the Nash product $\prod_{a \in A} |Y_f(a)|$ on feasible allocations Y_f on T with root flow f is nonzero, then a feasible allocation Y_f^* on T with root

flow f is leximin fair among feasible allocations Y_f on T with root flow f if and only if it maximizes the Nash product among such allocations.

Proof. Suppose Y_f maximizes the Nash product but is not leximin fair. Then it is possible to increase, in absolute value, the allocation of an agent without decreasing, in absolute value, the allocation of another agent with equal or smaller, in absolute value, allocation while maintaining a root flow f . This change would increase the Nash product, a contradiction. Similarly for the converse. \square

4.3 Main Theorems

We now formulate one of the primary novel contributions of our work, which is proving the existence of unique leximin fair congestion solutions on congestion trees together with their parameterization.

Theorem 4.1. *Given a congestion tree $T = (V, E, A)$ and a feasible root flow f , there exists a unique leximin fair allocation Y_f^* with root flow f on T . Moreover, if a root flow is not specified, there exists a unique global leximin fair allocation Y_{opt}^* on T .*

Theorem 4.2. *Let $T = (V, E, A)$ be a congestion tree, let f^{\min} and f^{\max} be the minimal and maximal feasible root flows, respectively, and let $f^{opt} = F(Y_{opt}^*)$. Then there exists a parameterization of unique leximin fair allocations Y_f^* for all feasible root flows f that is described by two piecewise linear increasing bijections*

$$\begin{aligned} P_+ &: [f^{\min}, f^{opt}] \rightarrow \mathcal{I} \\ P_- &: [f^{opt}, f^{\max}] \rightarrow \mathcal{J} \end{aligned} \tag{4.1}$$

with $\mathcal{I}, \mathcal{J} \subset \mathbb{R}$ closed intervals, such that for each agent $a \in A$ and for each feasible root flow f we can compute $Y_f^*(a)$ as a function of $P_{\pm}(f)$ and a in $\mathcal{O}(1)$ time.

4.4 Proof of the Main Theorems

4.4.1 Outline of the proof

We first set out to prove Theorem 4.1. The core of the proof revolves around finding unique leximin fair allocations for the maximum and minimum possible root flows of subtrees of a congestion tree. These provide for each agent an upper and lower bound between which all leximin fair allocations lie. The

bounds prevent congestion from occurring in the subtree, which then allows us to find the bounds recursively on subtrees. Finally, with the bounds we can find any unique leximin fair allocation. In Sections 4.5 and 4.6 we will show that this method results in a low computational complexity.

First, in Section 4.4.2, we introduce a method for finding unique leximin fair allocations for specific root flows by uniformly curtailing either consumers or producers within their determined bounds.

Subsequently, in Section 4.4.3, we break down congestion trees into equivalent (for congestion purposes) binary congestion trees. These facilitate the proof by allowing examination of smaller steps.

We then present our proof by induction. In Section 4.4.4, we formulate the induction hypothesis. This naturally includes Theorem 4.1, but also includes a number of other properties that will facilitate the proof and, later on, the mechanism. In Section 4.4.5 we verify the induction base, and in Section 4.4.6 we perform the induction step which takes up the majority of this section.

Finally, in Section 4.4.7, we prove Theorem 4.2 using Section 4.4.2 and the agent-specific bounds computed in the proof of Theorem 4.1.

4.4.2 Curtailment Allocations and Root Flow

Consider a congestion (sub)tree populated by agents with associated desired presumption values. To adjust the root flow of this congestion tree, we can curtail agents' presumptions by allocating new values to them that are smaller, in absolute value, than their desired values. In light of leximin fairness, it is never optimal to curtail both consumers and producers at the same time since an equal change in consumption and production results in a root flow change of zero. This is why we consider what we call positive and negative curtailment allocations, which are allocations that respectively curtail only consumers or producers. A positive curtailment solution results in a more negative root flow, while a negative curtailment solution results in a more positive root flow.

Within the context of leximin fairness we further specify positive and negative curtailment allocations to be of a specific form. As the next sections will show, we will look at intervals that specify not only the desired presumption value of an agent, but also the extent to which the agent may be curtailed. To find leximin fair allocations, we curtail a set of agents by choosing a single value w . Each agent is then allocated this value w , *bounded by its interval*.

Definition 4.2. Consider a set of agents A with for each agent $a \in A$ a lower bound l_a and upper bound u_a on its presumption, such that $l_a \leq u_a$ and $0 \notin (l_a, u_a)$. Then a **positive curtailment allocation** for l_a, u_a ($a \in A$) is

an allocation Y , for some $w \geq 0$, with

$$Y(a) = \begin{cases} \min(\max(w, l_a), u_a) & \forall a \in A^+ \\ l_a & \forall a \in A^- \end{cases} \quad (4.2)$$

while a **negative curtailment allocation** for l_a, u_a ($a \in A$) is an allocation Y , for some $w \leq 0$, with

$$Y(a) = \begin{cases} u_a & \forall a \in A^+ \\ \min(\max(w, l_a), u_a) & \forall a \in A^- \end{cases} \quad (4.3)$$

We will refer to an allocation that is a positive or negative curtailment allocation simply as a curtailment allocation.

The extreme cases of maximal curtailment and no curtailment result in straightforward allocations, as the following lemma shows.

Lemma 4.2. Positive curtailment allocations Y with $w \leq \min_{a \in A^+} l_a$ (including $w = 0$) are given by $Y(a) = l_a$ ($a \in A$), while negative curtailment allocations Y with $w \geq \max_{a \in A^-} u_a$ (including $w = 0$) are given by $Y(a) = u_a$ ($a \in A$). Moreover, positive curtailment allocations Y with $w \geq \max_{a \in A^-} u_a$ and negative curtailment allocations Y with $w \leq \min_{a \in A^+} l_a$ are all given by $Y(a) = \min(0, l_a) + \max(0, u_a)$ ($a \in A$).

Proof. This follows directly from Definition 4.2. □

When the root flow is increased or decreased using curtailment, then the value allocated to each agent also respectively increases or decreases (possibly by zero). The following lemma formalizes this.

Lemma 4.3. Given a set of agents A with lower and upper bounds $l_a \leq u_a$, $0 \notin (l_a, u_a)$ ($a \in A$), if Y and Y' are two curtailment allocations with respective root flows f and f' such that $f \leq f'$, then $Y(a) \leq Y'(a)$ for all agents $a \in A$.

Proof. First, note that $l_a \leq \min(\max(w, l_a), u_a) \leq u_a$ for any w . If Y and Y' are positive and negative curtailment allocations respectively, then $Y(a) \leq Y'(a)$ ($a \in A$) follows directly from Definition 4.2. The converse implies $f = f'$ and thus $Y(a) = Y'(a)$ ($a \in A$).

If Y and Y' are both positive curtailment allocations, then it must be that $w \leq w'$. This implies $Y(a) \leq Y'(a)$ ($a \in A$). Analogously when Y and Y' are both negative curtailment allocations. □

Finally, we show that curtailment allocations uniquely maximize the Nash product for their specific root flow.

Lemma 4.4. Given a set of agents A with lower and upper bounds l_a, u_a ($a \in A$) such that $l_a \leq u_a$, $0 \notin (l_a, u_a) \forall a \in A$ and a value f such that $\sum_{a \in A} l_a \leq f \leq \sum_{a \in A} u_a$, then there exists a unique solution $s = \{s_a \mid a \in A\}$ to the optimization problem:

$$\begin{aligned} & \text{maximize } \prod_{a \in A} |x_a| \\ & \text{subject to } \sum_{a \in A} x_a = f, \\ & \quad x_a \in [l_a, u_a] \forall a \in A. \end{aligned} \tag{4.4}$$

Furthermore, the unique solution s is a positive or negative curtailment allocation for l_a, u_a ($a \in A$).

Proof. We start by considering the simplified problem where $0 \leq l_a \leq u_a$ ($a \in A$). That a solution $s = \{s_a \mid a \in A\}$ exists follows from the given properties. Assume there does not exist a $w \geq 0$ such that $s_a = \min(\max(w, l_a), u_a)$ for all agents $a \in A$. For all $w \geq 0$, if $s_a = u_a$ for all agents a with $s_a < w$ and $s_a = l_a$ for all agents a with $s_a > w$, then $s_a = \min(\max(w, l_a), u_a)$ for all agents a . In particular, taking $w = 0$, this implies that there exists an agent a with $s_a > l_a$. Among such agents, let b be the agent with the highest value for s_b . Taking $w = s_b$ then implies that there exists an agent a with either $s_a > w$ and $s_a > l_a$, which would contradict the choice of agent b , or $s_a < w$ and $s_a < u_a$. In summary, there exist two agents a, b with $s_a < s_b$, $s_a < u_a$, and $s_b > l_b$.

Now let $\epsilon = \min(\frac{s_b - s_a}{2}, u_a - s_a, s_b - l_b)$. Because of the three inequalities involving agents a and b we know that $\epsilon > 0$, and that that $\{s'_a \mid a \in A\}$ with $s'_a = s_a + \epsilon$, $s'_b = s_b - \epsilon$, and $s'_c = s_c \forall c \neq a, b$ satisfies the optimization constraints. Then since

$$\begin{aligned} s_a s_b &= (s'_a - \epsilon)(s'_b + \epsilon) \\ &= s'_a s'_b + \epsilon(s'_a - s'_b) - \epsilon^2 \\ &< s'_a s'_b + \epsilon(s'_a - s'_b) \\ &= s'_a s'_b + \epsilon(s_a - s_b + 2\epsilon) \\ &\leq s'_a s'_b + \epsilon(s_a - s_b + s_b - s_a) \\ &= s'_a s'_b \end{aligned} \tag{4.5}$$

it follows that

$$\prod_{c \in A} s_c = \left(\prod_{c \neq a, b} s_c \right) \cdot s_a s_b < \left(\prod_{c \neq a, b} s'_c \right) \cdot s'_a s'_b = \prod_{c \in A} s'_c \quad (4.6)$$

and thus that $s = \{s_a \mid a \in A\}$ is not a solution. This proves that any solution can be written as $\{s_a = \min(\max(w, l_a), u_a) \mid a \in A\}$ for some $w \geq 0$, from which uniqueness follows.

Finally, we consider the original problem with both positive and negative bounds. Without loss of generality, assume that $\sum_{a \in A^-} l_a + \sum_{a \in A^+} u_a \geq f$. Since $\sum_{a \in A} l_a \leq f$, it follows that $\sum_{a \in A^+} l_a \leq f - \sum_{a \in A^-} l_a \leq \sum_{a \in A^+} u_a$. Thus, by the simplified version of Lemma 4.4 for exclusively positive bounds that we proved, there is a unique solution $\{s_a = \min(\max(w, l_a), u_a) \mid a \in A^+\}$ for a $w \geq 0$ such that $\sum_{a \in A^+} s_a = f - \sum_{a \in A^-} l_a$. A unique solution to the original problem is then $\{s_a = \min(\max(w, l_a), u_a) \mid a \in A^+\} \cup \{s_a = l_a \mid a \in A^-\}$ with this same $w \geq 0$, which is a positive curtailment allocation. \square

4.4.3 Breaking Down in Binary Congestion Trees

We break congestion trees down into equivalent (for congestion solutions) binary trees with less parameters (agents, children) per vertex. This reduces the complexity of induction.

Given a congestion tree $T = (V, E, A)$ we construct an equivalent binary tree T^{bin} called a binary congestion tree of T . T^{bin} allows capacity constraints to be infinite (technically, sufficiently large is enough, e.g. $\sum_{a \in A} |d_a|$). For each vertex v of T , create $n + \max(n + m - 2, 0)$ new vertices, where n is the number of agents at v and m is the number of children of v . See Figure 4.2 for an example. $\max(n + m - 2, 0)$ of these new vertices, called ‘connector vertices’, are consecutive children of v and each other, and all have infinite upwards capacity. The n remaining new vertices, called ‘agent vertices’, are children of v and the connector vertices, each containing exactly one agent a , and also with infinite upwards capacity. The original children of v in T become children of the rest of the connector vertices and retain their original upwards capacities.

Note that although a congestion tree T may have leaf vertices without agents, these can be (repeatedly) removed to obtain an equivalent (for congestion solutions) congestion tree which then translates to a binary congestion tree with agents located only at leaf vertices, and exactly one agent per leaf.

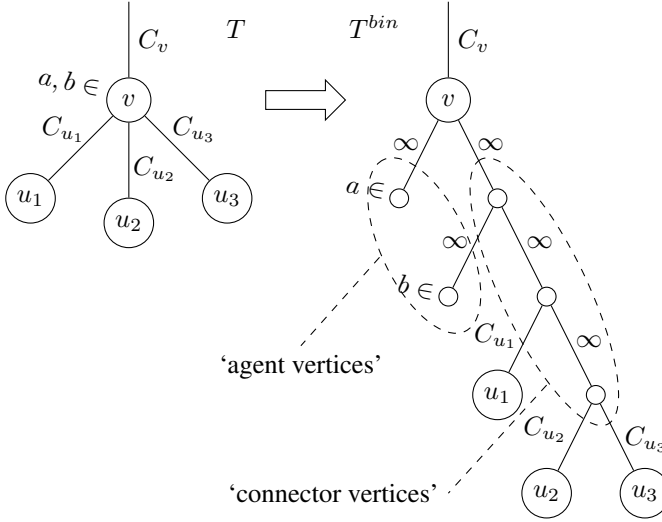


Figure 4.2 The congestion tree T and an equivalent binary congestion tree T^{bin} , depicted around a vertex v , its agents a, b and its children u_1, u_2, u_3

4.4.4 Formulating the Induction Hypothesis

We hypothesize two main properties and eight supporting properties. Given a (binary) congestion tree T , the following two main properties, corresponding to Theorem 4.1, hold:

- (a) For each feasible root flow f there exists a unique leximin fair allocation Y_f^* among feasible allocations with root flow f ,
- (b) There exists a unique leximin fair allocation Y_{opt}^* among feasible allocations.

Moreover, for all agents $a \in A$ there exist two unique values l_a and u_a such that the following supporting properties hold:

1. $l_a \leq u_a$ and $0 \notin (l_a, u_a)$ for all agents a ,
2. Any allocation Y with $Y(a) \in [l_a, u_a]$ for all agents a is feasible on T ,
3. For the minimal and maximal feasible root flows f^{\min}, f^{\max} we have $Y_{f^{\min}}^*(a) = l_a$ and $Y_{f^{\max}}^*(a) = u_a$ for all agents a ,
4. For any feasible root flow f , $Y_f^*(a) \in [l_a, u_a]$ for all agents a ,
5. $Y_{opt}^*(a) = \min(0, l_a) + \max(0, u_a)$ for all agents a .

Finally, we have the following supporting properties to facilitate the proof and later the parameterization and algorithmic mechanism:

6. For any feasible root flow f , Y_f^* is a positive or negative curtailment allocation for l_a, u_a ($a \in A$),
7. $f^{\min} < f^{\max}$ if and only if $C_r > 0$, and always $f^{\min} \leq 0 \leq f^{\max}$.
8. If $C_r > 0$, then $\max(|l_a|, |u_a|) > 0$ for all agents a .

4.4.5 Verifying the Induction base

Consider a leaf vertex v of a binary congestion tree T . Since v is a leaf vertex, there is exactly one agent a located at v . Without loss of generality this agent a can be assumed to be a consumer, i.e. $d_a > 0$. Since the upwards capacity is infinite, the only constraint on feasible root flows is desire compatibility. It follows that $f^{\min} = 0$ and $f^{\max} = d_a$.

Given the existence of a minimal and maximal feasible root flow, we can immediately see that a unique leximin fair allocation Y_{opt}^* is given by $Y_{opt}^*(a) = f^{\max}$. Furthermore, for any root flow value $f \in [f^{\min}, f^{\max}]$, we see that there exists a unique leximin fair allocation Y_f^* with root flow f which is given by $Y_f^*(a) = f$. Now take $l_a = Y_{f^{\min}}^*(a) = 0$ and $u_a = Y_{f^{\max}}^*(a) = d_a$. With this choice of l_a and u_a , the supporting properties 1–8 hold trivially.

4.4.6 Performing Induction on Subtrees

For the induction step, consider a binary congestion tree T with root vertex r and child vertices L and R . Assume that the induction hypothesis holds for T_L , with l'_a, u'_a ($a \in A_L$) denoting the associated values. Similarly, assume that the induction hypothesis holds for T_R , with l'_a, u'_a ($a \in A_R$) denoting the associated values. Since the vertex r is not a leaf vertex, there are no agents located at the vertex r . Thus $A = A_L \cup A_R$. This means that for all agents $a \in A$ we have two values l'_a and u'_a as described by the induction hypotheses for T_L and T_R . We will later find values l_a and u_a for all agents $a \in A$ such that the induction hypothesis also holds for T .

The proof is structured around steps labelled **I.** through **VI.**

- I.** Find the minimal and maximal feasible root flows f^{\min} and f^{\max} for T .
- II.** Show that a root flow f for T is feasible if and only if $f \in [f^{\min}, f^{\max}]$.
- III.** Show that property (a) of the hypothesis holds for T for:
 - III.a.** the edge cases $f^{\min} = f_L^{\min} + f_R^{\min}$ and $f^{\max} = f_L^{\max} + f_R^{\max}$,
 - III.b.** the other cases, where we can use the Nash product.
- IV.** Define the values l_a and u_a for all agents $a \in A$.
- V.** Show that properties 1 – 8 of the hypothesis hold for T .

VI. Show that property (b) of the hypothesis holds for T .

I. First, we establish the minimal and maximal feasible root flows f^{\min} and f^{\max} for T . Since there are no agents located at the vertex r , the minimal feasible root flow f^{\min} is determined by just two factors: the minimal feasible root flows of the subtrees T_L and T_R , and the minimal flow that the root capacity C_r of the tree T allows. By property 3 of the hypothesis for T_L , the minimal feasible root flow of T_L is given by $f_L^{\min} = \sum_{a \in A_L} l'_a$. Similarly, the minimal feasible root flow of T_R is given by $f_R^{\min} = \sum_{a \in A_R} l'_a$. Thus the minimal flow at r that the subtrees T_L and T_R can realise together this way equals $\sum_{a \in A} l'_a = f_L^{\min} + f_R^{\min}$. Meanwhile, the minimal flow at the root r that is allowed by the upwards capacity constraint is $-C_r$. Since both of these two factors provide a lower bound on the feasible root flow, the minimal feasible root flow f^{\min} for T is given by

$$f^{\min} = \max(-C_r, f_L^{\min} + f_R^{\min}). \quad (4.7)$$

Following an analogous argumentation, the maximal feasible root flow f^{\max} for T is given by

$$f^{\max} = \min(C_r, f_L^{\max} + f_R^{\max}). \quad (4.8)$$

II. By the definition of the minimal and maximal feasible root flows f^{\min} and f^{\max} for T , a root flow $f \notin [f^{\min}, f^{\max}]$ cannot be feasible. To see the converse, consider a root flow $f \in [f^{\min}, f^{\max}]$. By the definition of f^{\min} and f^{\max} we have $f \in [\sum_{a \in A} l'_a, \sum_{a \in A} u'_a]$. In addition, by property 1 for T_L and T_R , we have $l'_a \leq u'_a$ ($a \in A$). Thus there exists an allocation Y with $Y(a) \in [l'_a, u'_a]$ for all agents $a \in A$ such that $F(Y) = \sum_{a \in A} Y(a) = f$. By property 2 for T_L and T_R , this allocation Y is feasible on both T_L and T_R . By the definition of f^{\min} and f^{\max} we also have $f \in [-C_r, C_r]$. Thus the allocation Y is feasible on T , which makes f a feasible root flow.

III. Next, given a root flow $f \in [f^{\min}, f^{\max}]$, we will find a leximin fair allocation Y_f^* among feasible allocations on T with root flow f and prove that it is unique. Take a root flow $f \in [f^{\min}, f^{\max}]$. A feasible allocation X on T with $F(X) = f$ exists because f is a feasible root flow.

III.a. If $f = f_L^{\min} + f_R^{\min}$, then the only leximin fair allocations on T_L and T_R with root flows summing up to f are the unique leximin fair allocations for the root flows f_L^{\min} and f_R^{\min} on T_L and T_R respectively. Since an allocation on a (binary) congestion tree can only be leximin fair if it is leximin fair on each subtree, the allocation Y on T that consists of these two unique leximin fair allocations on T_L and T_R is the only allocation on T with root flow f that could be leximin fair.

To see that this allocation Y is leximin fair, we directly apply the definition of leximin fairness. Increasing the allocation of an agent in A_L while decreasing the allocation of an agent in A_R , both relative to Y , results in a root flow for T_R that is smaller than f_R^{\min} . Thus the resulting allocation would not be feasible on T_R and therefore not feasible on T . Then since Y is already uniquely leximin fair on T_L as well as T_R , we now see that no agent can have its (absolute) allocation increased without decreasing the (absolute) allocation of another agent with smaller (absolute) allocation, both relative to Y . It follows that Y is the unique leximin fair allocation Y_f^* on T with root flow f .

Consequently, in this case $Y_f^*(a) = l'_a$ ($a \in A$) and thus, by Lemma 4.2, Y_f^* is a positive curtailment allocation with $w = 0$ with respect to l'_a, u'_a ($a \in A$). Similarly if $f = f_L^{\max} + f_R^{\max}$.

III.b. Now consider $f_L^{\min} + f_R^{\min} < f < f_L^{\max} + f_R^{\max}$. Since L and R are children of r , we have $C_L > 0$ and $C_R > 0$. By property 7 for T_L and T_R this means that $f_L^{\min} < f_L^{\max}$ and $f_R^{\min} < f_R^{\max}$. Thus there exist an f_L and f_R such that $f = f_L + f_R$, $f_L^{\min} < f_L < f_L^{\max}$, and $f_R^{\min} < f_R < f_R^{\max}$. By property 6 for T_L , the allocations $Y_{f_L^{\min}}^*, Y_{f_L}^*, Y_{f_L^{\max}}^*$ are curtailment allocations for l'_a, u'_a ($a \in A$). So by property 3 for T_L and Lemma 4.2, $Y_{f_L^{\min}}^*$ and $Y_{f_L^{\max}}^*$ are positive and negative curtailment allocations with $w = 0$, respectively. By Lemma 4.3, this means that $Y_{f_L}^*$ is a curtailment allocation with $w \neq 0$. Similarly for $Y_{f_R}^*$. By property 8 for T_L and T_R , this implies that the Nash product on these two allocations is nonzero.

Let the allocation Y be a Nash-optimal allocation among feasible allocations on T with root flow f . In other words,

$$\begin{aligned} Y \in \operatorname{argmax}_{X \text{ on } T} \prod_{a \in A} |X(a)| \\ \text{subject to } F(X) = f, \\ X \text{ feasible.} \end{aligned} \quad (4.9)$$

We know that the Nash product is nonzero on this allocation Y since there exist feasible allocations on T_L and T_R with root flows summing up to f , on which the Nash product is nonzero. The allocation Y consists of two feasible allocations Y_L on T_L and Y_R on T_R with $F_L(Y_L) = f_L$ and $F_R(Y_R) = f_R$ for some f_L and f_R such that $f_L + f_R = f$. Since

$$\prod_{a \in A} |Y(a)| = \left(\prod_{a \in A_L} |Y_L(a)| \right) \cdot \left(\prod_{a \in A_R} |Y_R(a)| \right), \quad (4.10)$$

it must then be that Y_L and Y_R are Nash-optimal allocations for f_L on T_L and for f_R on T_R respectively.

Thus, by Lemma 4.1 and property (b) for T_L and T_R , it must be that the allocations Y_L and Y_R are the unique Nash-optimal allocations on T_L and T_R with root flows f_L and f_R respectively. By property 4 for T_L and T_R , it then holds that $Y_L \in [l'_a, u'_a]$ ($a \in A_L$) and $Y_R \in [l'_a, u'_a]$ ($a \in A_R$). Since this holds regardless of the values of f_L and f_R , it follows that $Y(a) \in [l'_a, u'_a]$ ($a \in A$).

Now consider the following optimization problem:

$$\begin{aligned} & \underset{X \text{ on } T}{\text{maximize}} && \prod_{a \in A} |X(a)| \\ & \text{subject to} && \sum_{a \in A} X(a) = f, \\ & && X(a) \in [l'_a, u'_a] \forall a \in A. \end{aligned} \tag{4.11}$$

As we have seen, the allocation Y is in the search space of problem (4.11). For an allocation X in this search space, $X(a) \in [l'_a, u'_a]$ ($a \in A$) implies that X is feasible on T_L and T_R and thus $F(X) = f \in [f^{\min}, f^{\max}] \subseteq [-C_r, C_r]$ implies that X is feasible on T as well. It follows that the search space of problem (4.11) is a subset of the search space of problem (4.9). Thus since Y is in the search space of problem (4.11) and is a solution to problem (4.9), Y must be a solution to problem (4.11) as well.

By definition of f^{\min} and f^{\max} , $f \in [f^{\min}, f^{\max}] \subseteq [\sum_{a \in A} l'_a, \sum_{a \in A} u'_a]$. By property 1 for T_L and T_R , we have $l'_a \leq u'_a$ and $0 \notin (l'_a, u'_a)$ for all agents $a \in A$. This means we can use Lemma 4.4 to conclude that problem (4.11) has a unique solution, which we know to be Y , and that Y is a positive or negative curtailment allocation with respect to l'_a, u'_a ($a \in A$). Since we have shown that any solution to problem (4.9) is a solution to problem (4.11), we can conclude that Y is also the unique solution to problem (4.9) and is thus the unique Nash-optimal allocation among feasible allocations on T with root flow f . By Lemma 4.1, this unique Nash-optimal allocation is the unique leximin fair allocation Y_f^* among feasible allocations on T with root flow f .

IV. Now that we have proven property (a) for T , we can use those unique leximin fair allocations for feasible root flows to define the values l_a and u_a described by the induction hypothesis for T as follows:

$$\begin{aligned} l_a &= Y_{f^{\min}}^*(a) \\ u_a &= Y_{f^{\max}}^*(a) \end{aligned} \quad \forall a \in A. \tag{4.12}$$

V. We can now prove properties 1 – 8 of the hypothesis for T :

- (3) To prove: for the minimal and maximal feasible root flows f^{\min} and f^{\max} we have $Y_{f^{\min}}^*(a) = l_a$ and $Y_{f^{\max}}^*(a) = u_a$ ($a \in A$).
This holds by definition of l_a and u_a ($a \in A$) in (4.12).
- (6) To prove: for any feasible root flow f , Y_f^* is a positive or negative curtailment allocation with respect to the values l_a, u_a ($a \in A$).
Let f be a feasible root flow. From **III.a.** and **III.b.** we know that Y_f^* is a positive or negative curtailment allocation with respect to l'_a, u'_a ($a \in A$). Since $f^{\min} \leq f \leq f^{\max}$, by property 3 for T and Lemma 4.3 for l'_a, u'_a ($a \in A$), we have $l_a \leq Y_f^*(a) \leq u_a$ ($a \in A$). It follows that Y_f^* is also a positive or negative curtailment allocation with respect to the values l_a, u_a ($a \in A$).
- (4) To prove: for any feasible root flow f , $Y_f^*(a) \in [l_a, u_a]$ ($a \in A$).
Let f be a feasible root flow. Since $f^{\min} \leq f \leq f^{\max}$, by properties 3 and 6 for T and Lemma 4.3 for l_a, u_a ($a \in A$) twice, it follows that $l_a \leq Y_f^*(a) \leq u_a$ ($a \in A$). Thus we have $Y_f^*(a) \in [l_a, u_a]$ ($a \in A$).
- (1) To prove: $l_a \leq u_a$ and $0 \notin (l_a, u_a)$ ($a \in A$).
Since $f^{\min} \leq f^{\max}$, by property 6 for T and Lemma 4.3 for l_a, u_a ($a \in A$), we have $l_a \leq u_a$ ($a \in A$). By property 4 for T_L and T_R , $Y_{f^{\min}}^*(a) \in [l'_a, u'_a]$ and $Y_{f^{\max}}^*(a) \in [l'_a, u'_a]$ ($a \in A$). Since $l_a = Y_{f^{\min}}^*(a)$ and $u_a = Y_{f^{\max}}^*(a)$ ($a \in A$) by property 3 for T , this implies that $[l_a, u_a] \subset [l'_a, u'_a]$ ($a \in A$). Then, by property 1 for T_L and T_R , we have $0 \notin (l'_a, u'_a)$ ($a \in A$) which now implies $0 \notin (l_a, u_a)$ ($a \in A$) as required.
- (2) To prove: any allocation Y with $Y(a) \in [l_a, u_a]$ ($a \in A$) is feasible on T .
As shown above, we have $[l_a, u_a] \subset [l'_a, u'_a]$ ($a \in A$). Thus, by property 2 for T_L and T_R , an allocation Y on T with $Y(a) \in [l_a, u_a]$ ($a \in A$) is feasible on both T_L and T_R . By property 1 for T , we have $l_a \leq u_a$ ($a \in A$) and thus $F(Y) \in [\sum_{a \in A} l_a, \sum_{a \in A} u_a] = [f^{\min}, f^{\max}]$. Then since $[f^{\min}, f^{\max}] \subset [-C_r, C_r]$, it follows that $F(Y) \in [-C_r, C_r]$. This means that the allocation Y is feasible on T as well.
- (7) To prove: $f^{\min} < f^{\max}$ if and only if $C_r > 0$, and always $f^{\min} \leq 0 \leq f^{\max}$.
We know from (4.7) and (4.8) that $f^{\max} = \min(C_r, f_L^{\max} + f_R^{\max})$ and $f^{\min} = \max(-C_r, f_L^{\min} + f_R^{\min})$. So by property 7 for T_L and T_R , we find that $f^{\min} \leq 0 \leq f^{\max}$. If $C_r = 0$, then the only feasible root flow is $f = 0$ and thus $f^{\min} = f^{\max}$. If $C_r > 0$, consider that by property 7 for T_L and T_R we have $f_L^{\min} + f_R^{\min} < 0$ and/or $0 < f_L^{\max} + f_R^{\max}$. It follows that $\max(-C_r, f_L^{\min} + f_R^{\min}) < 0$ and/or $0 < \min(C_r, f_L^{\max} + f_R^{\max})$ and thus that $f^{\min} < f^{\max}$.

(8) To prove: if $C_r > 0$, then $\max(|l_a|, |u_a|) > 0$ ($a \in A$).

By properties 1 and 8 for T_L and T_R , we have $u'_a > 0$ ($a \in A^+$). From **III.a.** and **III.b.** we know that $Y_{f^{\max}}^*$ is a positive or negative curtailment allocation with respect to l'_a, u'_a ($a \in A$). If $Y_{f^{\max}}^*$ is a negative curtailment allocation, then $u_a = u'_a$ ($a \in A^+$) and thus $u_a > 0$ ($a \in A^+$). If $Y_{f^{\max}}^*$ is not a negative curtailment allocation, then $f^{\max} = C_r$ and it is a positive curtailment solution for some $w > 0$ since $C_r > 0$. In this case, since $u'_a > 0$ ($a \in A^+$) and $w > 0$, it also follows that $u_a > 0$ ($a \in A^+$). Similarly for $Y_{f^{\min}}^*$ and $l_a < 0$ ($a \in A^-$).

(5) To prove: $Y_{opt}^*(a) = \min(0, l_a) + \max(0, u_a)$ ($a \in A$).

If $C_r = 0$ then, by property 7, $f^{\min} = f^{\max}$. In this case there is only one leximin fair allocation, which is $Y_{opt}^*(a) = l_a = u_a$ ($a \in A$). If $C_r > 0$ then, by property 8 for T , $l_a < 0$ ($a \in A^-$) and $u_a > 0$ ($a \in A^+$). In this case, the allocation Y with $Y(a) = u_a$ ($a \in A^+$) and $Y(a) = l_a$ ($a \in A^-$) is nonzero. By property 2, this Y is feasible. Thus, by property (a) and 4 for T , Y feasibly and uniquely maximizes the Nash product on T . Then by Lemma 4.1, Y is leximin fair on T and therefore $Y = Y_{opt}^*$. By property 1 for T , Y_{opt}^* can be written as $Y_{opt}^*(a) = \min(0, l_a) + \max(0, u_a)$ ($a \in A$).

VI. Property 5 shows that there exists a unique Nash-optimal allocation Y_{opt}^* among feasible allocations on T . \square

4.4.7 Parameterization of Allocations

The parameterization of leximin fair allocations follows easily at this point, since we have essentially already used it in the preceding proof. Let $T = (V, E, A)$ be a congestion tree and let l_a, u_a for each agent $a \in A$ be the individual lower and upper bounds defined by properties 1 through 5 of the induction hypothesis. With those, let $Y^+(w, a)$ and $Y^-(w, a)$ be the positive and negative curtailment functions from Equations (4.2) and (4.3) respectively. Then let

$$\begin{aligned} w_{\pm}^{\min} &= \max \{ w \mid F(Y^{\pm}(w)) = \min_x F(Y^{\pm}(x)) \} \\ w_{\pm}^{\max} &= \min \{ w \mid F(Y^{\pm}(w)) = \max_x F(Y^{\pm}(x)) \}. \end{aligned} \quad (4.13)$$

Now take P_+ the bijection from $[f^{\min}, f^{opt}]$ to $\mathcal{I} = [w_+^{\min}, w_+^{\max}]$ such that for all agents $a \in A$ and all root flows $f \in [f^{\min}, f^{opt}]$,

$$Y_f^*(a) = \begin{cases} \min(\max(P_+(f), l_a), u_a) & \forall a \in A^+ \\ l_a & \forall a \in A^- \end{cases} \quad (4.14)$$

and similarly for P_- from $[f^{opt}, f^{\max}]$ to $\mathcal{J} = [w_-^{\min}, w_-^{\max}]$.

Since, by property 6 of the induction hypothesis, all leximin fair allocations Y_f^* on T are curtailment allocations, these bijections P_+ and P_- exist. As Equation (4.14) shows, the parameterization P_+ lets us compute $Y_f^*(a)$ for any agent $a \in A$ and any root flow $f \in [f^{\min}, f^{\text{opt}}]$ in only $\mathcal{O}(1)$ time. This proves Theorem 4.2.

4.5 A Leximin Fair Mechanism

4.5.1 Curtailment using a Water Level Algorithm

Algorithm 5 presents our water level algorithm, which essentially implements Lemma 4.4 by taking a value f and a set of agents A with lower and upper bounds $0 \leq l'_a \leq u'_a$ ($a \in A$), and returning a value w such that $\sum_{a \in A} \min(\max(w, l'_a), u'_a) = f$ or as close to f as possible. We refer to w as the ‘water level’, see also Figure 4.3.

Algorithm 5: WaterLevel (f, A, I)

Input: A positive flow f , a set of agents A , and an associated set of bounds $l'_a, u'_a \in I$ with $0 \leq l'_a \leq u'_a$ ($a \in A$).

Output: A value w (also ‘water level’) such that

$$\sum_{a \in A} \min(\max(w, l'_a), u'_a) = \min(\max(f, \sum_{a \in A} l'_a), \sum_{a \in A} u'_a).$$

```

1 if  $f \geq \sum_{a \in A} u'_a$  then
2   |  $w \leftarrow \max_{a \in A} u'_a$ 
3 else if  $f \leq \sum_{a \in A} l'_a$  then
4   |  $w \leftarrow 0$ 
5 else
6   | lowerbounds  $\leftarrow \text{Sort}(l'_a \mid a \in A)$ 
7   | upperbounds  $\leftarrow \text{Sort}(u'_a \mid a \in A)$ 
   | // Satisfy lower bounds
8   | remainingflow  $\leftarrow f - \sum_{a \in A} l'_a$ 
   | // Divide remainingflow over agents
9   | while remainingflow  $> 0$  do
   |   | // Select next lowest bound as new  $w$  and
   |   |   allocate equal parts of remainingflow to
   |   |   agents  $a$  with  $l'_a < w \leq u'_a$ 
   |   | // If insufficient remainingflow is left, set  $w$ 
   |   |   to divide it over eligible agents
10 return  $w$ 

```

If f can be exactly divided over the agents within their individual bounds we first satisfy all the lower bounds. This leaves a remaining quantity $f - \sum_{a \in A} l'_a$ to be divided. Starting with a level at the lowest lower bound and proceeding through consecutively increasing lower and upper bounds, an equal part of this remaining flow is allocated to all agents for which that level lies between their bounds. When not enough remaining flow is left to allocate to all eligible agents up to the next level, the remaining flow is simply evenly divided over them. This procedure yields the desired ‘water level’ output w . See Figure 4.3 for a visualization.

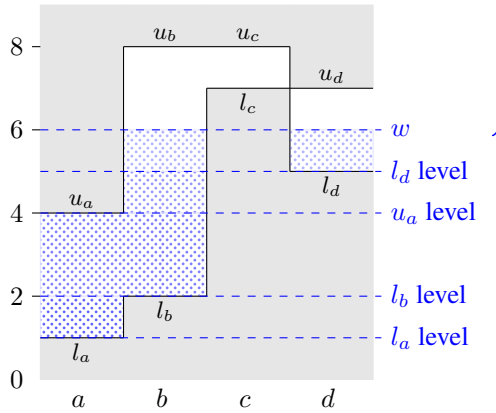


Figure 4.3 An Algorithm 5 example represented as ‘rising water in a cave’. Input is $f = 23$ and four agents with bounds as indicated. After meeting all lower bounds, 8 flow remains. This is divided by raising the water level, reducing the remaining flow to 7, 3, 2 and finally 0 at $w = 6$. Curtailment with the output w yields $Y(a) = 4$, $Y(b) = 6$, $Y(c) = 7$, $Y(d) = 6$.

4.5.2 A Mechanism for Leximin Fair Allocations

From Section 4.4 we know that for a (binary) congestion tree $T = (V, E, A)$ there exists a unique set of values l_a, u_a ($a \in A$) that between them contain all the information we need. Specifically, Section 4.4.7 summarizes that all unique leximin fair allocations, including Y_{opt}^* , are obtained as curtailment allocations for l_a, u_a ($a \in A$).

Importantly, we know that the values l_a, u_a ($a \in A$) for T can be recursively obtained from their subtree equivalents through the use of curtailment allocations. This lets us define a recursive algorithm on subtrees that almost automatically follows from the induction proof of Section 4.4. This algorithm, presented in Algorithm 6, takes the presumption desires of agents in a binary congestion tree T and returns the unique values l_a, u_a ($a \in A$) for T . As such,

Algorithm 6 constitutes a mechanism by taking agents' prosumption desires and returning the unique leximin fair congestion solution Y_{opt}^* .

From property 3 of the induction hypothesis we know that l_a, u_a ($a \in A$) are defined by $Y_{f_{min}}^*$ and $Y_{f_{max}}^*$. We also know that these allocations are curtailment allocations for subtree values l'_a, u'_a ($a \in A$). For agents a not at the root, l'_a and u'_a are assigned at lines 4–8 through subtree recursion. Having obtained l'_a, u'_a ($a \in A$), we can use curtailment with Algorithm 5 to find $Y_{f_{min}}^*$ and $Y_{f_{max}}^*$ and hence find l_a, u_a ($a \in A$). Since Algorithm 5 only works with consumers, we will compute our curtailment allocation in two parts: one for A^+ and one for A^- . For the A^- part, which is symmetric to the A^+ part, we can switch the roles of l'_a and u'_a at lines 14 and 18. We then change the sign of the output to get a negative w .

What is left to check is that the input flows for curtailment at lines 9, 13, 17 and 21 are correct, and that our two part curtailment approach indeed results in curtailment allocations. We check cases based on the upwards capacity C_r .

Consider $Y_{f_{max}}^*$ and the u_a ($a \in A$) which it defines.

- If $\sum_{a \in A^+} u'_a + \sum_{a \in A^-} l'_a \geq C_r$ then $f^{\max} = C_r$ and we need positive curtailment to obtain $\sum_{a \in A} u_a = C_r$. Since in this case $\sum_{a \in A^+} u'_a - C_r \geq -\sum_{a \in A^-} l'_a$, at line 14 we have $\text{maxnegflow} \geq \sum_{a \in A^-} -l'_a$ and thus $u_a = l'_a$ at line 16 for all $a \in A^-$ as required. The positive curtailment is performed at lines 9–12, resulting in $\sum_{a \in A^+} u_a = |\sum_{a \in A^-} l'_a| + C_r$ which equals $-\sum_{a \in A^-} u_a + C_r$ so that $\sum_{a \in A} u_a = C_r$.
- If $\sum_{a \in A} u'_a > C_r > \sum_{a \in A^+} u'_a + \sum_{a \in A^-} l'_a$ then $f^{\max} = C_r$ and we need negative curtailment to obtain $\sum_{a \in A} u_a = C_r$. Since in this case $-\sum_{a \in A^-} l'_a + C_r > \sum_{a \in A^+} u'_a$, at line 10 we have $\text{maxposflow} > \sum_{a \in A^+} u'_a$ and thus $u_a = u'_a$ at line 12 for all $a \in A^+$ as required. The negative curtailment is performed at lines 13–16, resulting in $\sum_{a \in A^-} u_a = -(\sum_{a \in A^+} u'_a - C_r)$ which equals $-\sum_{a \in A^+} u_a + C_r$.
- If $C_r \geq \sum_{a \in A} u'_a$ then $f^{\max} = \sum_{a \in A} u'_a$ and we need negative curtailment to obtain $\sum_{a \in A} u_a = \sum_{a \in A} u'_a$. Since $-\sum_{a \in A^-} u'_a + C_r \geq \sum_{a \in A^+} u'_a$ and thus $-\sum_{a \in A^-} l'_a + C_r \geq \sum_{a \in A^+} u'_a$, at line 10 we have $\text{maxposflow} \geq \sum_{a \in A^+} u'_a$ and thus $u_a = u'_a$ at line 12 for all $a \in A^+$ as required. Moreover, since in this case $\sum_{a \in A^+} u'_a - C_r \leq -\sum_{a \in A^-} u'_a$, at line 14 we have $\text{maxnegflow} \leq \sum_{a \in A^-} -u'_a$ and thus maximal negative curtailment $u_a = u'_a$ at line 16 for all $a \in A^-$.

Analogously for $Y_{f_{min}}^*$ and the l_a ($a \in A$) which it defines.

Since the upwards capacities of the vertices that are added when constructing a binary congestion tree from a congestion tree (Section 4.4.3) are infinite, the values l, u at those vertices do not change from the input values l', u' . Thus, Algorithm 6 can also be directly applied to (non-binary) congestion trees.

Algorithm 6: Mechanism (T)

Input: A (binary) congestion tree $T = (V, E, A)$
Output: The unique l_a and u_a for all $a \in A$

```

// Initialize with agent desires
1 for agents  $a$  at root  $r$  do
2   |  $l'_a \leftarrow \min(0, d_a)$ 
3   |  $u'_a \leftarrow \max(0, d_a)$ 

// Recursion on child vertices
4 for children  $c$  of root  $r$  do
5   | lowers, uppers  $\leftarrow$  Mechanism( $T_c$ )
6   | for agents  $a \in A_c$  do
7     | |  $l'_a \leftarrow$  lowers[ $a$ ]
8     | |  $u'_a \leftarrow$  uppers[ $a$ ]

// Now we have  $l'_a$  and  $u'_a \forall a \in A$ 

// Compute positive  $u_a$  values
9 maxposflow  $\leftarrow |\sum_{a \in A^-} l'_a| + C_r$ 
10  $w \leftarrow$  WaterLevel(maxposflow,  $A^+$ ,  $\{l'_a, u'_a\}_{a \in A^+}$ )
11 for  $a \in A^+$  do
12   |  $u_a \leftarrow \min(\max(w, l'_a), u'_a)$ 

// Compute negative  $u_a$  values
13 maxnegflow  $\leftarrow |\sum_{a \in A^+} u'_a| - C_r$ 
14  $w \leftarrow -$ WaterLevel(maxnegflow,  $A^-$ ,  $\{-u'_a, -l'_a\}_{a \in A^-}$ )
15 for  $a \in A^-$  do
16   |  $u_a \leftarrow \min(\max(w, l'_a), u'_a)$ 

// Compute negative  $l_a$  values
17 minnegflow  $\leftarrow |\sum_{a \in A^+} u'_a| + C_r$ 
18  $w \leftarrow -$ WaterLevel(minnegflow,  $A^-$ ,  $\{-u'_a, -l'_a\}_{a \in A^-}$ )
19 for  $a \in A^-$  do
20   |  $l_a \leftarrow \min(\max(w, l'_a), u'_a)$ 

// Compute positive  $l_a$  values
21 minposflow  $\leftarrow |\sum_{a \in A^-} l'_a| - C_r$ 
22  $w \leftarrow$  WaterLevel(minposflow,  $A^+$ ,  $\{l'_a, u'_a\}_{a \in A^+}$ )
23 for  $a \in A^+$  do
24   |  $l_a \leftarrow \min(\max(w, l'_a), u'_a)$ 

// Now we have  $l_a$  and  $u_a \forall a \in A$ 
25 return  $\{l_a\}_{a \in A}, \{u_a\}_{a \in A}$ 

```

4.6 Properties of the Mechanism

4.6.1 The Computational Complexity is Low

Algorithm 5 sorts the u and l values in $\mathcal{O}(m \cdot \log(m))$ time, where m is the number of agents in the input set A . It then enters a while loop which takes at most $m - 1$ iterations. Thus the computational complexity of Algorithm 5 is $\mathcal{O}(m \cdot \log(m))$.

Algorithm 6 visits each vertex in V exactly once, and at each vertex calls Algorithm 5 four times. Therefore, the worst case computational complexity of Algorithm 6 is $\mathcal{O}(n \cdot m \cdot \log(m))$. Here n is the number of vertices in V and m is the number of agents in A .

4.6.2 The Mechanism is Individually Rational

The mechanism is individually rational if it satisfies the following two conditions: it must not ask production from an agent that plans to consume and vice versa, and it must not ask an agent to exceed its planned consumption or production.

Proposition 4.1. Algorithm 6 computing Y_{opt}^* , as a mechanism, is individually rational.

Proof. Y_{opt}^* is leximin fair and thus a feasible allocation. Since feasible means desire compatible, the proposition follows. \square

4.6.3 The Mechanism is Incentive Compatible

Especially important in mechanisms that deal with fairness is incentive compatibility. Without this property, fairness is undermined as agents may ‘game the system’ by reporting strategically to try to improve their own situation relative to that of other agents.

Proposition 4.2. Algorithm 6 computing Y_{opt}^* , as a mechanism, is incentive compatible.

Proof. Let $T = (V, E, A)$ be a congestion tree containing an agent a with reported desire d_a and true desire d_a^* . Without loss of generality, assume that $a \in A^+$, i.e. $d_a^* > 0$.

If agent a reports $d_a = d_a^*$ and the mechanism returns $Y_{opt}^*(a) < d_a^*$, then there is exactly one subtree T_v such that for u_a, u'_a in Mechanism(T_v), we have $Y_{opt}^*(a) = u_a < u'_a$. This means that congestion occurred at v , i.e. $|\sum_{a \in A^-} l'_a| + C_v < \sum_{a \in A^+} u'_a$ at line 9.

From line 12 we know that for the water level w computed at line 10 we have $w < u'_a$. Since this w does not depend on $\sum_{a \in A^+} u'_a$ the only way a reduction $u_a < u'_a$ at v is avoided, is if $u'_a \leq Y_{opt}^*(a)$.

This would require agent a to report a desire d_a no larger than $Y_{opt}^*(a)$, which it would then be allocated. This shows that agent a cannot improve its situation by reporting anything other than its true desire d_a^* , i.e., being truthful is a weakly dominant strategy. \square

4.6.4 The Mechanism as a Distributed Algorithm

Algorithm 6 executes a recursion on subtrees. At each step, only information pertaining to its subtree is used. Thus computations pertaining to disjoint subtrees may be performed in parallel. As such, the algorithm may be implemented as a distributed algorithm.

4.7 Conclusions

In this chapter we investigated the hierarchical interactions of congestion in flow subnetworks. We naturally extended the notion of leximin fairness to capture local balancing of commodity flows. We then proved that for each feasible root flow interaction with the higher-level network, there exists a unique leximin fair congestion solution within the subnetwork. Furthermore, we fully parameterized these unique leximin fair congestion solutions. This lets us find the allocation of any agent in a unique leximin fair congestion solution with any possible root flow in just $\mathcal{O}(1)$ time. We then demonstrated the ease of applicability of these results by proposing an insightful algorithmic mechanism that computes such unique leximin fair congestion solutions. Finally, we proved that this mechanism is of low computational complexity, individually rational and incentive compatible, and implementable as a distributed algorithm, making it highly applicable in e.g. the energy domain.

5 Fairness, Welfare, and a Congestion Aftermarket

In the previous chapters we considered primarily fairness. In this chapter we introduce the comparison with welfare and study the combination of the two.

5.1 Introduction

With the changes in the electric grid brought about by the energy transition new challenges arise, many of which concern flexibility of users (Fang et al., 2012; Pepermans et al., 2005; Ramchurn et al., 2011; Vytelingum et al., 2010). A central challenge is that of grid congestion, which traditionally only occurs at the transmission system level, but now also occurs at the local distribution system level (Verzijlbergh et al., 2014). With the rapidly increasing penetration of distributed and renewable energy resources these problems can be expected to extend into the future (Clement-Nyns et al., 2010), especially since research indicates that neither grid expansion nor storage are solutions on the short to mid term (Härtel et al., 2016; Spiliotis et al., 2016).

A popular congestion management approach at the transmission system level is locational marginal pricing (LMP), which is part of the modern standardized market design (Ma et al., 2003). LMP determines the area price based on the marginal cost of producers that can actually deliver in that area within the transmission network capacity constraints. Nevertheless, since the price set at the transmission system level does not reflect local distribution grid constraints, local congestion management is required (Philipsen et al., 2016). LMP can fill this role as well, implementing what is essentially scarcity pricing, i.e. raising local prices until users sufficiently reduce their network usage. However, LMP is not budget balanced and does little to utilize local flexibility.

In this chapter we propose alternative mechanisms for local congestion management. We model capacity constrained distribution networks as rooted weighted trees whose edge weights represent line capacities and whose root represents a connection to a higher-level network, e.g. the transmission system. Such non-cyclic graphs accurately model the active radial structure of

most distribution networks (Sallam and Malik, 2018). The vertices are populated by agents that represent prosumers: users of the network that can both consume and/or produce power. Based on the price \hat{p} set by the higher-level market mechanism, these agents have desired prosumptions that may cause congestion in the local network. Our congestion management mechanisms follow the common approach to active power congestion management, which is curtailment (Bach Andersen et al., 2012; Hu et al., 2014; Rivera et al., 2015; Tonkoski et al., 2011; Verzijlbergh et al., 2014). Curtailment entails allocating to agents a prosumption that is a reduction of their desired prosumption in order to resolve congestion.

Mechanisms for local congestion management can be designed to focus on different aspects, the most straightforward being economic efficiency as expressed by utilitarian welfare. Mechanisms such as LMP aim to achieve this through price signals that result in allocation of capacity to the most competitive agents. However, such solutions do not consider another aspect that has become prominent in energy networks: fairness (Hekkelman and La Poutré, 2019). Fairness deserves explicit consideration because, when resolving grid congestion, the question arises how congestion solutions affect different users, not just individually but also relative to each other. Since energy has become a basic need for full participation in modern society, this issue of fairness between users must be addressed (European Parliament, 2016).

In this chapter we consider both approaches. On the one hand, congestion solutions should be fair to all users of the network. On the other hand, congestion implies a market limitation and has economic consequences for the users. Users may differ in how they value fairness versus welfare, which suggests that congestion solutions ideally are able to reconcile these differing viewpoints. However, since fair congestion solutions consider relative prosumptions instead of individual demand curves, it is unlikely that a fair curtailment solution also maximizes the total (utilitarian) welfare.

We first consider congestion solutions in the form of curtailment that maximize the agents' total welfare. We present an algorithmic mechanism that computes the maximal welfare curtailment allocation that is feasible for all agents at the price \hat{p} set by the higher-level market mechanism. Our algorithm considers the demand curves that the agents submit to the higher-level market mechanism to determine maximal welfare, and is purely a curtailment solution in the sense that it does not send price signals.

For fair congestion solutions we turn to the literature (Hekkelman and La Poutré, 2020; Megiddo, 1977). Rather than defining fairness ourselves we work with the abstract concept of agents' 'fair shares' that can be determined by any fair congestion solution of choice. With these generalised fair shares we then go on to propose an algorithmic mechanism that computes a hybrid

congestion solution that combines fairness and welfare. Our mechanism provides agents with the choice to either claim their fair share or to participate in welfare maximization, resulting in a hybrid congestion solution that focuses on the goals of the two sets of agents in parallel to each other.

Finally, we provide incentives to organize this hybrid solution. We present a pricing scheme for the welfare maximization aspect of our hybrid solution that lets us define a congestion aftermarket. Our congestion aftermarket operates on top of the fair shares and the higher-level price \hat{p} , letting agents trade portions of their allocated fair shares at locally marginal prices. This principle of local aftermarket prices bears similarities to LMP, with the two most important differences being that our aftermarket is budget balanced and does not expose agents that choose to claim their fair share to congestion prices. Our aftermarket incentivizes participation as we prove that it is individually rational for agents to participate, but does so without imposing economical consequences on agents that choose not to participate. Since agents are always free to claim their fair share at the higher-level price \hat{p} , we argue that our hybrid congestion management mechanism constitutes a fair mechanism.

The main contribution of this chapter is our congestion management solution, provided in algorithmic form, that both allows individual agents to claim their fair share and simultaneously maximizes welfare for the other agents in our novel congestion aftermarket.

5.1.1 Related Work

Integral network management frameworks, such as that proposed by (Kok et al., 2005), generally rely on congestion pricing like LMP. (Esmat and Usaola, 2016) and (Haque et al., 2017) address congestion with demand side management which also usually relies on price signals. (Parhizi et al., 2016, 2017) focuses on the distribution market operator to implement settlement or penalties for congestion management. (Khodabakhsh et al., 2019) consider fairness in energy rates with respect to how the burden of network overhead costs are divided over prosumers. (Philipsen et al., 2016) propose time-slot auctions for EV charging to resolve congestion and promote fairness among asymmetric parties. (Ardakanian et al., 2013) address fairness in EV charging through curtailment with a fair allocation that is found by optimizing a fair objective function under feasibility constraints. (Hekkelman and La Poutré, 2020) consider different notions of fairness for greedy local matching under capacity constraints. (Bashir et al., 2017) evaluate which factors in power networks, specifically PV, should be subject to fairness considerations. In comparison to the discussed work, we propose opt-in fairness alongside welfare maximization and emphasise self-contained local market resolution of congestion.

5.2 Preliminaries

In this chapter we consider congestion in tree graphs modelled on electricity distribution grids. Given a situation where the users of the network collectively cause congestion, we seek solutions for this congestion. In particular, we investigate both fair and welfare maximizing solutions. Fair solutions can be given as curtailment of users, while welfare solutions often involve a market mechanism.

We start by modelling commodity flow in congested tree networks populated by agents. Our model resembles a standard source-sink flow network with edge constraints, except that we have multiple sources and sinks in the form of agents. The agents act as either consumers or producers of the commodity based on a demand (or supply) curve that they submit to a higher-level (e.g. national) market. The price \hat{p} set by this market then determines each agent's desired prosumption. These desired prosumptions may cause congestion in the local network that we consider. We will compute allocations that resolve congestion given such an initial situation. We can have two objectives for these allocations: fair division of capacity over the agents, or welfare expressed by demand curves.

Let a **congestion tree** $T = (V, E, A)$ be a rooted weighted tree (V, E) with a set of agents A located at the vertices. Let a virtual edge at the root r represent the connection to a virtual parent that represents an external network. Let the edge weights be strictly positive, representing flow capacities, and denote the weight of an edge between vertex $v \in V$ and its parent as the **capacity** C_v .

Let supply and demand both be represented by **prosumption**; respectively by negative and positive prosumption. Let an agent's prosumption induce a matching flow from the external network to the vertex of that agent. Using this terminology, a maximal production is synonymous with a minimal (negative) flow.

Let each agent $a \in A$ submit a **demand curve** $d_a(p)$ which is a strictly monotonically decreasing function of price p . A positive demand $d_a(p) > 0$ indicates a consumer, while a negative demand $d_a(p) < 0$ indicates a producer. Given a price \hat{p} let $A^+ \subseteq A$ and $A^- \subseteq A$ denote the sets of consumers and producers, respectively.

The inverse relation of the demand curve expresses the marginal price or **marginal** $m_a(q)$ of an agent a , which is a strictly monotonically decreasing function of its prosumption q . Note that the marginal $m_a(q)$ of a production, i.e. $q < 0$, represents a marginal cost. The **welfare** $W_a(p, q)$ of an agent a is then given by its prosumption surplus $\int_0^q m_a(x) - p \, dx$. Note that for a producer this expression takes the equivalent form $\int_q^0 p - m_a(x) \, dx$, see also Figure 5.1.

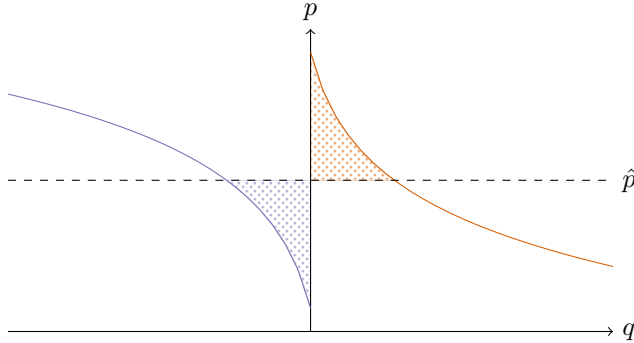


Figure 5.1 Marginal functions for a consumer (red, right) and a producer (blue, left) with their respective presumption surplus (hatched areas) corresponding to a price \hat{p} .

Since the flow in the entire network is induced by the presumptions of all individual agents, it will be convenient to work with **allocations** $Y : A \rightarrow \mathbb{R}$ that allocate a presumption to each agent. Because we consider trees, an allocation Y then defines the flow over each edge as the sum of presumptions in the subtree under that edge. This way, the **root flow** $F(Y)$ is given by the sum of allocated presumptions over all agents, i.e. $F(Y) = \sum_{a \in A} Y(a)$.

Congestion occurs when commodity flow induced by an allocation Y , e.g. of desired presumptions $Y(a) = d_a(\hat{p})$ ($a \in A$) given a price \hat{p} , results in the tree's edge capacities being exceeded. An allocation Y on a congestion tree $T = (V, E, A)$ is **congestion free** if for each vertex v the root flow $F_v(Y)$ of the subtree $T_v = (V_v, E_v, A_v)$ with root v does not exceed the capacity C_v , i.e. $|F_v(Y)| \leq C_v$. An allocation Y is **desire compatible** if each agent a is allocated a presumption between zero and its desire $d_a(\hat{p})$. An allocation Y is **feasible** if it is both congestion free and desire compatible. Finally, we say that a root flow f for a congestion tree T is feasible if there exists a feasible allocation Y on T with $F(Y) = f$.

We define a **congestion solution** Y as a feasible and Pareto allocation on T . We impose these restrictions since a non-feasible allocation does not satisfy the boundary conditions of the congestion problem, and if an allocation is not Pareto then it is possible to improve the allocation for an agent while still resolving congestion.

Congestion only occurs under certain circumstances, and when no congestion occurs a computed congestion solution will simply allocate the desired presumptions. In contrast to the uncertainty of congestion occurring, the agents

always participate in the higher-level market through their submitted demand curves. Thus, if congestion occurs relatively infrequently, agents' higher-level market participation dominates their participation in congestion management mechanisms. Therefore, if the higher-level market is incentive compatible we assume that agents truthfully submit their demand curves. Independent of this assumption curtailment solutions are **individually rational** since, at the same price \hat{p} , any prosumption less than the desired prosumption still has positive surplus.

5.3 Fair Share Congestion Management

When, for a price \hat{p} , congestion occurs in a congestion tree $T = (V, E, A)$ due to agents' desired prosumptions $d_a(\hat{p})$, there are often different congestion solutions Y possible. As stated before, an allocation Y must at least be feasible and Pareto to qualify as a congestion solution. However, an allocation Y may be required to have additional desirable properties. One such property is that of fairness, which may uniquely determine the allocation Y .

A fair congestion solution Y_{fair} allocates to each agent a 'fair share' of the available capacity. Previous work considers different notions of fairness for this setting (Hekkelman and La Poutré, 2020). Unique fair allocations are also provided for similar settings, such as by (Megiddo, 1977). For the present work it is only important that a fair allocation is unique, feasible, and Pareto. Going forward, when we refer to 'the fair allocation' or 'the fair shares' we will refer to the egalitarian fair allocation discussed in (Megiddo, 1977). However, any other notion of fairness that constitutes a unique feasible and Pareto allocation on T , such as those from (Hekkelman and La Poutré, 2020), can be substituted for egalitarian fairness.

5.4 Maximal Welfare Solutions

As opposed to fair congestion solutions we may aim to find congestion solutions in the form of allocations that maximize the welfare of the agents. The problem can be formulated as follows: given a congestion tree $T = (V, E, A)$ and a price \hat{p} , find a feasible allocation Y that maximizes the total welfare $\sum_{a \in A} W_a(\hat{p}, Y(a))$.

This problem decomposes into local division problems. Consider a congestion tree $T = (V, E, A)$ consisting of a single vertex r . When consumption congestion (i.e. flow is positive and exceeds capacity) occurs at r then the available capacity C_r has to be divided over the consumers $a \in A^+$. The available consumption capacity C^+ that is to be divided is given by the capacity C_r increased by the (maximal) production of the producers $a \in A^-$, which can

meet local demand, i.e. $C^+ = C_r + \sum_{a \in A^-} -d_a(\hat{p})$. This means that the aggregated consumption $\sum_{a \in A^+} Y(a)$ must equal C^+ for the allocation to be Pareto. Looking at the welfare of the consumers, we see that

$$\begin{aligned}
 \sum_{a \in A^+} W(\hat{p}, Y(a)) &= \sum_{a \in A^+} \left(\int_0^{Y(a)} m_a(x) - \hat{p} \, dx \right) \\
 &= \sum_{a \in A^+} \left(\int_0^{Y(a)} m_a(x) \, dx - \hat{p} \cdot Y(a) \right) \\
 &= \sum_{a \in A^+} \left(\int_0^{Y(a)} m_a(x) \, dx \right) - \hat{p} \cdot \sum_{a \in A^+} Y(a) \\
 &= \sum_{a \in A^+} \left(\int_0^{Y(a)} m_a(x) \, dx \right) - \hat{p} \cdot C^+.
 \end{aligned} \tag{5.1}$$

Equation (5.1) tells us that optimization of the total welfare among consumers does not depend on the price \hat{p} , because it represents a constant factor independent of the division. Hence, the local problem reduces to finding a feasible allocation Y that maximizes

$$\sum_{a \in A^+} \int_0^{Y(a)} m_a(x) \, dx \quad \text{s.t.} \quad \sum_{a \in A^+} Y(a) = C^+ \tag{5.2}$$

The above optimization problem is a standard market efficiency optimization problem, the solution to which is an allocation that minimizes the difference between agents' marginals $m_a(Y(a))$ at their allocated prosomptions. Indeed, when for two consumers a and b we have $m_a(Y(a)) > m_b(Y(b))$, that means consumer a can obtain more welfare from a marginal unit of consumption than consumer b does. Therefore, the allocation Y may be improved in total welfare by shifting an amount of consumption x from consumer b to consumer a such that $m_a(Y(a)+x) = m_b(Y(b)-x)$. A market obtains the unique solution by setting a single (scarcity) price p so that the demands sum up to the available consumption capacity, i.e. p such that $\sum_{a \in A^+} d_a(p) = C^+$. This is the principle under which locational marginal pricing (LMP) works; increasing the local price at r for all agents there reduces consumption while ensuring equal marginals among prosumers.

Our approach bears similarities to locational marginal pricing in that we compute allocations that exactly divide the available consumption capacity C^+ by setting a single marginal p among the consumers. However, we merely use this to compute the allocations and do not alter the actual price \hat{p} . As stated

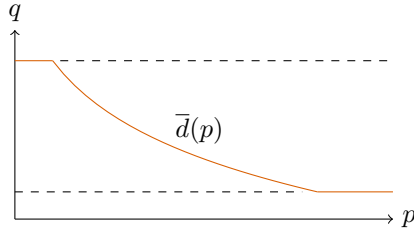


Figure 5.2 A bounded demand curve \bar{d} (red). The dashed lines indicate the maximum and minimum feasible demands.

earlier, the price \hat{p} does not affect the composition of the maximal welfare allocation.

In our setting of congestion trees with multiple nodes, we must account for possible congestion in subtrees. As such, we may not be able to find a feasible allocation Y for which all consumers' marginals are equal. To solve this problem we will compute feasible welfare maximizing allocations for both minimal and maximal local root flows recursively on subtrees of T . These minimal and maximal local flow allocations then define agent-specific bounds of feasibility for the subtrees of T . The initial agent-specific bounds prior to consideration of subtrees are 0 and the agent's desired presumption $d_a(\hat{p})$. This is, respectively, because consumers should not be allocated production and vice versa, and because the price \hat{p} makes it so that any units of presumption in excess of the agent's desired presumption $d_a(\hat{p})$ are negative welfare for that agent.

The agent-specific bounds are used to bound the agents' demand curves. See Figure 5.2 for a visual representation. The resulting bounded demand curves \bar{d}_a let us determine presumption levels based on marginals p , within the constraints of feasibility. Given a set of bounded demand curves \bar{d}_a for the consumers $a \in A^+$, we can look at the aggregated bounded demand curve to determine a marginal p such that $\sum_{a \in A^+} \bar{d}_a(p) = C^+$. This way we find the consumer allocation $Y(a) = \bar{d}_a(p)$ ($a \in A^+$) that feasibly divides the available consumption capacity C^+ over the consumers A^+ and maximizes their total welfare. In this case of consumption congestion the producers would be allocated maximal production within their individual bounds, i.e. $Y(a) = \bar{d}_a(\hat{p})$ ($a \in A^-$).

5.4.1 Maximal Welfare Congestion Algorithm

Consider a congestion tree $T = (V, E, A)$ and a price \hat{p} . The recursive algorithm presented here in Algorithm 7 computes for each agent $a \in A$

a unique lower bound l_a and upper bound u_a for which the following is an invariant of subtrees:

Theorem 5.1. *Given a congestion tree $T = (V, E, A)$ and a price \hat{p} , there exists a unique set of lower bounds $\{l_a\}_{a \in A}$ and upper bounds $\{u_a\}_{a \in A}$ such that for any feasible root flow f , a feasible allocation Y_f^\dagger uniquely maximizes the welfare among feasible allocations Y_f with root flow f if and only if Y_f^\dagger is given by either*

$$Y_f^\dagger(a) = \begin{cases} \min(\max(d_a(p), l_a), u_a) & a \in A^+ \\ l_a & a \in A^- \end{cases} \quad (5.3)$$

for all $a \in A$ or

$$Y_f^\dagger(a) = \begin{cases} u_a & a \in A^+ \\ \min(\max(d_a(p), l_a), u_a) & a \in A^- \end{cases} \quad (5.4)$$

for all $a \in A$, for some marginal p .

Proof. The proof is by induction. For the induction basis, consider a tree T with only one vertex r . For the agents $a \in A$ we define initial lower and upper bounds l'_a and u'_a that ensure desire compatibility, i.e. 0 and $d_a(p)$ ordered such that $l'_a \leq u'_a$. Consider a feasible root flow f . We can formulate three properties of the feasible allocation Y_f^\dagger that maximizes the welfare among feasible allocations with root flow f :

- Y_f^\dagger must be bounded by the bounds l'_a and u'_a to be feasible, i.e. $l'_a \leq Y_f^\dagger(a) \leq u'_a$ ($a \in A$),
- Consumers must be allocated maximal consumption or producers must be allocated maximal production (i.e. minimal flow). If not, then we can increase both consumption and production by some amount, which means Y_f^\dagger is not Pareto,
- Consumer and producer welfare is maximal when the difference between respective agents' marginals is minimal, as discussed earlier in Section 5.4.

From these three properties we can see that Y_f^\dagger is given by either

$$Y_f^\dagger(a) = \begin{cases} \min(\max(d_a(p), l'_a), u'_a) & a \in A^+ \\ l'_a & a \in A^- \end{cases} \quad (5.5)$$

for all $a \in A$ or

$$Y_f^\dagger(a) = \begin{cases} u'_a & a \in A^+ \\ \min(\max(d_a(p), l'_a), u'_a) & a \in A^- \end{cases} \quad (5.6)$$

for all $a \in A$, for some marginal p . Here, the first property imposes the minimum of l_a and maximum of u_a on all agents $a \in A$, the second property requires allocating either l_a to all producers $a \in A^-$ or u_a to all consumers $a \in A^+$, and the third property leads to a single marginal p across all other agents.

Now consider $Y_{f_{\max}}^\dagger$ and $Y_{f_{\min}}^\dagger$ for the maximal and minimal feasible root flows f_{\max} and f_{\min} . By Equations (5.5) and (5.6) for all agents, if $f \leq f'$, then $Y_f^\dagger(a) \leq Y_{f'}^\dagger(a)$ ($a \in A$). Thus, for every feasible root flow f we have $Y_{f_{\min}}^\dagger(a) \leq Y_f^\dagger(a) \leq Y_{f_{\max}}^\dagger(a)$ ($a \in A$). This means we can take $l_a = Y_{f_{\min}}^\dagger(a)$ ($a \in A$) and $u_a = Y_{f_{\max}}^\dagger(a)$ ($a \in A$) to obtain the unique bounds described by Theorem 5.1.

For the induction step, assume that the theorem holds for all subtrees T_c of T with c a child of the root vertex r . Thus for all agents a not at the root r we have lower and upper bounds l'_a and u'_a by the induction hypothesis. For agents a at the root r we again define initial lower and upper bounds l'_a and u'_a that ensure desire compatibility, i.e. 0 and $d_a(p)$ ordered such that $l'_a \leq u'_a$.

From here we follow the same argumentation as for the induction basis. Since an allocation that feasibly maximizes the welfare on T must also maximize the welfare on each subtree for that subtree's root flow, a maximal welfare allocation on T must be bounded by the bounds l'_a and u'_a . So again, for any feasible root flow f , Y_f^\dagger is given by either Equation (5.5) or Equation (5.6) for some marginal p (minimize differences across subtrees). The unique bounds are again obtained by taking $l_a = Y_{f_{\min}}^\dagger(a)$ ($a \in A$) and $u_a = Y_{f_{\max}}^\dagger(a)$ ($a \in A$). \square

Corollary 5.1. *For $\{l_a\}_{a \in A}$ and $\{u_a\}_{a \in A}$ as in Theorem 5.1,*

$$Y_{wel}(a) = \begin{cases} u_a & a \in A^+ \\ l_a & a \in A^- \end{cases} \quad (5.7)$$

is the unique feasible allocation that maximizes the total welfare among all feasible allocations on T .

Proof. Since Equation (5.7) is of the form of Equation (5.3) for some marginal p (and of the form of Equation (5.4) for some other marginal p), Y_{wel} is feasible. In addition, Y_{wel} uniquely maximizes the prosumption of consumers within the bounds (u_a) and uniquely minimizes the prosumption

of producers within the bounds (l_a). Thus Y_{wel} is the unique feasible maximal surplus allocation. \square

Our algorithm essentially implements the proof of Theorem 5.1.

Each recursion step of the algorithm considers a different subtree T_v of T . We will denote the unique lower and upper bounds defined by Theorem 5.1 for direct subtrees T_c of T_v , i.e. with c a child of vertex v , as l'_a and u'_a respectively for agents $a \in A_c$. In addition, for agents a located at the vertex v , we will similarly denote the initial bounds as l'_a and u'_a . This way, when considering T_v , we have l'_a, u'_a ($a \in A_v$) denoting the bounds that ensure desire compatibility for all agents and ensure congestion freeness on all strict subtrees of T_v . As such, the information contained in these bounds l'_a, u'_a lets us focus exclusively on the root capacity C_v of T_v when computing the unique bounds l_a, u_a defined for T_v by Theorem 5.1.

Algorithm 7 presents the algorithmic mechanism that computes the unique bounds l_a, u_a for the congestion tree T , and with them, by Corollary 5.1, the feasible allocation that maximizes the total welfare on T . On Lines 1 to 8 we compute the bounds l'_a, u'_a for all agents $a \in A$; in Lines 1 to 3 for agents at the root by setting the initial bounds of 0 and their desired prosumption (as obtained from their demand curve and the price \hat{p}), and in Lines 4 to 8 for all other agents through recursion on subtrees.

On Lines 9 and 10 we introduce a help function for clarity. This function takes any value x and ‘bounds’ it by the summed l'_a and summed u'_a of a set of agents $B \subset A$. The result is that any value x bounded by this function can be computed as the sum (aggregate) of a set of values that are feasibly allocated to the agents $a \in B$. In essence, the bounding function lets us apply the information contained in the l'_a, u'_a in a simple and straightforward way.

On Lines 11 to 14 we aggregate the agents’ demand curves into an aggregated bounded demand and supply curve. The aggregated bounded demand curve indicates for each marginal p what the combined demand of the consumers is. Because of the bounds, these aggregated demands are guaranteed to not cause feasibility or congestion issues in strict subtrees of T . See also Figure 5.2 for a visual representation.

On Lines 15 to 22 we compute the bounds u_a ($a \in A$) that constitute the unique maximal welfare allocation $Y_{f_{\max}}^\dagger$ that feasibly maximizes the root flow. For the consumers, the aggregated demand can at most equal the capacity C_r plus the maximal production $-\sum_{a \in A^-} l'_a$, which means a maximal positive flow as seen on Line 15. If the aggregated demand does not exceed this, i.e. there is no consumption congestion, then the maximal consumption is simply $\sum_{a \in A^+} u'_a$. This last case is caught by the bounded function.

On Line 16, since the aggregated demand is a continuous decreasing (thus invertible) function that takes values between $\sum_{a \in A^+} l'_a$ and $\sum_{a \in A^+} u'_a$, we

Algorithm 7: MaxWelfare (T, \hat{p})

Input: A congestion tree $T = (V, E, A)$ and a price \hat{p}
Output: The unique l_a and u_a for all $a \in A$

```

// Initialize with agent desires
1 for agents a at root r do
2   |  $l'_a \leftarrow \min(0, d_a(\hat{p}))$ 
3   |  $u'_a \leftarrow \max(0, d_a(\hat{p}))$ 

// Recursion on child vertices
4 for children c of root r do
5   | lowers, uppers  $\leftarrow \text{maxwelfare}(T_c, \hat{p})$ 
6   | for agents a  $\in A_c$  do
7     | |  $l'_a \leftarrow \text{lowers}[a]$ 
8     | |  $u'_a \leftarrow \text{uppers}[a]$ 

// Now we have  $l'_a$  and  $u'_a \forall a \in A$ 
// Add a bounding function for clarity:
9 Function bounded( $x, B$ )
   | Input: A value  $x \in \mathbb{R}$  and a subset of agents  $B \subset A$ 
   | Output: The value closest to  $x$  between the combined lower and
   | upper bounds of the agent(s) in  $B$ 
10 | return  $\min(\max(x, \sum_{a \in B} l'_a), \sum_{a \in B} u'_a)$ 

// Aggregate bounded demand curves:
11 Function demand( $p$ )
12 | return  $\sum_{a \in A^+} \text{bounded}(d_a(p), a)$ 
13 Function supply( $p$ )
14 | return  $\sum_{a \in A^-} \text{bounded}(d_a(p), a)$ 

// Compute maximum flow values in two steps:
// Compute positive  $u_a$  values
15 maxposflow  $\leftarrow \text{bounded}(C_r - \sum_{a \in A^-} l'_a, A^+)$ 
16 Select marginal  $p$  s.t. demand( $p$ ) = maxposflow
17 for  $a \in A^+$  do
18 |  $u_a \leftarrow \text{bounded}(d_a(p), a)$ 

// Compute negative  $u_a$  values
19 maxnegflow  $\leftarrow \text{bounded}(C_r - \sum_{a \in A^+} u'_a, A^-)$ 
20 Select marginal  $p$  s.t. supply( $p$ ) = maxnegflow
21 for  $a \in A^-$  do
22 |  $u_a \leftarrow \text{bounded}(d_a(p), a)$ 

```

```

// Compute minimum flow values in two steps:
// Compute negative  $l_a$  values
23 minnegflow  $\leftarrow$  bounded( $-C_r - \sum_{a \in A^+} u'_a, A^-$ )
24 Select marginal  $p$  s.t. supply( $p$ ) = minnegflow
25 for  $a \in A^-$  do
26 |  $l_a \leftarrow$  bounded( $d_a(p), a$ )

// Compute positive  $l_a$  values
27 minposflow  $\leftarrow$  bounded( $-C_r - \sum_{a \in A^-} l'_a, A^+$ )
28 Select marginal  $p$  s.t. demand( $p$ ) = minposflow
29 for  $a \in A^+$  do
30 |  $l_a \leftarrow$  bounded( $d_a(p), a$ )

// Now we have  $l_a$  and  $u_a \forall a \in A$ 
31 return  $\{l_a\}_{a \in A}, \{u_a\}_{a \in A}$ 

```

can select a marginal p that corresponds to the determined maximum positive flow. Then on Lines 17 and 18 we compute the upper bounds u_a for individual consumers $a \in A^+$ by breaking down the aggregated demand at the selected marginal p .

For the producers, the aggregated supply must at least match the maximal consumption $\sum_{a \in A^+} u'_a$ minus the capacity C_r , which means a minimal negative flow as seen on Line 19. Again, the bounded function on Line 19 catches the case where there is no production congestion.

On Lines 20 to 22 we select the marginal p that corresponds to the minimal supply (i.e. maximal negative flow) and subsequently compute u_a for all producers, analogous to Lines 16 to 18.

Analogously to the bounds u_a ($a \in A$) before, on Lines 23 to 30 we compute bounds l_a ($a \in A$) that constitute the unique maximal welfare allocation $Y_{f_{\min}}^{\dagger}$ that feasibly minimizes the root flow.

With the output of Algorithm 7, the maximal welfare allocation Y_{wel} on T can now be found by taking for each agent $a \in A$ their most extreme bound, i.e. for each consumer $a \in A^+$ their upper bound u_a and for each producer $a \in A^-$ their lower bound l_a .

Theorem 5.2. *Given a congestion tree $T = (V, E, A)$ and a price \hat{p} , Algorithm 7 computes the unique feasible maximal welfare allocation Y_{wel} described by Corollary 5.1.*

Proof. We showed how Algorithm 7 computes the bounds l_a, u_a described by Theorem 5.1. The maximal welfare allocation Y_{wel} on T is then found by simply taking for each consumer $a \in A^+$ their upper bound u_a and for each producer $a \in A^-$ their lower bound l_a , as described by Corollary 5.1. \square

5.5 A Hybrid Solution With A Choice

In Sections 5.3 and 5.4 respectively, we introduced the fair allocation Y_{fair} and the maximal welfare allocation Y_{wel} , both of which are congestion solutions (i.e. feasible and Pareto). Being curtailment solutions, the price per unit that the agents trade their allocated prosumptions for is the higher-level market price \hat{p} .

Each agent carries a private preference for one of the two solutions. Such preferences may be principled or based on individual circumstance. Where the fair solution provides agents with prosumptions that are fair relative to each other at the cost of welfare, the maximal welfare solution rewards economical efficiency at the cost of some welfare among less efficient agents. This last situation can be observed with LMP as well, where scarcity pricing pushes some less competitive agents out of the market.

In this section we propose a way for the two solutions to exist in parallel, and for individual agents to choose in which one they want to participate. We do this by allowing agents to either "claim their fair share" or participate in a congestion aftermarket. More specifically, we first curtail agents with the fair allocation Y_{fair} . Then the subset $A^{fair} \subset A$ of agents that indicate that they want to participate in the fair solution are allocated as determined by Y_{fair} . Subsequently, we compute a feasible maximal welfare allocation for the remaining agents, with the prosumption already allocated to the agents in A^{fair} fixed. This gives us a hybrid allocation Y_{hyb} .

Algorithm 8 presents a modified version of Algorithm 7 that computes the hybrid solution Y_{hyb} . The modification is small: for the agents $a \in A^{fair}$ that choose to claim their fair share, we initialize both their lower and upper bounds at this fair share on Lines 2 to 4. For the other agents Algorithm 8 then proceeds identical to Algorithm 7.

Theorem 5.3. *Given a congestion tree $T = (V, E, A)$, a price \hat{p} , a fair congestion solution Y_{fair} , and a subset of agents A^{fair} , Algorithm 8 computes the unique feasible allocation Y_{hyb} that allocates the fair share $Y_{fair}(a)$ to agents $a \in A^{fair}$ and maximizes the total welfare of the agents $a \in A \setminus A^{fair}$ given the fair shares allocated to A^{fair} .*

Proof. Algorithm 8 initializes $l'_a = u'_a = Y_{fair}(a)$ for agents $a \in A^{fair}$. If $l'_a = u'_a$ for an agent a then $l'_a = l_a = u_a = u'_a$ because of the use of

Algorithm 8: Hybrid $(T, \hat{p}, Y_{fair}, A^{fair})$

Input: A congestion tree $T = (V, E, A)$, a price \hat{p} , a congestion solution Y_{fair} and a subset of agents A^{fair}

Output: Unique l_a and u_a for all $a \in A$

```

// Initialize with fair shares or agent desires
1 for agents a at root r do
2   if a ∈ Afair then
3     | l'_a ← Yfair(a)
4     | u'_a ← Yfair(a)
5   else
6     | l'_a ← min(0, d_a(ĥp))
7     | u'_a ← max(0, d_a(ĥp))

// Recursion on child vertices
8 for children c of root r do
9   lowers, uppers ← Hybrid(T_c, ĥp, Yfair, Afair)
10  for agents a ∈ A_c do
11    | l'_a ← lowers[a]
12    | u'_a ← uppers[a]

// Now we have l'_a and u'_a ∀ a ∈ A
13 From here proceeds identical to Algorithm 7

```

the bounded function when computing l_a and u_a . Therefore, the lower and upper bounds of such an agent will stay constant through recursive steps of the algorithm. As a result, $Y_{hyb}(a) = Y_{fair}(a)$ for $a \in A^{fair}$.

For agents $a \notin A^{fair}$, the bounds l_a and u_a are computed identically to Algorithm 7. Since the fair shares $Y_{fair}(a)$ claimed by agents $a \in A^{fair}$ are part of the feasible allocation Y_{fair} , we know that fixing these presumptions does not render it impossible to find a feasible allocation. In other words, for agents $a \notin A^{fair}$, Algorithm 8 can be regarded as Algorithm 7 on a congestion tree with its capacities adjusted for the fair shares $Y_{fair}(a)$ claimed by agents $a \in A^{fair}$. Thus, Y_{hyb} maximizes the total welfare among agents $a \notin A^{fair}$ given that $Y_{hyb}(a) = Y_{fair}(a)$ for $a \in A^{fair}$. \square

The computation of the hybrid solution Y_{hyb} from a fair allocation Y_{fair} gives rise to the **difference allocation** Y_{diff} :

$$Y_{diff}(a) = Y_{hyb}(a) - Y_{fair}(a) \quad a \in A. \quad (5.8)$$

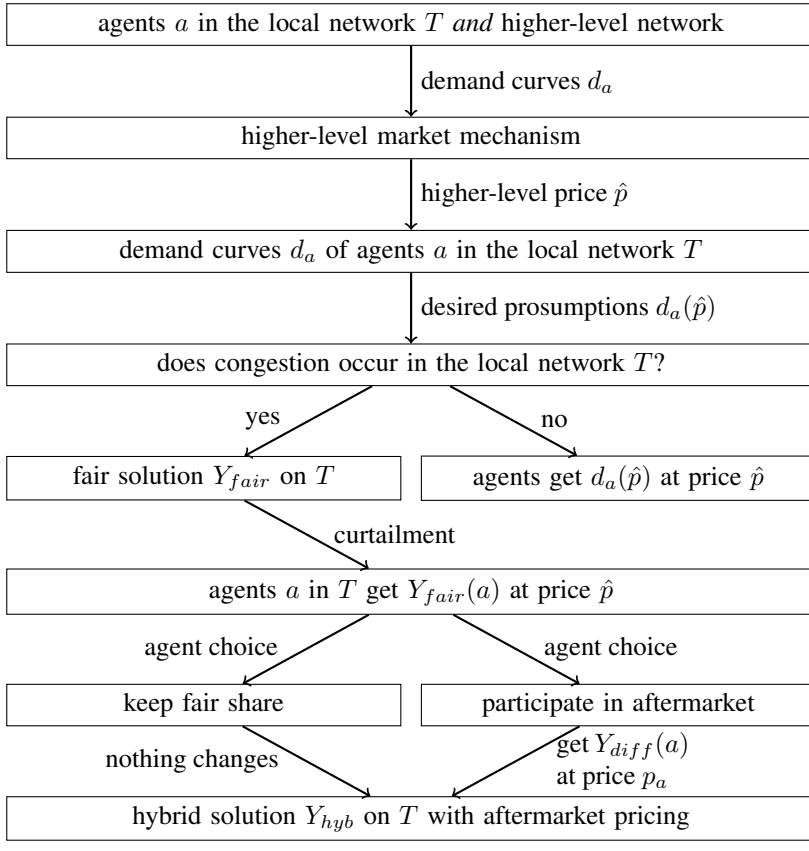


Figure 5.3 Visualization of the construction of a hybrid congestion solution: a fair curtailment solution with optional participation in a welfare-maximizing priced aftermarket.

This difference allocation indicates how the presumptions allocated to maximize welfare deviate from the fair shares. Accordingly, $Y_{diff}(a) = 0$ ($a \in A^{fair}$). Y_{diff} essentially tells us how units of presumption are transferred between agents relative to their fair shares, and thus will form the basis for the congestion aftermarket.

For some agents $a \in A$ the change $Y_{diff}(a)$ from their fair share $Y_{fair}(a)$ to $Y_{hyb}(a)$ moves them away from their desired presumption $d_a(\hat{p})$. In order to incentivize these agents to still choose to participate in welfare maximization we can implement a pricing scheme. In Section 5.6 we will lay out the specifics

of such a pricing scheme, including computation of agent-specific prices p_a . What is important is that we can interpret the difference allocation Y_{diff} as a congestion aftermarket in the following way. Each agent $a \in A$ gets to trade its allocated fair share $Y_{fair}(a)$ at the higher-level market price \hat{p} . Subsequently, agents $a \in A$ can choose to enter the competitive congestion aftermarket to trade an amount of prosumption equal to $Y_{diff}(a)$ at a certain price p_a (defined in Section 5.6). This may mean either selling a portion of their allocated fair share or purchasing additional prosumption from other agents participating in the aftermarket.

Note that when no congestion occurs, then both Y_{fair} and Y_{hyb} simply allocate the desired prosumptions $d_a(p)$ to all agents $a \in A$. This means that Y_{diff} is zero and hence that the congestion aftermarket does not exist. In other words, the congestion aftermarket only serves to let agents efficiently divide the available capacity among themselves when congestion occurs. Importantly, the aftermarket approach does not interfere in the higher-level market mechanism when no congestion occurs.

Figure 5.3 visualizes the different steps taken to arrive at the hybrid congestion solution Y_{hyb} with a congestion aftermarket. First the interaction between the agents' demand curves d_a , the higher-level market and its price \hat{p} , and the desired prosumptions $d_a(\hat{p})$ is indicated. We then turn our attention to the local network T where the desired prosumptions $d_a(\hat{p})$ may cause congestion. If no congestion occurs in T , the higher-level market mechanism can operate as intended in T . If, however, congestion does occur in T we must curtail the desired prosumptions $d_a(\hat{p})$, for which we use a fair congestion solution Y_{fair} . Now, based on agents' private preferences, agents $a \in A$ either claim their fair share $Y_{fair}(a)$ at price \hat{p} or participate in a welfare-maximizing congestion aftermarket. The resulting hybrid solution Y_{hyb} provides every agent $a \in A$ with their fair share $Y_{fair}(a)$ at the price \hat{p} , but on top of that provides those agents $a \notin A^{fair}$ that chose to participate in the aftermarket with a prosumption $Y_{diff}(a)$ traded at a certain (averaged) price p_a .

Figure 5.4 shows how first welfare is maximized by market clearing, then fair shares are allocated to resolve congestion, and finally welfare is increased again through the congestion aftermarket.

5.6 An Aftermarket Pricing Scheme

In Section 5.5 we discussed a congestion aftermarket where units of prosumption are traded according to a difference allocation Y_{diff} at individual prices p_a , without specifying these prices. In this section we present an explicit pricing scheme for the congestion aftermarket. This pricing scheme will ensure budget balance and individual rationality. Individual rationality means

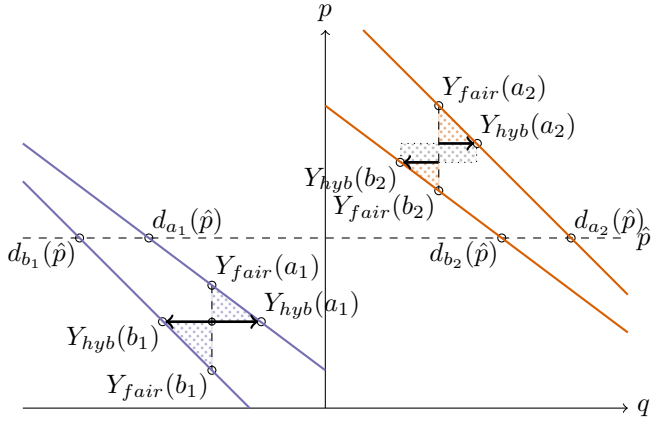


Figure 5.4 Two examples of transitions from higher-level market to fair curtailment to aftermarket, shown on marginal functions of two producers a_1, b_1 (blue, left) that share a connection and two consumers a_2, b_2 (red, right) that share a connection. Indicated values refer to q -coordinates. Desires of all four agents have equal (on a horizontal line) marginals, namely the price \hat{p} . Subsequently, the producers are allocated an equal (on a vertical line) fair share as indicated, as are the consumers. Finally, the producers return to equal (on a horizontal line) marginals in the aftermarket. The consumers, however, are constrained by some intermediate capacity in this example, so in the aftermarket their marginals only approach each other. The black arrows indicate the sign and magnitude of aftermarket trades Y_{diff} , and the hatched areas indicate the prosumers' aftermarket surplus. The efficiency gap between the consumers is shown in gray, with any price between the two marginals being acceptable for both parties.

that it will be economically beneficial for agents to participate in the aftermarket, independent of whether they buy or sell units of prosumption there.

Our goal for this pricing scheme is to put prices on trades of prosumption between agents that are given by the difference allocation Y_{diff} . By putting prices on bilateral trades rather than on purchases and sales individually, we will automatically satisfy budget balance. For individual rationality, the price for a trade should be such that both agents get positive surplus out of each unit transferred between them in the trade. If the post-aftermarket marginals $m_a(Y_{hyb}(a))$ and $m_b(Y_{hyb}(b))$ of an aftermarket buyer a and seller b are equal, then we can put a price equal to that marginal on the trade between

the two agents. In this case the aftermarket operates as regular unconstrained market-clearing.

However, due to the capacity constraints, not all trades attain maximal efficiency. For example, consider two consumers a and b that are curtailed by Y_{fair} because of a congestion they both contributed to. Now, consumer a would like to purchase a number of units x from consumer b because consumer b has a smaller marginal $m_b(Y_{fair}(b)) < m_a(Y_{fair}(a))$. Ideally, this number of units x is such that the two consumers' marginals become equal, i.e. $m_b(Y_{fair}(b) - x) = m_a(Y_{fair}(a) + x)$, and consumer a would pay a price per unit equal to that marginal. However, in our example we have this trade cause new congestion on one of the edges between the vertices where consumers a and b are located. Thus, a reduced number of units is traded. This means that after the trade, the marginal of consumer a will still be higher than that of consumer b , i.e. $m_b(Y_{fair}(b)) < m_b(Y_{hyb}(b)) < m_a(Y_{hyb}(a)) < m_a(Y_{fair}(a))$. For both consumers, any price between the two marginals results in a positive surplus. The gap indicates an economic inefficiency caused by the capacity constraints. This is shown in Figure 5.4 for agents a_2 and b_2 .

To deal with pricing under these capacity constraints we need to consider three aspects. First, we need to determine which trades can happen in which part of the network. Second, we need to determine what price to put on a trade when multiple prices yield a mutually beneficial trade. Third, we need to determine which agents trade with which agents and in what quantity.

For the **first aspect**, we look at what we call bottlenecks and congestion regions. In short, a congestion tree $T = (V, E, A)$ consists of alternating positive and negative congestion regions separated by congestion bottlenecks, and aftermarket trades are confined to these congestion regions. There may also be a single uncongested region containing the root r of T where all agents a are allocated their desired prosumption $d_a(\hat{p})$ by both Y_{hyb} and Y_{fair} , which therefore does not play a role in the aftermarket. See also Figure 5.5.

Definition 5.1. Given a congestion tree $T = (V, E, A)$, a price \hat{p} , and a congestion solution Y , we say that a vertex v is a positive **intermediate obstruction** if $F_v(Y) = C_v$ and there exists a consumer a with $Y(a) \neq d_a(\hat{p})$ at a vertex $u \in V_v$ such that for all vertices $z \neq v$ on the path from u to v , $F_z(Y) < C_z$. Analogously for a negative intermediate obstruction.

Definition 5.2. Given a congestion tree $T = (V, E, A)$, a price \hat{p} , and a congestion solution Y , we say that a positive intermediate obstruction v is a positive **bottleneck** if not the closest other intermediate obstruction on the root path of v is also a positive intermediate obstruction. Analogously for a negative bottleneck.

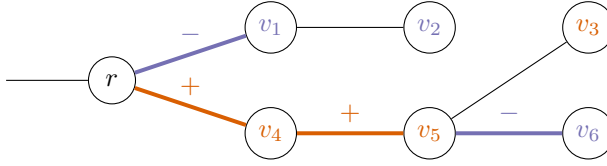


Figure 5.5 A congestion tree T with three congestion regions. Edges from v towards r marked with $+$ indicate $F_v(Y) = C_v$ (red) while $-$ indicates $F_v(Y) = -C_v$ (blue). Vertices v_1, v_4, v_5, v_6 are intermediate obstructions, but only v_1, v_4, v_6 are bottlenecks. $\{v_1, v_2\}$ and $\{v_6\}$ are negative congestion regions while $\{v_3, v_4, v_5\}$ is a positive congestion region. $\{r\}$ is uncongested.

Definition 5.3. Given a congestion tree $T = (V, E, A)$, a price \hat{p} , and a congestion solution Y , we say that a subgraph R^+ of T is a positive **congestion region** if it is a connected component of the forest obtained from T by removing all edges between bottlenecks and their parents, and it contains a positive bottleneck. Analogously for a negative congestion region R^- .

Lemma 5.1. Given a congestion tree $T = (V, E, A)$, a price \hat{p} , and a congestion solution Y , all producers in a positive congestion region R^+ are allocated their desired prosumption. i.e. $Y(a) = d_a(\hat{p})$ for all $a \in A^-$ in R^+ . Analogously for consumers.

Proof. Assume, without loss of generality, that there is a positive congestion region R^+ and a producer $a \in A^-$ with $Y(a) \neq d_a(\hat{p})$ located at a vertex u in R^+ . Consider, on the root path of u , the intermediate obstruction v closest to u . Such v exists and is in R^+ because a congestion region contains a bottleneck that is on the root path of every vertex in that congestion region. Since a positive congestion region, by definition, contains no negative intermediate obstructions, v is a positive intermediate obstruction. Hence there exists a vertex u' with a consumer $b \in A^+$ for which $Y(b) \neq d_b(\hat{p})$ such that, on the root path of u' , v is the vertex closest to u' that is at positive capacity (i.e. $F_v(Y) = C_v$). Therefore, since v is the closest intermediate obstruction to both u and u' on their root paths, there exists an $\epsilon > 0$ such that $Y(a)$ and $Y(b)$ can be feasibly decreased and increased, respectively, by ϵ without causing congestion on the root paths of u and u' , and thus anywhere in T . We conclude that the congestion solution Y is not Pareto, which is a contradiction. \square

Lemma 5.2. Given a congestion tree $T = (V, E, A)$ and a price \hat{p} , bottlenecks are independent of the congestion solution Y .

Proof. The proof is by induction. For the induction basis, consider without loss of generality a positive bottleneck v with no other bottlenecks in its subtree T_v . Since $F_v(Y) = C_v$, a congestion solution Y' for which v is not a bottleneck must have $F_v(Y') < C_v$. Since, by Lemma 5.1, $Y(a) = d_a(\hat{p})$ ($a \in A_v^-$), it must also be that $Y'(a) = d_a(\hat{p})$ ($a \in A_v^-$) for Y' to be Pareto. Thus it must be that $\sum_{a \in A_v^+} Y'(a) < \sum_{a \in A_v^+} Y(a)$. Now we consider two cases.

If there exist no other bottlenecks on the root path of v , then v connects to an uncongested region R where $Y(a) = d_a(\hat{p})$ (a in R). In this case, Y' cannot feasibly allocate more to any agent a in R than Y does, i.e. $Y'(a) = Y(a)$ (a in R). But Y' allocates less to consumers $a \in A_v^+$ than Y does, so Y' is not Pareto.

If there do exist other bottlenecks on the root path of v , then among these the closest to v must be a negative bottleneck u . u is in a negative congestion region R^- that v connects to. Since $F_u(Y) = -C_u$ and $F_v(Y') < F_v(Y)$, it must be that Y' allocates more to agents in R^- than Y does, i.e. $\sum_{a \in B} Y'(a) > \sum_{a \in B} Y(a)$ for B the set of agents in R^- (agents $a \in A_u \setminus (A_v \cup B)$ are in subtrees of positive bottlenecks in T_u). But because, by Lemma 5.1, $Y(a) = d_a(\hat{p})$ ($a \in B^+$), this means that $\sum_{a \in B^-} Y'(a) > \sum_{a \in B^-} Y(a)$. So in this case not only are the consumers $a \in A_v^+$ now allocated less (less consumption) by Y' than by Y , also the producers $a \in B^-$ in R^- are allocated more (less production) by Y' than by Y , which means that Y' is not Pareto.

For the induction step, assume that the lemma holds for all bottlenecks other than v in T_v . Hence we know that $F_u(Y') = F_u(Y)$ for any of these other bottlenecks u , for all congestion solutions Y' . Therefore the can follow the same argumentation as in the induction basis for the congestion region R^+ containing v (instead of T_v). \square

With these definitions and lemmas we formalized the fact that congestion trees consist of alternating positive and negative congestion regions invariant across congestion solutions, which are the only allocations that we are interested in. Differences between congestion solutions only appear among consumers in positive congestion regions and producers in negative congestion regions.

Lemma 5.3. The aftermarket trades Y_{diff} in the subtree T_v of a bottleneck v result in a net-zero root flow, i.e. $F_v(Y_{diff}) = 0$.

Proof. Since Y_{diff} is the difference between two congestion solutions, Y_{hyb} and Y_{fair} , this follows from Lemma 5.2. \square

Lemma 5.4. The aftermarket trades Y_{diff} of all producers in positive congestion regions R^+ are zero, i.e. $Y_{diff}(a) = 0$ for producers $a \in A^-$ in R^+ . Analogously for consumers.

Proof. This follows from Lemmas 5.1 and 5.3. □

Lemmas 5.3 and 5.4 show how aftermarket trades are confined to congestion regions, and that consumers only trade with other consumers (within the same positive congestion region) while producers only trade with other producers (within the same negative congestion region). Consequently, both consumers and producers only interact in one of two ways in the aftermarket: they either buy from, or sell to, agents with the same presumption sign (\pm).

For the **second aspect** (what price to put on a trade when multiple prices yield a mutually beneficial trade), we distinguish between two types of aftermarket participants which we call strainers and relievers. The strainers' aftermarket trades Y_{diff} are aligned with their presumption (i.e. $d_a(\hat{p}) > 0 < Y_{diff}(a)$ or $d_a(\hat{p}) < 0 > Y_{diff}(a)$ for a strainer a), moving them closer to their desired presumption. Because their presumption was curtailed by Y_{fair} to resolve congestion, these movements strain the line capacities. To relieve this strain on capacity, the relievers accept aftermarket trades that move them further away from their desired presumption (i.e. $d_a(\hat{p}) > 0 > Y_{diff}(a)$ or $d_a(\hat{p}) < 0 < Y_{diff}(a)$ for a reliever a). In a positive congestion region R^+ , the strainers are consumers $a \in A^+$ in R^+ that buy (i.e. $d_a(\hat{p}) > 0 < Y_{diff}(a)$) and the relievers are consumers $a \in A^+$ in R^+ that sell (i.e. $d_a(\hat{p}) > 0 > Y_{diff}(a)$), while in a negative congestion region R^- the strainers are producers $a \in A^-$ in R^- that sell (i.e. $d_a(\hat{p}) < 0 > Y_{diff}(a)$) and the relievers are producers $a \in A^-$ in R^- that buy (i.e. $d_a(\hat{p}) < 0 < Y_{diff}(a)$).

This leads us to choose the marginal of the strainer as the price for every trade, for two reasons. Firstly and objectively, in the aftermarket a strainer a attains, with $Y_{fair}(a) + Y_{diff}(a)$, at most its desired presumption $d_a(\hat{p})$ for which the marginal is \hat{p} for all agents, while a reliever a is bounded by a presumption of zero for which the marginal $m_a(0)$ exclusively depends on the submitted demand curve. This makes the strainer's marginal the most consistent choice of price since it always reflects the real marginal value in the aftermarket at its vertex and not a bounded value. Secondly and subjectively, we established that the role of the relievers is to enable additional presumption for strainers by essentially resolving some congestion. Since any price between the marginals of the reliever and strainer is acceptable for a trade, we may want to maximally reward the role that works to resolve congestion by setting the price at the strainer's marginal (instead of at for example the midpoint).

For the **third aspect** (which agents trade with which agents and in what quantity), we look at the matching of supply and demand in the aftermarket.

When a congestion region contains multiple strainers and relievers in the aftermarket, it is not yet clear which strainer trades with which reliever. Note that this matching merely labels indistinguishable units whose flows are already determined by Y_{diff} . Since strainers always pay a price equal to their marginals, the matching does not impact them. For relievers, however, it can matter which strainer they are said to trade with. To resolve the ambiguity, we simply proportionally match all strainers and relievers that could trade with each other. As a result, relievers' prices are set at the proportional average of the accessible strainers' marginals.

We implement this proportional matching recursively on subtrees, which, by Lemma 5.3, results in matching within congestion regions. This recursive approach where relievers and strainers in the aftermarket are maximally matched locally within each subtree before moving up the tree to larger subtrees (i.e. greedy local-first matching), is equivalent to our goal of matching every reliever with every strainer accessible to them. These are equivalent because if there is a difference in marginals between strainers in a congestion region, then there must be an intermediate obstruction between them that necessitates local matching or else Y_{hyb} would not maximize welfare. If there is no difference between the strainer's marginals then all matchings with relievers are equivalent.

Algorithm 9 presents our pricing scheme in accordance with the three discussed aspects. On Lines 1 to 8 we set for each agent $a \in A$ the price and quantity already traded in strict subtrees of T . For agents at the root r of T this must be zero, as seen on Lines 1 to 3. For the other agents we find their quantities already traded in strict subtrees T_c through subtree recursion on Lines 4 to 8.

With this information we can identify the quantity that each agent will still trade outside strict subtrees of T . On Line 11 we compute these untraded quantities for each agent $a \in A$ as the difference between the total eventual aftermarket participation $Y_{diff}(a)$ and the quantity q_a already traded in strict subtrees of T . The untraded quantities tell us the total remaining aftermarket demand and supply in the subtree T , computed on Lines 12 and 13. The maximal matching of supply and demand, given on Line 14, later determines the portions q'_a of the untraded quantities that will be traded within T among agents $a \in A$. Note that if the root r of T is a bottleneck then, by Lemma 5.3, $F_r(Y_{diff}) = 0$ meaning that all aftermarket supply and demand in T is matched and untraded quantities are 0.

On Line 15 we check if any trades can be made in T . If not, we simply return the prices p_a and quantities q_a traded in strict subtrees of T . Otherwise, we can assign new quantities traded in T . On Lines 16 to 19 we proportionally assign quantities of supply and demand. For the supply or demand side

Algorithm 9: PricingScheme (T, Y_{hyb}, Y_{fair})

Input: A congestion tree $T = (V, E, A)$ and the two congestion solutions Y_{hyb} and Y_{fair}

Output: For each $a \in A$, a price p_a and a quantity traded q_a

```

// Initialize prices and quantities
1 for agents a at root r do
2   |  $p_a \leftarrow 0$ 
3   |  $q_a \leftarrow 0$ 
// Recursion on child vertices
4 for children c of root r do
5   | prices, quantities  $\leftarrow$  PricingScheme( $T_c, Y_{hyb}, Y_{fair}$ )
6   | for agents  $a \in A_c$  do
7     | |  $p_a \leftarrow$  prices[ $a$ ]
8     | |  $q_a \leftarrow$  quantities[ $a$ ]
// Infer total aftermarket quantities
9 for agents  $a \in A$  do
10  |  $Y_{diff}(a) \leftarrow Y_{hyb}(a) - Y_{fair}(a)$ 
// Identify still untraded quantities (of  $Y_{diff}$ )
11 untraded  $\leftarrow \{Y_{diff}(a) - q_a\}_{a \in A}$ 
// Total still untraded demand and supply
12 demand  $\leftarrow \sum_{a \in A} \max(0, \text{untraded}[a])$ 
13 supply  $\leftarrow \sum_{a \in A} \min(0, \text{untraded}[a])$ 
// Maximally match untraded demand and supply
14 matching  $\leftarrow \min(\text{demand}, -\text{supply})$ 
15 if matching  $\neq 0$  then
// Proportionally assign new trade quantities
16 for  $a \in A$  with untraded[ $a$ ]  $> 0$  do
17   |  $q'_a \leftarrow \frac{\text{matching}}{\text{demand}} \cdot \text{untraded}[a]$ 
18 for  $a \in A$  with untraded[ $a$ ]  $< 0$  do
19   |  $q'_a \leftarrow \frac{\text{matching}}{-\text{supply}} \cdot \text{untraded}[a]$ 
// Identify strainers by trades aligned with
// presumption, and relievers as complement
20 strainers  $\leftarrow \{a \in A \mid \text{untraded}[a] \cdot Y_{hyb}(a) > 0\}$ 
21 relievers  $\leftarrow \{a \in A \mid \text{untraded}[a] \neq 0\} \setminus \text{strainers}$ 
// Proportional average of strainers' marginals
// sets relief price of new trades
22  $p_{rel} \leftarrow (\sum_{a \in \text{strainers}} q'_a \cdot m_a(Y_{hyb}(a))) / \sum_{a \in \text{strainers}} q'_a$ 

```

```

23   // Update price and quantity with new trades
23   for  $a \in \text{strainers}$  do
24     |    $p_a \leftarrow m_a(Y_{hyb}(a))$ 
25     |    $q_a \leftarrow q'_a + q_a$ 
26   for  $a \in \text{relievers}$  do
27     |    $p_a \leftarrow (p_{\text{rel}} \cdot q'_a + p_a \cdot q_a) / (q'_a + q_a)$ 
28     |    $q_a \leftarrow q'_a + q_a$ 
29 return  $\{p_a\}_{a \in A}, \{q_a\}_{a \in A}$ 

```

or both, the newly traded quantities q'_a equal the previously untraded quantities $\text{untraded}[a]$, i.e. we maximally match aftermarket supply and demand within T .

On Lines 20 and 21 we identify which agents are strainers and which are relievers. We can identify the strainers by the alignment of their aftermarket trades with their presumption, i.e. consumers with aftermarket demand and producers with aftermarket supply. The relievers then necessarily coincide with the complementary set of agents with nonzero untraded quantities in the aftermarket.

Having identified the strainers, we can compute a single price p_{rel} for the the reliever side of the newly matched trades in T . Since we want to proportionally match each reliever with each strainer, we compute this price as the proportional average of the strainers' prices p_a , i.e. their marginals $m_a(Y_{hyb}(a))$.

We are now ready to set a price p_a for each agent $a \in A$. For the strainers, on Lines 23 and 24, the price p_a always equals their marginal $m_a(Y_{hyb}(a))$ as discussed before. For the relievers, on Lines 26 and 27, the price of the new trades q'_a in T is set by the price p_{rel} . However, a reliever a may have already traded a quantity q_a in a strict subtree of T at a different price p_a . So, in order to get the price p_a for the total quantity traded in T , we take the weighted average of the two prices. On Lines 25 and 28 we add the newly traded quantities q'_a to the previously traded quantities q_a .

The output of the pricing scheme presented in Algorithm 9 on the full congestion tree $T = (V, E, A)$ on which Y_{fair} and Y_{hyb} are defined is, for each agent $a \in A$, a price p_a and a quantity traded q_a . This final quantity traded q_a equals the predetermined aftermarket trade $Y_{diff}(a)$. The price p_a is the price at which the agent a trades the quantity q_a on the congestion aftermarket, with the total amount paid given by $q_a \cdot p_a$.

For the hybrid allocation Y_{hyb} as a whole, each agent $a \in A$ receives its fair share of prosumption $Y_{fair}(a)$ at a market clearing price \hat{p} , and on top of that optionally trades $Y_{diff}(a)$ prosumption in the congestion aftermarket at a price p_a . The total payment for the final prosumption $Y_{hyb}(a)$ is given by $Y_{fair}(a) \cdot \hat{p} + Y_{diff}(a) \cdot p_a$.

Theorem 5.4. *Given the allocations Y_{hyb} , Y_{fair} , and their difference Y_{diff} on the congestion tree $T = (V, E, A)$, Algorithm 9 computes prices p_a for the aftermarket trades $Y_{diff}(a)$ ($a \in A$) such that*

$$\sum_{a \in A} Y_{diff}(a) \cdot p_a = 0 \quad (\text{budget balance})$$

$$\int_{Y_{fair}(a)}^{Y_{hyb}(a)} m_a(x) - p_a \, dx \geq 0 \quad (a \in A) \quad (\text{individual rationality}).$$

Proof. The aftermarket is easily seen to be budget balanced since we showed how Algorithm 9 computes the reliever prices by aggregating and exactly distributing the constrainer prices. We also showed how we chose constrainer prices to always equal their marginals in Y_{hyb} , and how reliever prices are consequently equal to or better than their marginals in Y_{hyb} . Therefore, each unit traded between agents in the aftermarket has positive prosumption surplus for both strainer and reliever, resulting in individual rationality. \square

Note that since aftermarket trades only occur between agents in the same congestion region, as stated by Lemma 5.3, Theorem 5.4 also holds for any subtree T_v of a bottleneck v .

5.7 Full Aftermarket Participation

Since it is individually rational for agents to participate in the congestion aftermarket, it may be that all agents choose to participate in it. In this case, since A^{fair} is empty, $Y_{hyb} = Y_{wel}$. Moreover, since the aftermarket is budget-balanced and Y_{wel} maximizes total welfare, Y_{hyb} with the aftermarket also maximizes the total welfare. The difference between the two is the distribution of the welfare among the agents. The aftermarket, relative to Y_{wel} , increases welfare for some agents while decreasing it for some others.

The deciding factor in how the welfare is redistributed among the agents between Y_{wel} and Y_{hyb} with the aftermarket is the fair allocation Y_{fair} . If $Y_{fair} = Y_{wel}$ then nothing is traded in the aftermarket and welfare is distributed identically among agents both with and without aftermarket, but any other Y_{fair} results in a redistribution of welfare relative to Y_{wel} . What is interesting to note is that the choice of fair shares through Y_{fair} thus translates to a

choice of welfare distribution among the agents, even if no agent claims their fair share and all agents participate in the aftermarket.

5.8 Numerical Example

Figure 5.6 and Table 5.1 show a simple numerical example illustrating the benefits of our hybrid solution over straightforward LMP.

Figure 5.6 shows marginal functions of three consumers a , b , and c with simple linear demand curves $d_a(p) = 8 - 2p$, $d_b(p) = 8 - p$, and $d_c(p) = 14 - 2p$. These consumers share a total capacity of 15, and the higher-level price \hat{p} is currently 1. The consumers' demands at $p = 1$ lead to a total demand of $6 + 7 + 12 = 25$, exceeding the capacity by 10. LMP sets the local congestion price at 3 to reduce total demand to $2 + 5 + 8 = 15$. An equitable solution allocates a total consumption of $5 + 5 + 5 = 15$. A subsequent aftermarket trade at $p = 4$, in which consumer a does not participate, allocates a total consumption of $5 + 4 + 6 = 15$.

Table 5.1 shows values corresponding to different congestion solutions. Allocation Y is the result of the higher-level market and is the preferred outcome, but causes congestion since its total consumption of 25 exceeds the capacity of 15. The LMP solution Y_{LMP} simply raises the price to find the efficient allocation of 2, 5, 8. However, the total payment now exceeds the cost of the total consumption in the higher-level market, causing a budget imbalance of 30 (i.e. the consumers pay the network). The simple equity allocation Y_{fair} curtails all three consumers to 5 but keeps the original price \hat{p} and is therefore budget balanced. However, Y_{fair} can be improved in terms of efficiency. Our aftermarket, in which agent a chooses not to participate, allows an extra trade between agents b and c at a price $p = 4$ which increases the efficiency while maintaining budget balance (i.e. the consumers pay each other). Finally, Y'_{hyb} shows full aftermarket participation (at a price $p = 3$). As we can see, this results in the maximally efficient allocation 2, 5, 8 that was also found by LMP. However, the aftermarket has avoided congestion pricing and drastically reduced consumer's costs. The difference in surplus between full aftermarket participation Y'_{hyb} and Y_{LMP} is exactly the 30 congestion overpayment. Notice that for Y'_{hyb} agent a bought 5 units at $\hat{p} = 1$ each and sold 3 units at $p = 3$ each in the aftermarket, netting an *income* of 4. A different notion of fairness, e.g. proportional, would affect aftermarket payments, but Y'_{hyb} would always allocate 2, 5, 8.

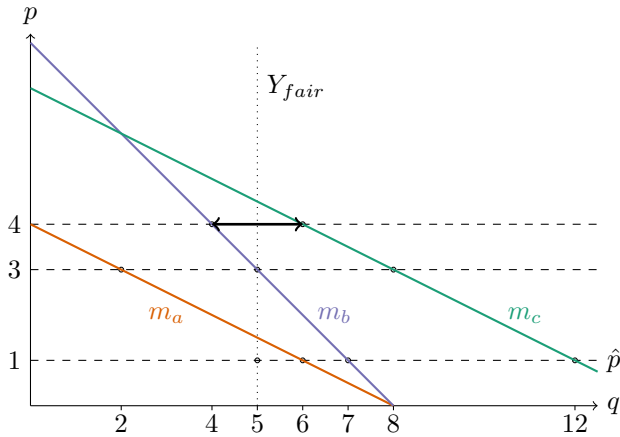


Figure 5.6 An example congestion problem, showing marginal functions for three consumers a (red, left), b (blue, middle), and c (green, right).

	Y	Y_{LMP}	Y_{fair}	Y_{hyb}	Y'_{hyb}
price p	1	3	1	1 & 4	1 & 3
consumption _{a}	6	2	5	5	2
consumption _{b}	7	5	5	4	5
consumption _{c}	12	8	5	6	8
consumption_{total}	25	15	15	15	15
payment _{a}	6	6	5	5	-4
payment _{b}	7	15	5	1	5
payment _{c}	12	24	5	9	14
payment_{total}	25	45	15	15	15
surplus _{a}	9	1	8.75	8.75	11
surplus _{b}	24.5	12.5	22.5	23	22.5
surplus _{c}	36	16	23.75	24	26
surplus_{total}	69.5	29.5	55	55.75	59.5

Table 5.1 Values corresponding to the example congestion problem from Figure 5.6.

5.9 Conclusions

In this chapter we considered congestion management in systems modelled after electricity distribution networks. In particular, we considered the welfare expressed by agents' demand curves in relation to quantities allocated by congestion-resolving curtailment mechanisms. We differentiated between congestion solutions focused on fairness and those focused on maximizing welfare.

We presented an algorithmic mechanism to find such welfare-maximizing congestion solutions for tree networks populated by both consumers and producers. These model e.g. common prosumer-oriented low- and medium-voltage electricity networks. We then went on to propose a hybrid congestion solution that provides agents with the choice between fairness and welfare maximization. We argued that giving agents the choice to claim a fair share of the available capacity at the original higher-level market clearing price is sufficient to constitute a fair congestion management mechanism. In such a mechanism we can let agents that do not choose to claim their fair share engage in welfare-maximizing activity amongst themselves. We achieved this choice-based hybrid congestion solution by applying our welfare-maximizing mechanism after locking in the fair shares of agents that decided to claim their fair share.

We then went on to define the welfare-maximizing part of our hybrid congestion solution as a congestion aftermarket by presenting a pricing scheme for the changes relative to the fair solution. We showed that this pricing scheme makes participation in the aftermarket an individually rational choice, and defines in a budget-balanced aftermarket. As a result, in contrast to popular congestion management mechanisms such as locational marginal pricing (LMP) where scarcity prices generate income for the mechanism, our hybrid solution gives agents the option of receiving a fair share at a non-scarcity price while still incentivizing welfare maximization through participation in a budget-balanced internal market.

Our Theorems 5.3 and 5.4, supported by Algorithms 8 and 9 respectively, provide local prosumer networks of arbitrary size with a way of becoming autarkic in their congestion management by offering internally-defined fair shares in parallel with a completely internal congestion aftermarket that feasibly maximizes welfare.

6 Conclusions

With the ongoing energy transition a new, more flexible and distributed, energy system is emerging. The changes brought about by this transition are especially disruptive for distribution systems as found in the built environment. The role of these local grids is changing from an essentially passive one-way system for distributing energy to end-users, to complex autarkic-like prosumer-based grids with internal supply-demand matching.

With the increased complexity of energy systems in the built environment, in part due to the increased penetration of distributed energy resources at the local level, come challenges previously only faced at the transmission system level as well as new challenges altogether. The most significant of these new challenges is the adequate and but also *fair* matching of supply and demand between prosumers. A major constraint here is the limited capacity of energy systems in the built environment that were not designed for today's usage patterns.

In our research we set out to investigate the modern need for local congestion management at the distribution level. Since these systems are predominantly comprised of end-users taking the role of prosumers, the role of resolving congestion through flexibility falls more and more with these individual users. As a result, fairness among prosumers in congestion management becomes a leading concern for supply-demand matching mechanisms. Therefore, we sought to answer the following overarching research question:

In a constrained multi-agent power flow system, how can we define fairness and how can we subsequently design supply-demand matching mechanisms that manage congestion fairly?

6.1 Answering the Research Questions

We will try to answer our overarching research question by answering the four narrower research questions that we formulated in Chapter 1. For each question, we will highlight two aspects. First, we discuss how we arrived at relevant notions of fairness in the setting. Second, we discuss how we

were able to implement those notions of fairness in supply-demand matching mechanisms geared towards congestion management.

Which factors determine how agents compare themselves to other agents, and how can we duly define, measure, and compute fairness among them?

In Chapter 2 we turned to social comparison theory to gain insight into what utility functions of agents in our setting may look like when we explicitly incorporate notions of fairness. We discovered how the topological aspect of congestion may lead agents to compare themselves primarily to their local peers. Moreover, we found how the contrast between production and consumption separates the agents into two groups each with internal envy that are, however, supportive rather than envious towards the other group. Finally, we devised a way to incorporate notions of fairness that were found inherently in humans by research in behavioural economics, directly into the utility functions.

In this chapter we also designed a simple congestion management mechanism for arriving at an egalitarian allocation for radial networks populated by prosumers of only one type, i.e. only consumers or only producers. We proved that our mechanism's allocations maximize the Nash product which also implies Pareto efficiency. The radial structure found in most energy systems in the built environment allowed us to efficiently compute allocations using subtree recursion, leading to a low computational complexity. Individual rationality was a constraint of the mechanism, and we proved the mechanism to be incentive compatible as well.

How does local supply-demand matching affect congestion and how can we incorporate local balancing into notions of fairness?

In Chapter 3 we explored the implications of local balancing opportunities for fairness. We postulated that it may be considered fair if an agent can use local prosumption for local balancing, as opposed to local prosumption being used for balancing further away in the network and leaving the local agent with congestion issues. From this starting point we were able to formulate the concept of local, outer matching for radial networks. For division among local agents we considered three principal notions of fair division: proportional, egalitarian, and non-discriminatory. Moreover, we showed how the local, outer matching aspect of fairness is compatible with any notion of fair division. This broad compatibility exists because we incorporate fairness in two disjoint aspects: one aspect is the selection of agents for supply-demand matching (where we apply local, outer matching), and the other aspect is the division of power flow among selected producers and among selected consumers separately (where we apply one of three principal notions of fair division).

In this chapter we incorporated the interaction between producers and consumers by designing a mechanism that takes a subtree-first approach to supply-demand matching. Again we made use of subtree recursion, this time to realize local, outer matching with the added benefit of keeping computational complexity low. We designed this mechanism with a plug-and-play module for division, making it readily compatible with any notion of fair division. Individual rationality was a constraint of the mechanism again, but this time we proved that incentive compatibility depends on the choice of fair division.

How can we adapt fair congestion solutions within a subnetwork to its collective interaction with a higher-level network?

In Chapter 4, instead of considering local fairness, we considered global fairness. Local notions of fairness have the potential to prefer matching only the producers and consumers closest to each other, at the expense of prosumers further away in the same congested region. This distinction of distance within the region be viewed as irrelevant when all prosumers in the region are constrained by congestion at the same bottleneck. We can reduce emphasis on the arbitrary distances between prosumers within congested regions by requiring a global fairness criterion to hold. A global fairness criterion can be any criterion that does not explicitly include the network topology in its definition. To this end we took the very strict notion of max-min fairness and adapted it to supply-demand matching settings by defining the max-min property on absolute values of prosumption.

In this chapter we expanded our focus from the local network to the hierarchical interactions between different levels of the energy system. Since a radial subnetwork interacts with the higher-level network through its flow at the root connection, we designed a mechanism that computes an entire parameterized family of globally max-min fair allocations for all possible root flows. With this parameterization of fair allocations the local network can quickly decide on an allocation once the intended interaction with the higher-level network is decided. Same as the mechanism from Chapter 2, this mechanism is of low computational complexity and is both individually rational and incentive compatible.

How do fairness and welfare relate to each other in the presence of limitations imposed by congestion, and is there a way to reconcile them?

In Chapter 5, we investigated the trade-off between maximizing for fairness or for welfare. Congestion always introduces a market limitation, resulting in economical inefficiencies. To arrive at an optimal market solution under congestion, many approaches such as locational marginal pricing (LMP) introduce

scarcity pricing to reduce consumption. Such an approach, however, raises the price of all the units traded indiscriminately and results in budget-imbalance. To avoid such consequences and pursue fairness, we first allocate fair shares which then serve as a starting point for maximizing welfare through exchanges between agents. Importantly, this lets us allow agents to choose between either their fair share or welfare maximization. With our approach, agents' fair shares are protected from scarcity pricing while agents retain the option to be flexible and maximize welfare in a budget-balanced aftermarket. This results in a choice-based notion of fairness that can exist in conjunction with welfare maximization, combining the best of both worlds while avoiding detrimental cross-effects.

In this chapter we first designed a straightforward congestion management mechanism that maximizes the total welfare, or surplus, of the agents. We then designed a hybrid mechanism by fixing the fair shares of agents that chose this option and subsequently applying the welfare maximizing mechanism to the rest of the agents. The fair shares can be defined by any fair mechanism, and which agents then actually 'claim' their fair share is up to the preferences of the agents themselves. With this hybrid mechanism we were able to define a congestion aftermarket with a budget-balanced pricing scheme. We prove that this pricing scheme makes the mechanism individually rational. We thus showed that participation in the aftermarket is economically optimal for each agent, while agents also always retain the option to claim their fair share at its pre-congestion price.

6.1.1 Answering our overarching research question

With these answers we can now formulate an answer to our overarching research question:

In a constrained multi-agent power flow system, how can we define fairness and how can we subsequently design supply-demand matching mechanisms that manage congestion fairly?

We found that notions of fairness regarding congested commodity flow networks can either focus on local or global fairness, and that agents can have differing opinions on the two depending on how wide they draw the circle of their peers. Furthermore, we found that it is possible to combine notions of fairness with welfare maximization by letting individual agents decide which of the two is more important, and protecting their fair shares if so desired.

We were able to use the radial structure prevalent in energy systems in the built environment to design algorithmic mechanisms of consistently low computational complexity. Besides individual rationality, we showed that egalitarian notions of fairness also led to incentive compatible mechanisms, while

some other common notions of fairness did not. Finally, we showed that it is possible to combine fairness and welfare in choice-based hybrid mechanisms through the addition of a budget-balanced congestion aftermarket for which participation is voluntary.

6.2 Outlook

In this work we focused on smart energy systems in the built environment and their typically radial structures. Some low- and medium-voltage distribution networks, however, have a mesh structure. It would be interesting to see how our results can be expanded to those networks and general network topologies. Max-min fair allocations can be computed for general isolated networks, but our concepts such as local, outer matching, root flow parameterized allocations, and aftermarket pricing schemes do not immediately translate to the general setting.

Many of the notions of fairness that we considered are based on a version of egalitarian fairness. A feature of these notions of fairness is that all agents are considered equal. In reality, however, we may want to discern between different agents, for example between a single solar panel on a roof and a complete solar field, or between a residential house and a hospital. A straightforward way to implement such distinctions is by assigning weights to the agents. Many of our mechanisms can be directly extended with weights, since weights are essentially just a way to split agents into “unit agents”.

A potentially interesting direction for weights, however, is to change them over time. One could imagine gradually raising the weight of an agent that is repeatedly allocated a small share until it is allocated a sufficient share. This would introduce the interesting possibility to dynamically change agents’ shares based on shares allocated in the recent past. For example, an agent that never asks for a large share may be accommodated when it sporadically does ask for a larger share. Another approach would be to periodically have an auction for weights, allowing users that expect to cause a lot of congestion to compensate other users up front instead of being curtailed or suffering penalties.

In general, adding a time dimension to fair allocation problems is an interesting approach. Next to the dynamic allocations mentioned before, there are settings in which agents require consecutive allocations over certain periods of time. For example, it is detrimental for electric vehicle batteries to start and stop charging often. This raises a question of fairness when a new car arrives at a charging station with limited capacity: do we re-allocate charging capacity from an already almost-charged car or do we let it finish charging uninterrupted before serving the newcomer?

Another technical extension of our mechanisms that could be made is to incorporate line losses. Even though line losses are usually not a deciding factor in low-voltage networks, their impact can be noticed. This raises an interesting question regarding fair shares: does an agent's fair share concern the amount of power generated or passing through a certain bottleneck, or does it concern the amount of power that arrives? Both alternatives could be incorporated, by reasonable approximation, in most of our mechanisms by multiplying all prosumption values that remain unmatched in a subtree, with the loss ratio of the subtree's upward connecting edge. This would crudely account for the losses incurred by the locally unmatched prosumption as the power flows required for matching from the rest of the network are inflated to compensate these losses.

Finally, it would be interesting to see if the local and global aspects of fairness that we found can be reconciled into one multi-level integrated system. A goal could be to have a system of interconnected fair mechanisms for the whole energy system, with appropriate notions of fairness for each level but also with consideration for the impact on other levels.

A The SES-BE Program

The research in this dissertation contributed to the Smart Energy Systems in the Built Environment (SES-BE) research program. This appendix provides a brief overview of the program as a whole and in relation to our work.

A.1 Research Program Outline

With the energy transition, a new sustainable energy supply chain is emerging due to the increasing deployment of renewable energy sources (RES) such as wind and solar, together with new types of appliances such as electric vehicles, heat pumps and other domestic electric usage technologies. Even though buildings are a major consumer, accounting for up to one-third of total energy consumption, they can potentially offer a high degree of flexibility in the way energy is used, stored, and locally produced with their on-site distributed energy resources (DER). The challenge is to design, operate and organize the future energy system to fully exploit the flexibility potential of smart buildings and their environment, while maintaining a high standard of system reliability and quality of energy supply. (Kling, 2014)

The main scientific challenges of the SES-BE program are to identify and explore the synergies between the built environment and emerging smart grid technologies and applications. Concretely, the scientific developments that the program aims for are: from traditional conversion to new more efficient technology in buildings; from single commodity to a multi-commodity energy management; from passive to active participation of end users (prosumers); and from a centralized to a decentralized market based system operation. (Kling, 2014)

The smart energy systems in the built environment (SES-BE) research program (Kling, 2014) formulates the following research question:

How can the complex multi-layered, multi-actor energy system within the built environment improve its overall sustainability and efficiency while satisfying user requirements such as cost, reliability and performance?

A.2 Individual Research Projects

The SES-BE research program couples different disciplines (power systems, built environment, conversion technologies, automation, computational intelligence and multi-agent decision making) and different aspects (interactive energy management and control, markets and services, and policies and entrepreneurship) of so-called smart energy systems (Kling, 2014).

The program was designed around multiple interconnected research projects that operate in different layers, from physical infrastructure to digital infrastructure to products and services. This has resulted in a cohesive body of research on smart energy systems in the built environment.

- On the topic of energy conversion, storage and distribution technologies for smart buildings, research focused on local voltage control in smart grids (Zhang et al., 2017, 2018) and optimization of future microgrids (Bandyopadhyay et al., 2018, 2020).
- Research on interactive energy management systems and lifecycle performance design for energy infrastructures of local communities worked towards net-zero energy neighbourhoods (Shafiullah et al., 2018, 2019) and studied life cycle performance design for clusters of buildings (Walker et al., 2018, 2021).
- Data-driven monitoring, prediction and real-time control for the smart grid was employed to analyze reduced price volatility in Germany since the energy transition (Khoshrou and Pauwels, 2018; Khoshrou et al., 2019) and provide solutions to power-flow problems (Sereeter et al., 2017, 2019).
- A framework for demand-supply matching and ancillary service provision through distributed energy resources was developed by research on fair congestion management matching mechanisms (Hekkelman and La Poutré, 2019, 2020) and distributed energy resources in coupled local and central markets (Farrokhsersht et al., 2020, 2021), and by surveying the participation of distributed energy resources in balancing markets (Farrokhsersht et al., 2018).
- Finally, a modelling lab for smart grids, smart policies and smart entrepreneurship was developed through research on the design of energy transitions (Moncada et al., 2017; Nava Guerrero et al., 2019) and the effectiveness and profitability of aggregation (Okur et al., 2018b; Özge Okur et al., 2019), and by reviewing aggregator business models (Özge Okur et al., 2021).

The common theme across these projects is the need for **flexibility** in energy networks, playing a role at all levels of energy systems in the built environment: device, building, community, aggregator, and regional grid. (Kling, 2014)

Bibliography

- Rediet Abebe, Jon Kleinberg, and David C. Parkes. Fair division via social comparison. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, page 281–289, Richland, SC, 2017. International Foundation for Autonomous Agents and Multiagent Systems.
- Noga Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247 – 253, 1987. ISSN 0001-8708. doi: [https://doi.org/10.1016/0001-8708\(87\)90055-7](https://doi.org/10.1016/0001-8708(87)90055-7). URL <http://www.sciencedirect.com/science/article/pii/0001870887900557>.
- Omid Ardakanian, Catherine Rosenberg, and S. Keshav. Distributed control of electric vehicle charging. In *Proceedings of the Fourth International Conference on Future Energy Systems*, e-Energy '13, page 101–112, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450320528. doi: 10.1145/2487166.2487178. URL <https://doi.org/10.1145/2487166.2487178>.
- Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 454–464, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341325. doi: 10.1145/2897518.2897522. URL <https://doi.org/10.1145/2897518.2897522>.
- P. Bach Andersen, J. Hu, and K. Heussen. Coordination strategies for distribution grid congestion management in a multi-actor, multi-objective setting. In *2012 3rd IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*, pages 1–8, 2012.
- Soumya Bandyopadhyay, Laura Ramirez-Elizondo, and Pavol Bauer. Stability constrained gain optimization of droop controlled converters in dc nanogrids. In *2018 International Power Electronics Conference (IPEC-Niigata 2018 -ECCE Asia)*, pages 1426–1434, 2018. doi: 10.23919/IPEC.2018.8507814.
- Soumya Bandyopadhyay, Gautham Ram Chandra Mouli, Zian Qin, Laura Ramirez Elizondo, and Pavol Bauer. Techno-economical model based optimal sizing of pv-battery systems for microgrids. *IEEE Transactions on Sustainable Energy*, 11(3): 1657–1668, 2020. doi: 10.1109/TSTE.2019.2936129.
- Salvador Barbarà and Matthew Jackson. Maximin, leximin, and the protective criterion: Characterizations and comparisons. *Journal of Economic Theory*, 46(1):34–44, 1988. ISSN 0022-0531. doi: [https://doi.org/10.1016/0022-0531\(88\)90148-2](https://doi.org/10.1016/0022-0531(88)90148-2). URL <https://www.sciencedirect.com/science/article/pii/0022053188901482>.
- Noman Bashir, David Irwin, Prashant Shenoy, and Jay Taneja. Enforcing fair grid energy access for controllable distributed solar capacity. In *Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments*, BuildSys '17, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450355445. doi: 10.1145/3137133.3137147. URL <https://doi.org/10.1145/3137133.3137147>.

- Marcus Berliant, William Thomson, and Karl Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992.
- Dimitri P Bertsekas, Robert G Gallager, and Pierre Humblet. *Data networks*, volume 2. Prentice-Hall International New Jersey, 1992.
- Steven J. Brams and Alan D. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995. doi: 10.1080/00029890.1995.11990526. URL <https://doi.org/10.1080/00029890.1995.11990526>.
- Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996. doi: 10.1017/CBO9780511598975.
- Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. *Multi-agent systems*, pages 213–283, 2012.
- Niv Buchbinder and Joseph Naor. Improved bounds for online routing and packing via a primal-dual approach. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 293–304, 2006. doi: 10.1109/FOCS.2006.39.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Trans. Econ. Comput.*, 7(3), September 2019. ISSN 2167-8375. doi: 10.1145/3355902. URL <https://doi.org/10.1145/3355902>.
- R. Ciavarella, M. Di Somma, G. Graditi, and M. Valenti. Congestion management in distribution grid networks through active power control of flexible distributed energy resources. In *2019 IEEE Milan PowerTech*, pages 1–6, June 2019. doi: 10.1109/PTC.2019.8810585.
- Kristien Clement-Nyns, Edwin Haesen, and Johan Driesen. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Transactions on Power Systems*, 25(1):371–380, 2010. doi: 10.1109/TPWRS.2009.2036481.
- Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 629–646, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450345279. doi: 10.1145/3033274.3085125. URL <https://doi.org/10.1145/3033274.3085125>.
- Hugh Dalton. The measurement of the inequality of incomes. *The Economic Journal*, 30(119):348–361, 1920. ISSN 00130133, 14680297. URL <http://www.jstor.org/stable/2223525>.
- Dominik Danner, Wolfgang Duschl, and Hermann de Meer. Fair charging service allocation for electric vehicles in the power distribution grid. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, e-Energy '19, page 406–408, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450366717. doi: 10.1145/3307772.3330168. URL <https://doi.org/10.1145/3307772.3330168>.
- L. E. Dubins and E. H. Spanier. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1P1):1–17, 1961. doi: 10.1080/00029890.1961.11989615. URL <https://doi.org/10.1080/00029890.1961.11989615>.
- Ayman Esmat and Julio Usaola. DSO congestion management using demand side flexibility. pages 197 – 201. IET, jun 2016. doi: 10.1049/cp.2016.0797.
- European Parliament. Clean energy for all europeans, nov 2016.

- Xi Fang, Satyajayant Misra, Guoliang Xue, and Dejun Yang. Smart grid - the new and improved power grid: A survey. 14:944–980, 01 2012.
- Mana Farrokhseresht, Nikolaos G. Paterakis, Madeleine Gibescu, Yaser Tohidi, and J. G. Slootweg. A survey on the participation of distributed energy resources in balancing markets. In *2018 15th International Conference on the European Energy Market (EEM)*, pages 1–5, 2018. doi: 10.1109/EEM.2018.8469948.
- Mana Farrokhseresht, Nikolaos G. Paterakis, Han Slootweg, and Madeleine Gibescu. Enabling market participation of distributed energy resources through a coupled market design. *IET Renewable Power Generation*, 14(4):539–550, February 2020. doi: 10.1049/iet-rpg.2019.0598. URL <https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/iet-rpg.2019.0598>.
- Mana Farrokhseresht, Han Slootweg, and Madeleine Gibescu. Day-ahead bidding strategies of a distribution market operator in a coupled local and central market. *Smart Energy*, 2:100021, 2021. ISSN 2666-9552. doi: <https://doi.org/10.1016/j.segy.2021.100021>. URL <https://www.sciencedirect.com/science/article/pii/S2666955221000216>.
- Ernst Fehr and Simon Gächter. Fairness and retaliation: The economics of reciprocity. *The Journal of Economic Perspectives*, 14(3):159–181, 2000. ISSN 08953309.
- Ernst Fehr and Klaus M. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868, 1999. ISSN 00335533, 15314650.
- Leon Festinger. A theory of social comparison processes. *Human Relations*, 7(2): 117–140, 1954. doi: 10.1177/001872675400700202. URL <https://doi.org/10.1177/001872675400700202>.
- Duncan K. Foley. Resource allocation and the public sector. *Yale economic essays*, 7 (1):45 – 98, January 1967.
- Francesca Fossati, Sahar Hoteit, Stefano Moretti, and Stefano Secci. Fair resource allocation in systems with complete information sharing. *IEEE/ACM Transactions on Networking*, 26(6):2801–2814, 2018. doi: 10.1109/TNET.2018.2878644.
- George Gamow and Marvin Stern. *Puzzle-Math*. Viking Press, New York, 1958.
- Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Nsdi*, volume 11, pages 24–24, 2011.
- Jonathan Goldman and Ariel D. Procaccia. Spliddit: Unleashing fair division algorithms. *SIGecom Exch.*, 13(2):41–46, jan 2015. doi: 10.1145/2728732.2728738. URL <https://doi.org/10.1145/2728732.2728738>.
- Avital Gutman and Noam Nisan. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '12*, page 719–728, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 0981738125.
- A.N.M.M. Haque, D.S. Shafullah, P.H. Nguyen, and F.W. Bliet. Real-time congestion management in active distribution network based on dynamic thermal overloading cost. In *2016 Power Systems Computation Conference (PSCC)*, pages 1–7, 2016. doi: 10.1109/PSCC.2016.7540985.

- A.N.M.M. Haque, P.H. Nguyen, F.W. Bliiek, and J.G. Slootweg. Demand response for real-time congestion management incorporating dynamic thermal overloading cost. *Sustainable Energy, Grids and Networks*, 10:65 – 74, 2017. ISSN 2352-4677. doi: <https://doi.org/10.1016/j.segan.2017.03.002>.
- Philipp Härtel, Michael Doering, Mareike Jentsch, Carsten Pape, Karsten Burges, and Rena Kuwahata. Cost assessment of storage options in a region with a high share of network congestions. *J. of Energy Storage*, 8:358 – 367, 2016. ISSN 2352-152X. doi: <https://doi.org/10.1016/j.est.2016.05.010>.
- Brinn Hekkelman and Han La Poutré. Fairness in smart grid congestion management. In *2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*, pages 1–5. IEEE, 2019. doi: 10.1109/ISGTEurope.2019.8905496. URL <https://doi.org/10.1109/ISGTEurope.2019.8905496>.
- Brinn Hekkelman and Han La Poutré. Fairness in power flow network congestion management with outer matching and principal notions of fair division. In *Proceedings of the Eleventh ACM International Conference on Future Energy Systems, e-Energy '20*, page 106–115, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380096. doi: 10.1145/3396851.3397701. URL <https://doi.org/10.1145/3396851.3397701>.
- Brinn Hekkelman and Han La Poutré. Fairness vs welfare: A hybrid congestion aftermarket. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, e-Energy '22*, page 93–104, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393973. doi: 10.1145/3538637.3538843. URL <https://doi.org/10.1145/3538637.3538843>.
- J. Hu, S. You, M. Lind, and J. Østergaard. Coordinated charging of electric vehicles for congestion prevention in the distribution grid. *IEEE Transactions on Smart Grid*, 5(2):703–711, 2014.
- Junjie Hu, Arshad Saleem, Shi You, Lars Nordström, Morten Lind, and Jacob Østergaard. A multi-agent system for distribution grid congestion management with electric vehicles. *Engineering Applications of Artificial Intelligence*, 38:45 – 58, 2015. ISSN 0952-1976. doi: <https://doi.org/10.1016/j.engappai.2014.10.017>. URL <http://www.sciencedirect.com/science/article/pii/S0952197614002577>.
- Matthew O. Jackson. Mechanism theory. *SSRN Electronic Journal*, 2014. doi: 10.2139/ssrn.2542983. URL <https://ssrn.com/abstract=2542983>.
- J. Jaffe. Bottleneck flow control. *IEEE Transactions on Communications*, 29(7):954–962, 1981. doi: 10.1109/TCOM.1981.1095081.
- Ali Khodabakhsh, Jimmy Horn, Evdokia Nikolova, and Emmanouil Pountourakis. Prosumer pricing, incentives and fairness. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems, e-Energy '19*, page 116–120, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450366717. doi: 10.1145/3307772.3328304. URL <https://doi.org/10.1145/3307772.3328304>.
- Abdolrahman Khoshrou and Eric J. Pauwels. Quantifying volatility reduction in german day-ahead spot market in the period 2006 through 2016. In *2018 IEEE Power Energy Society General Meeting (PESGM)*, pages 1–5, 2018. doi: 10.1109/PESGM.2018.8586020.

- Abdolrahman Khoshrou, André B. Dorsman, and Eric J. Pauwels. The evolution of electricity price on the german day-ahead market before and after the energy switch. *Renewable Energy*, 134:1–13, 2019. ISSN 0960-1481. doi: <https://doi.org/10.1016/j.renene.2018.10.101>. URL <https://www.sciencedirect.com/science/article/pii/S0960148118313077>.
- Daniel S Kirschen and Goran Strbac. *Fundamentals of power system economics*. John Wiley & Sons, 2nd edition, 2018.
- J. Kleinberg, Y. Rabani, and E. Tardos. Fairness in routing and load balancing. In *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039)*, pages 568–578, 1999. doi: 10.1109/SFFCS.1999.814631.
- Wil L. Kling. Smart energy systems in the built environment (SES-BE). Project Proposal, NWO project number 14182 P13-21, 2014.
- J. K. Kok, C. J. Warmer, and I. G. Kamphuis. Powermatcher: Multiagent control in the electricity infrastructure. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '05*, page 75–82, New York, NY, USA, 2005. Association for Computing Machinery. ISBN 1595930930. doi: 10.1145/1082473.1082807. URL <https://doi.org/10.1145/1082473.1082807>.
- David Kurokawa, Ariel D. Procaccia, and Nisarg Shah. Leximin allocations in the real world. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15*, page 345–362, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450334105. doi: 10.1145/2764468.2764490. URL <https://doi.org/10.1145/2764468.2764490>.
- Jérôme Lang and Jörg Rothe. Fair division of indivisible goods. In *Economics and Computation*, pages 493–550. Springer, 2016.
- Ruoyang Li, Qiuwei Wu, and Shmuel S. Oren. Distribution locational marginal pricing for optimal electric vehicle charging management. *IEEE Transactions on Power Systems*, 29(1):203–211, 07 2014.
- R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce, EC '04*, page 125–131, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581137710. doi: 10.1145/988772.988792. URL <https://doi.org/10.1145/988772.988792>.
- Xingwang Ma, D.I. Sun, and K.W. Cheung. Evolution toward standardized market design. *IEEE Transactions on Power Systems*, 18(2):460–469, 2003. doi: 10.1109/TPWRS.2003.810673.
- Nimrod Megiddo. Optimal flows in networks with multiple sources and sinks. *Mathematical Programming*, 7(1):97–107, 1974. doi: 10.1007/bf01585506. URL <https://lens.org/006-415-014-252-712>.
- Nimrod Megiddo. A good algorithm for lexicographically optimal flows in multi-terminal networks. *Bulletin of the American Mathematical Society*, 83(3):407–409, May 1977. doi: 10.1090/s0002-9904-1977-14298-5. URL <https://lens.org/100-404-996-870-897>.
- A. Merlin and H. Back. Search for a minimal-loss operating spanning tree configuration in an urban power distribution system. In *Proceedings of the Fifth Power Systems*

- Computer Conference*, PSCC 1975, pages 1–18, Cambridge, UK, September 1975. URL <https://psc-central.epfl.ch/repo/papers/1975/18.pdf>.
- Georgios Methenitis, Michael Kaisers, and Han La Poutré. Forecast-based mechanisms for demand response. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '19, page 1600–1608, Richland, SC, 2019a. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450363099.
- Georgios Methenitis, Michael Kaisers, and Han La Poutré. Degrees of rationality in agent-based retail markets. *Computational Economics*, 56(4):953–973, December 2019b. doi: 10.1007/s10614-019-09955-2. URL <https://link.springer.com/article/10.1007%2Fs10614-019-09955-2>.
- J. A. Moncada, E. H. Park Lee, G.D.C. Nava Guerrero, O. Okur, S.T. Chakraborty, and Z. Lukszo. Complex systems engineering: designing in sociotechnical systems for the energy transition. *EAI Endorsed Transactions on Energy Web*, 3(11), 7 2017. doi: 10.4108/eai.11-7-2017.152762.
- Hervé Moulin. *Fair division and collective welfare*. MIT press, 2003. ISBN 0262633116. doi: 10.7551/mitpress/2954.001.0001. URL <https://doi.org/10.7551/mitpress/2954.001.0001>.
- Christian Nabe, Marie-Louise Arlt, Michael Döring, Uwe Holzhammer, and Norman Gerhardt. Smart-market-design in deutschen verteilnetzen. Technical report, Ecofys and Fraunhofer IWES, requested by Agora Energiewende, 2017.
- John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1907266>.
- Graciela del Carmen Nava Guerrero, Gijsbert Korevaar, Helle Hvid Hansen, and Zofia Lukszo. Agent-based modeling of a thermal energy transition in the built environment. *Energies*, 12(5), 2019. ISSN 1996-1073. doi: 10.3390/en12050856. URL <https://www.mdpi.com/1996-1073/12/5/856>.
- Netbeheer Nederland. Capaciteitskaart elektriciteitsnet. 2022. URL <https://capaciteitskaart.netbeheernederland.nl/>.
- Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007. doi: 10.1017/CBO9780511800481.
- Włodzimierz Ogryczak, Hanan Luss, Michał Pióro, Dritan Nace, and Artur Tomaszewski. Fair optimization and networks: A survey. *Journal of Applied Mathematics*, 2014, 2014.
- Özge Okur, Robin Brouwer, Pieter Bots, and Fredrik Troost. Aggregated flexibility to support congestion management. In *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, pages 1–6, 2018a. doi: 10.1109/ISGTEurope.2018.8571791.
- Özge Okur, Nina Voulis, Petra Heijnen, and Zofia Lukszo. Critical analysis of the profitability of demand response for end-consumers and aggregators with flat-rate retail pricing. In *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, pages 1–6, 2018b. doi: 10.1109/ISGTEurope.2018.8571538.
- Sina Parhizi, Amin Khodaei, and Shay Bahramirad. Distribution market clearing and settlement. In *2016 IEEE Power and Energy Society General Meeting (PESGM)*,

- pages 1–5, 2016. doi: 10.1109/PESGM.2016.7741959.
- Sina Parhizi, Alireza Majzoobi, and Amin Khodaei. Net-zero settlement in distribution markets. In *2017 IEEE Power Energy Society General Meeting*, pages 1–5, 2017. doi: 10.1109/PESGM.2017.8274018.
- David C. Parkes, Ariel D. Procaccia, and Nisarg Shah. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12*, page 808–825, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450314152. doi: 10.1145/2229012.2229075. URL <https://doi.org/10.1145/2229012.2229075>.
- Elisha A. Pazner and David Schmeidler. Egalitarian Equivalent Allocations: A New Concept of Economic Equity*. *The Quarterly Journal of Economics*, 92(4):671–687, 11 1978. ISSN 0033-5533. doi: 10.2307/1883182. URL <https://doi.org/10.2307/1883182>.
- Q. Peng and S. H. Low. Distributed algorithm for optimal power flow on a radial network. In *53rd IEEE Conference on Decision and Control*, pages 167–172, Dec 2014. doi: 10.1109/CDC.2014.7039376.
- G. Pepermans, J. Driesen, D. Haeseldonckx, R. Belmans, and W. D’haeseleer. Distributed generation: definition, benefits and issues. *Energy Policy*, 33(6):787–798, 2005. ISSN 0301-4215. doi: <https://doi.org/10.1016/j.enpol.2003.10.004>. URL <https://www.sciencedirect.com/science/article/pii/S0301421503003069>.
- R.M. Philipsen, M.M. de Weerd, and L.J. de Vries. Auctions for congestion management in distribution grids. *13th International Conference on the European Energy Market*, pages 1–5, 2016. doi: 10.1109/EEM.2016.7521326.
- Arthur Cecil Pigou. *Wealth and welfare*. Macmillan and Company, limited, 1912.
- Michał Pióro and Deepankar Medhi. Chapter 8 - fair networks. In *Routing, Flow, and Capacity Design in Communication and Computer Networks*, The Morgan Kaufmann Series in Networking, pages 307–349. Morgan Kaufmann, San Francisco, 2004. doi: <https://doi.org/10.1016/B978-012557189-0/50011-1>. URL <https://www.sciencedirect.com/science/article/pii/B9780125571890500111>.
- Ariel D. Procaccia. Cake cutting: Not just child’s play. *Commun. ACM*, 56(7): 78–87, jul 2013. ISSN 0001-0782. doi: 10.1145/2483852.2483870. URL <https://doi.org/10.1145/2483852.2483870>.
- Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nick Jennings. Agent-based control for decentralised demand side management in the smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, AAMAS '11*, page 5–12, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 0982657153.
- Jose Rivera, Christoph Goebel, and Hans-Arno Jacobsen. A distributed anytime algorithm for real-time ev charging congestion control. In *Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems, e-Energy '15*, page 67–76, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450336093. doi: 10.1145/2768510.2768544. URL <https://doi.org/10.1145/2768510.2768544>.
- Jack M. Robertson and William A. Webb. *Cake Cutting Algorithms: Be Fair If You Can*. CRC Press, July 1998. ISBN 1568810768. URL <https://openlibrary.org/books/>

- OL693984M/Cake-cutting_algorithms.
- Abdelhay A. Sallam and Om P. Malik. *Electric Distribution Systems*. Wiley, 2nd edition, October 2018. ISBN 978-1-119-50932-5.
- Roland Saur, Neil Yorke-Smith, and Han La Poutré. Combined heat and power markets by double-sided auction mechanisms. In *2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*, pages 1–5, 2019. doi: 10.1109/ISGTEurope.2019.8905714. URL <https://ieeexplore.ieee.org/document/8905714>.
- Nripsuta Ani Saxena, Karen Huang, Evan DeFilippis, Goran Radanovic, David C. Parkes, and Yang Liu. How do fairness definitions fare? examining public attitudes towards algorithmic definitions of fairness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, AIES '19*, page 99–106, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450363242. doi: 10.1145/3306618.3314248. URL <https://doi.org/10.1145/3306618.3314248>.
- Pieter Schavemaker and Lou van der Sluis. *Electrical Power System Essentials*. Wiley, 2nd edition, August 2017. ISBN 978-1-118-80347-9.
- Erel Segal-Halevi and Balázs R Sziklai. Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory*, 68(2):363–401, 2019.
- Baljinnyam Sereeter, Kees Vuik, and Cees Witteveen. Newton power flow methods for unbalanced three-phase distribution networks. *Energies*, 10(10), 2017. ISSN 1996-1073. doi: 10.3390/en10101658. URL <https://www.mdpi.com/1996-1073/10/10/1658>.
- Baljinnyam Sereeter, Cornelis Vuik, Cornelis Witteveen, and Peter Palensky. Optimal power flow formulations and their impacts on the performance of solution methods. In *2019 IEEE Power Energy Society General Meeting (PESGM)*, pages 1–5, 2019. doi: 10.1109/PESGM40551.2019.8973585.
- D. S. Shafiqullah, T.H. Vo, and P. H. Nguyen. A heuristic search approach for sizing battery storage of a neighbourhood including nearly zero energy buildings (nzeb). In *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, pages 1–6, 2018. doi: 10.1109/ISGTEurope.2018.8571488.
- D. S. Shafiqullah, A.N.M.M. Haque, and P.H. Ngyuen. Effects of uncertainty characterization of energy demand of a neighborhood on stochastic day-ahead scheduling. In *2019 IEEE International Conference on Environment and Electrical Engineering and 2019 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I CPS Europe)*, pages 1–6, 2019. doi: 10.1109/EEEIC.2019.8783801.
- C. Shao, X. Wang, M. Shahidehpour, X. Wang, and B. Wang. Partial decomposition for distributed electric vehicle charging control considering electric power grid congestion. *IEEE Transactions on Smart Grid*, 8(1):75–83, Jan 2017. ISSN 1949-3061. doi: 10.1109/TSG.2016.2595494.
- Konstantinos Spiliotis, Ariana Isabel Ramos Gutierrez, and Ronnie Belmans. Demand flexibility versus physical network expansions in distribution grids. *Applied Energy*, 182:613–624, 2016. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2016.08.145>. URL <https://www.sciencedirect.com/science/article/pii/S0306261916312545>.
- Hugo Steinhaus. Report of the washington meeting, september 6-18, 1947. *Econometrica*, 16(1):101–104, 1948. ISSN 00129682, 14680262. URL <http://www.jstor.org/>

- stable/1914289.
- Walter Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980. doi: 10.1080/00029890.1980.11995109. URL <https://doi.org/10.1080/00029890.1980.11995109>.
- Reinaldo Tonkoski, Luiz A. C. Lopes, and Tarek H. M. El-Fouly. Coordinated active power curtailment of grid connected pv inverters for overvoltage prevention. *IEEE Transactions on Sustainable Energy*, 2(2):139–147, 2011. doi: 10.1109/TSTE.2010.2098483.
- Jurre van den Berg. Elektriciteitsnetwerk kan stroom uit lokale groene projecten niet aan. *De Volkskrant*, 2019. URL <https://www.volkskrant.nl/nieuws-achtergrond/elektriciteitsnetwerk-kan-stroom-uit-lokale-groene-projecten-niet-aan~b5a9fc1b/>.
- Stijn Vandael, Nelis Boucké, Tom Holvoet, Klaas De Craemer, and Geert Deconinck. Decentralized coordination of plug-in hybrid vehicles for imbalance reduction in a smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '11, page 803–810, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 0982657161.
- Hal R Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974. ISSN 0022-0531. doi: [https://doi.org/10.1016/0022-0531\(74\)90075-1](https://doi.org/10.1016/0022-0531(74)90075-1). URL <https://www.sciencedirect.com/science/article/pii/0022053174900751>.
- Remco A. Verzijlbergh, L.J. Vries, and Z Lukszo. Renewable energy sources and responsive demand. Do we need congestion management in the distribution grid? *IEEE Transactions on Power Systems*, 29(5):2119–2128, 09 2014.
- Perukrishnen Vytelingum, Sarvapali D. Ramchurn, Thomas D. Voice, Alex Rogers, and Nicholas R. Jennings. Trading agents for the smart electricity grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, page 897–904, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9780982657119.
- Shalika Walker, Timilehin Labeodan, Gert Boxem, Wim Maassen, and Wim Zeiler. An assessment methodology of sustainable energy transition scenarios for realizing energy neutral neighborhoods. *Applied Energy*, 228:2346–2360, 2018. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2018.06.149>. URL <https://www.sciencedirect.com/science/article/pii/S0306261918310195>.
- Shalika Walker, Vince Bergkamp, Dajuan Yang, T.A.J. van Goch, Katarina Katic, and Wim Zeiler. Improving energy self-sufficiency of a renovated residential neighborhood with heat pumps by analyzing smart meter data. *Energy*, 229:120711, 2021. ISSN 0360-5442. doi: <https://doi.org/10.1016/j.energy.2021.120711>. URL <https://www.sciencedirect.com/science/article/pii/S0360544221009592>.
- G. Weiss, editor. *Multiagent Systems*. MIT Press, 2nd edition, 2013.
- Zhenyu Xue, D. Z. Fang, Wen Sun, and Shiqiang Yuan. A novel approach for dynamic congestion management in open power markets. ICEET '09, pages 11–14, Washington, DC, USA, 2009. IEEE Computer Society. ISBN 978-0-7695-3819-8. doi: 10.1109/ICEET.2009.239. URL <http://dx.doi.org/10.1109/ICEET.2009.239>.

132 Bibliography

- Ya Zhang, Marcel A. M. Hendrix, Maurice G. L. Roes, and Jorge L. Duarte. Grid-connected converters with voltage support using only local measurements. In *2017 IEEE 8th International Symposium on Power Electronics for Distributed Generation Systems (PEDG)*, pages 1–6, 2017. doi: 10.1109/PEDG.2017.7972445.
- Ya Zhang, Maurice G. L. Roes, Marcel A. M. Hendrix, and Jorge L. Duarte. Voltage harmonic suppression by means of grid-connected converters using only local measurements. *Energies*, 11(10), 2018. ISSN 1996-1073. doi: 10.3390/en11102515. URL <https://www.mdpi.com/1996-1073/11/10/2515>.
- Özge Okur, Nina Voulis, Petra Heijnen, and Zofia Lukszo. Aggregator-mediated demand response: Minimizing imbalances caused by uncertainty of solar generation. *Applied Energy*, 247:426–437, 2019. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2019.04.035>. URL <https://www.sciencedirect.com/science/article/pii/S0306261919306683>.
- Özge Okur, Petra Heijnen, and Zofia Lukszo. Aggregator’s business models in residential and service sectors: A review of operational and financial aspects. *Renewable and Sustainable Energy Reviews*, 139:110702, 2021. ISSN 1364-0321. doi: <https://doi.org/10.1016/j.rser.2020.110702>. URL <https://www.sciencedirect.com/science/article/pii/S1364032120309837>.

SIKS Dissertation Series

- 2016 01 Syed Saiden Abbas (RUN), Recognition of Shapes by Humans and Machines
- 02 Michiel Christiaan Meulendijk (UU), Optimizing medication reviews through decision support: prescribing a better pill to swallow
- 03 Maya Sappelli (RUN), Knowledge Work in Context: User Centered Knowledge Worker Support
- 04 Laurens Rietveld (VU), Publishing and Consuming Linked Data
- 05 Evgeny Sherkhonov (UVA), Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers
- 06 Michel Wilson (TUD), Robust scheduling in an uncertain environment
- 07 Jeroen de Man (VU), Measuring and modeling negative emotions for virtual training
- 08 Matje van de Camp (TiU), A Link to the Past: Constructing Historical Social Networks from Unstructured Data
- 09 Archana Nottamkandath (VU), Trusting Crowdsourced Information on Cultural Artefacts
- 10 George Karafotias (VUA), Parameter Control for Evolutionary Algorithms
- 11 Anne Schuth (UVA), Search Engines that Learn from Their Users
- 12 Max Knobbout (UU), Logics for Modelling and Verifying Normative Multi-Agent Systems
- 13 Nana Baah Gyan (VU), The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach
- 14 Ravi Khadka (UU), Revisiting Legacy Software System Modernization
- 15 Steffen Michels (RUN), Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments
- 16 Guangliang Li (UVA), Socially Intelligent Autonomous Agents that Learn from Human Reward
- 17 Berend Weel (VU), Towards Embodied Evolution of Robot Organisms
- 18 Albert Meroño Peñuela (VU), Refining Statistical Data on the Web
- 19 Julia Efremova (Tu/e), Mining Social Structures from Genealogical Data
- 20 Daan Odijk (UVA), Context & Semantics in News & Web Search
- 21 Alejandro Moreno Celleri (UT), From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground
- 22 Grace Lewis (VU), Software Architecture Strategies for Cyber-Foraging Systems

134 Bibliography

- 23 Fei Cai (UVA), Query Auto Completion in Information Retrieval
- 24 Brend Wanders (UT), Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach
- 25 Julia Kiseleva (TU/e), Using Contextual Information to Understand Searching and Browsing Behavior
- 26 Dilhan Thilakarathne (VU), In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains
- 27 Wen Li (TUD), Understanding Geo-spatial Information on Social Media
- 28 Mingxin Zhang (TUD), Large-scale Agent-based Social Simulation - A study on epidemic prediction and control
- 29 Nicolas Höning (TUD), Peak reduction in decentralised electricity systems - Markets and prices for flexible planning
- 30 Ruud Mattheij (UvT), The Eyes Have It
- 31 Mohammad Khelghati (UT), Deep web content monitoring
- 32 Eelco Vriezolkolk (UT), Assessing Telecommunication Service Availability Risks for Crisis Organisations
- 33 Peter Bloem (UVA), Single Sample Statistics, exercises in learning from just one example
- 34 Dennis Schunselaar (TUE), Configurable Process Trees: Elicitation, Analysis, and Enactment
- 35 Zhaochun Ren (UVA), Monitoring Social Media: Summarization, Classification and Recommendation
- 36 Daphne Karreman (UT), Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies
- 37 Giovanni Sileno (UvA), Aligning Law and Action - a conceptual and computational inquiry
- 38 Andrea Minuto (UT), Materials that Matter - Smart Materials meet Art & Interaction Design
- 39 Merijn Bruijnes (UT), Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect
- 40 Christian Detweiler (TUD), Accounting for Values in Design
- 41 Thomas King (TUD), Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance
- 42 Spyros Martzoukos (UVA), Combinatorial and Compositional Aspects of Bilingual Aligned Corpora
- 43 Saskia Koldijk (RUN), Context-Aware Support for Stress Self-Management: From Theory to Practice
- 44 Thibault Sellam (UVA), Automatic Assistants for Database Exploration
- 45 Bram van de Laar (UT), Experiencing Brain-Computer Interface Control
- 46 Jorge Gallego Perez (UT), Robots to Make you Happy

- 47 Christina Weber (UL), Real-time foresight - Preparedness for dynamic innovation networks
- 48 Tanja Buttler (TUD), Collecting Lessons Learned
- 49 Gleb Polevoy (TUD), Participation and Interaction in Projects. A Game-Theoretic Analysis
- 50 Yan Wang (UVT), The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains
-
- 2017 01 Jan-Jaap Oerlemans (UL), Investigating Cybercrime
- 02 Sjoerd Timmer (UU), Designing and Understanding Forensic Bayesian Networks using Argumentation
- 03 Daniël Harold Telgen (UU), Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines
- 04 Mrunal Gawade (CWI), Multi-core Parallelism in a Column-store
- 05 Mahdieh Shadi (UVA), Collaboration Behavior
- 06 Damir Vandić (EUR), Intelligent Information Systems for Web Product Search
- 07 Roel Bertens (UU), Insight in Information: from Abstract to Anomaly
- 08 Rob Konijn (VU), Detecting Interesting Differences: Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery
- 09 Dong Nguyen (UT), Text as Social and Cultural Data: A Computational Perspective on Variation in Text
- 10 Robby van Delden (UT), (Steering) Interactive Play Behavior
- 11 Florian Kunneman (RUN), Modelling patterns of time and emotion in Twitter #anticipointment
- 12 Sander Leemans (TUE), Robust Process Mining with Guarantees
- 13 Gijs Huisman (UT), Social Touch Technology - Extending the reach of social touch through haptic technology
- 14 Shoshannah Tekofsky (UvT), You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior
- 15 Peter Berck (RUN), Memory-Based Text Correction
- 16 Aleksandr Chuklin (UVA), Understanding and Modeling Users of Modern Search Engines
- 17 Daniel Dimov (UL), Crowdsourced Online Dispute Resolution
- 18 Ridho Reinanda (UVA), Entity Associations for Search
- 19 Jeroen Vuurens (UT), Proximity of Terms, Texts and Semantic Vectors in Information Retrieval
- 20 Mohammadbashir Sedighi (TUD), Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility
- 21 Jeroen Linssen (UT), Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)

- 22 Sara Magliacane (VU), Logics for causal inference under uncertainty
- 23 David Graus (UVA), Entities of Interest — Discovery in Digital Traces
- 24 Chang Wang (TUD), Use of Affordances for Efficient Robot Learning
- 25 Veruska Zamborlini (VU), Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search
- 26 Merel Jung (UT), Socially intelligent robots that understand and respond to human touch
- 27 Michiel Joosse (UT), Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors
- 28 John Klein (VU), Architecture Practices for Complex Contexts
- 29 Adel Alhuraibi (UvT), From IT-Business Strategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT"
- 30 Wilma Latuny (UvT), The Power of Facial Expressions
- 31 Ben Ruijl (UL), Advances in computational methods for QFT calculations
- 32 Thaer Samar (RUN), Access to and Retrievability of Content in Web Archives
- 33 Brigit van Loggem (OU), Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity
- 34 Maren Scheffel (OU), The Evaluation Framework for Learning Analytics
- 35 Martine de Vos (VU), Interpreting natural science spreadsheets
- 36 Yuanhao Guo (UL), Shape Analysis for Phenotype Characterisation from High-throughput Imaging
- 37 Alejandro Montes Garcia (TUE), WiBAF: A Within Browser Adaptation Framework that Enables Control over Privacy
- 38 Alex Kayal (TUD), Normative Social Applications
- 39 Sara Ahmadi (RUN), Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR
- 40 Altaf Hussain Abro (VUA), Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems
- 41 Adnan Manzoor (VUA), Minding a Healthy Lifestyle: An Exploration of Mental Processes and a Smart Environment to Provide Support for a Healthy Lifestyle
- 42 Elena Sokolova (RUN), Causal discovery from mixed and missing data with applications on ADHD datasets
- 43 Maaïke de Boer (RUN), Semantic Mapping in Video Retrieval
- 44 Garm Lucassen (UU), Understanding User Stories - Computational Linguistics in Agile Requirements Engineering
- 45 Bas Testerink (UU), Decentralized Runtime Norm Enforcement
- 46 Jan Schneider (OU), Sensor-based Learning Support

- 47 Jie Yang (TUD), Crowd Knowledge Creation Acceleration
- 48 Angel Suarez (OU), Collaborative inquiry-based learning
-
- 2018 01 Han van der Aa (VUA), Comparing and Aligning Process Representations
- 02 Felix Mannhardt (TUE), Multi-perspective Process Mining
- 03 Steven Bosems (UT), Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction
- 04 Jordan Janeiro (TUD), Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks
- 05 Hugo Huurdeman (UVA), Supporting the Complex Dynamics of the Information Seeking Process
- 06 Dan Ionita (UT), Model-Driven Information Security Risk Assessment of Socio-Technical Systems
- 07 Jieting Luo (UU), A formal account of opportunism in multi-agent systems
- 08 Rick Smetsers (RUN), Advances in Model Learning for Software Systems
- 09 Xu Xie (TUD), Data Assimilation in Discrete Event Simulations
- 10 Julienka Mollee (VUA), Moving forward: supporting physical activity behavior change through intelligent technology
- 11 Mahdi Sargolzaei (UVA), Enabling Framework for Service-oriented Collaborative Networks
- 12 Xixi Lu (TUE), Using behavioral context in process mining
- 13 Seyed Amin Tabatabaei (VUA), Computing a Sustainable Future
- 14 Bart Joosten (UVT), Detecting Social Signals with Spatiotemporal Gabor Filters
- 15 Naser Davarzani (UM), Biomarker discovery in heart failure
- 16 Jaebok Kim (UT), Automatic recognition of engagement and emotion in a group of children
- 17 Jianpeng Zhang (TUE), On Graph Sample Clustering
- 18 Henriette Nakad (UL), De Notaris en Private Rechtspraak
- 19 Minh Duc Pham (VUA), Emergent relational schemas for RDF
- 20 Manxia Liu (RUN), Time and Bayesian Networks
- 21 Aad Slootmaker (OUN), EMERGO: a generic platform for authoring and playing scenario-based serious games
- 22 Eric Fernandes de Mello Araújo (VUA), Contagious: Modeling the Spread of Behaviours, Perceptions and Emotions in Social Networks
- 23 Kim Schouten (EUR), Semantics-driven Aspect-Based Sentiment Analysis
- 24 Jered Vroon (UT), Responsive Social Positioning Behaviour for Semi-Autonomous Telepresence Robots
- 25 Riste Gligorov (VUA), Serious Games in Audio-Visual Collections
- 26 Roelof Anne Jelle de Vries (UT), Theory-Based and Tailor-Made: Motivational Messages for Behavior Change Technology

- 27 Maikel Leemans (TUE), Hierarchical Process Mining for Scalable Software Analysis
 - 28 Christian Willemse (UT), Social Touch Technologies: How they feel and how they make you feel
 - 29 Yu Gu (UVT), Emotion Recognition from Mandarin Speech
 - 30 Wouter Beek, The "K" in "semantic web" stands for "knowledge": scaling semantics to the web
-
- 2019 01 Rob van Eijk (UL), Web privacy measurement in real-time bidding systems. A graph-based approach to RTB system classification
 - 02 Emmanuelle Beauxis Aussalet (CWI, UU), Statistics and Visualizations for Assessing Class Size Uncertainty
 - 03 Eduardo Gonzalez Lopez de Murillas (TUE), Process Mining on Databases: Extracting Event Data from Real Life Data Sources
 - 04 Ridho Rahmadi (RUN), Finding stable causal structures from clinical data
 - 05 Sebastiaan van Zelst (TUE), Process Mining with Streaming Data
 - 06 Chris Dijkshoorn (VU), Nichesourcing for Improving Access to Linked Cultural Heritage Datasets
 - 07 Soude Fazeli (TUD), Recommender Systems in Social Learning Platforms
 - 08 Frits de Nijs (TUD), Resource-constrained Multi-agent Markov Decision Processes
 - 09 Fahimeh Alizadeh Moghaddam (UVA), Self-adaptation for energy efficiency in software systems
 - 10 Qing Chuan Ye (EUR), Multi-objective Optimization Methods for Allocation and Prediction
 - 11 Yue Zhao (TUD), Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs
 - 12 Jacqueline Heinerma (VU), Better Together
 - 13 Guanliang Chen (TUD), MOOC Analytics: Learner Modeling and Content Generation
 - 14 Daniel Davis (TUD), Large-Scale Learning Analytics: Modeling Learner Behavior & Improving Learning Outcomes in Massive Open Online Courses
 - 15 Erwin Walraven (TUD), Planning under Uncertainty in Constrained and Partially Observable Environments
 - 16 Guangming Li (TUE), Process Mining based on Object-Centric Behavioral Constraint (OCBC) Models
 - 17 Ali Hurriyetoglu (RUN), Extracting actionable information from microtexts
 - 18 Gerard Wagenaar (UU), Artefacts in Agile Team Communication
 - 19 Vincent Koeman (TUD), Tools for Developing Cognitive Agents
 - 20 Chide Groenouwe (UU), Fostering technically augmented human collective intelligence

- 21 Cong Liu (TUE), Software Data Analytics: Architectural Model Discovery and Design Pattern Detection
 - 22 Martin van den Berg (VU), Improving IT Decisions with Enterprise Architecture
 - 23 Qin Liu (TUD), Intelligent Control Systems: Learning, Interpreting, Verification
 - 24 Anca Dumitrache (VU), Truth in Disagreement - Crowdsourcing Labeled Data for Natural Language Processing
 - 25 Emiel van Miltenburg (VU), Pragmatic factors in (automatic) image description
 - 26 Prince Singh (UT), An Integration Platform for Synchromodal Transport
 - 27 Alessandra Antonaci (OUN), The Gamification Design Process applied to (Massive) Open Online Courses
 - 28 Esther Kuindersma (UL), Cleared for take-off: Game-based learning to prepare airline pilots for critical situations
 - 29 Daniel Formolo (VU), Using virtual agents for simulation and training of social skills in safety-critical circumstances
 - 30 Vahid Yazdanpanah (UT), Multiagent Industrial Symbiosis Systems
 - 31 Milan Jelisavcic (VU), Alive and Kicking: Baby Steps in Robotics
 - 32 Chiara Sironi (UM), Monte-Carlo Tree Search for Artificial General Intelligence in Games
 - 33 Anil Yaman (TUE), Evolution of Biologically Inspired Learning in Artificial Neural Networks
 - 34 Negar Ahmadi (TUE), EEG Microstate and Functional Brain Network Features for Classification of Epilepsy and PNES
 - 35 Lisa Facey-Shaw (OUN), Gamification with digital badges in learning programming
 - 36 Kevin Ackermans (OUN), Designing Video-Enhanced Rubrics to Master Complex Skills
 - 37 Jian Fang (TUD), Database Acceleration on FPGAs
 - 38 Akos Kadar (OUN), Learning visually grounded and multilingual representations
-
- 2020 01 Armon Toubman (UL), Calculated Moves: Generating Air Combat Behaviour
 - 02 Marcos de Paula Bueno (UL), Unraveling Temporal Processes using Probabilistic Graphical Models
 - 03 Mostafa Deghani (UvA), Learning with Imperfect Supervision for Language Understanding
 - 04 Maarten van Gompel (RUN), Context as Linguistic Bridges
 - 05 Yulong Pei (TUE), On local and global structure mining

140 Bibliography

- 06 Preethu Rose Anish (UT), Stimulation Architectural Thinking during Requirements Elicitation - An Approach and Tool Support
- 07 Wim van der Vegt (OUN), Towards a software architecture for reusable game components
- 08 Ali Mirsoleimani (UL), Structured Parallel Programming for Monte Carlo Tree Search
- 09 Myriam Traub (UU), Measuring Tool Bias and Improving Data Quality for Digital Humanities Research
- 10 Alifah Syamsiyah (TUE), In-database Preprocessing for Process Mining
- 11 Sepideh Mesbah (TUD), Semantic-Enhanced Training Data Augmentation- Methods for Long-Tail Entity Recognition Models
- 12 Ward van Breda (VU), Predictive Modeling in E-Mental Health: Exploring Applicability in Personalised Depression Treatment
- 13 Marco Virgolin (CWI), Design and Application of Gene-pool Optimal Mixing Evolutionary Algorithms for Genetic Programming
- 14 Mark Raasveldt (CWI/UL), Integrating Analytics with Relational Databases
- 15 Konstantinos Georgiadis (OUN), Smart CAT: Machine Learning for Configurable Assessments in Serious Games
- 16 Ilona Wilmont (RUN), Cognitive Aspects of Conceptual Modelling
- 17 Daniele Di Mitri (OUN), The Multimodal Tutor: Adaptive Feedback from Multimodal Experiences
- 18 Georgios Methenitis (TUD), Agent Interactions & Mechanisms in Markets with Uncertainties: Electricity Markets in Renewable Energy Systems
- 19 Guido van Capelleveen (UT), Industrial Symbiosis Recommender Systems
- 20 Albert Hankel (VU), Embedding Green ICT Maturity in Organisations
- 21 Karine da Silva Miras de Araujo (VU), Where is the robot?: Life as it could be
- 22 Maryam Masoud Khamis (RUN), Understanding complex systems implementation through a modeling approach: the case of e-government in Zanzibar
- 23 Rianne Conijn (UT), The Keys to Writing: A writing analytics approach to studying writing processes using keystroke logging
- 24 Lenin da Nóbrega Medeiros (VUA/RUN), How are you feeling, human? Towards emotionally supportive chatbots
- 25 Xin Du (TUE), The Uncertainty in Exceptional Model Mining
- 26 Krzysztof Leszek Sadowski (UU), GAMBIT: Genetic Algorithm for Model-Based mixed-Integer optimization
- 27 Ekaterina Muravyeva (TUD), Personal data and informed consent in an educational context
- 28 Bibeg Limbu (TUD), Multimodal interaction for deliberate practice: Training complex skills with augmented reality

- 29 Ioan Gabriel Bucur (RUN), Being Bayesian about Causal Inference
- 30 Bob Zadok Blok (UL), Creatief, Creatieve, Creatiefst
- 31 Gongjin Lan (VU), Learning better – From Baby to Better
- 32 Jason Rhuggenaath (TUE), Revenue management in online markets: pricing and online advertising
- 33 Rick Gilsing (TUE), Supporting service-dominant business model evaluation in the context of business model innovation
- 34 Anna Bon (MU), Intervention or Collaboration? Redesigning Information and Communication Technologies for Development
- 35 Siamak Farshidi (UU), Multi-Criteria Decision-Making in Software Production
-
- 2021 01 Francisco Xavier Dos Santos Fonseca (TUD), Location-based Games for Social Interaction in Public Space
- 02 Rijk Mercur (TUD), Simulating Human Routines: Integrating Social Practice Theory in Agent-Based Models
- 03 Seyyed Hadi Hashemi (UVA), Modeling Users Interacting with Smart Devices
- 04 Ioana Jivet (OU), The Dashboard That Loved Me: Designing adaptive learning analytics for self-regulated learning
- 05 Davide Dell'Anna (UU), Data-Driven Supervision of Autonomous Systems
- 06 Daniel Davison (UT), "Hey robot, what do you think?" How children learn with a social robot
- 07 Armel Lefebvre (UU), Research data management for open science
- 08 Nardie Fanchamps (OU), The Influence of Sense-Reason-Act Programming on Computational Thinking
- 09 Cristina Zaga (UT), The Design of Robothings. Non-Anthropomorphic and Non-Verbal Robots to Promote Children's Collaboration Through Play
- 10 Quinten Meertens (UvA), Misclassification Bias in Statistical Learning
- 11 Anne van Rossum (UL), Nonparametric Bayesian Methods in Robotic Vision
- 12 Lei Pi (UL), External Knowledge Absorption in Chinese SMEs
- 13 Bob R. Schadenberg (UT), Robots for Autistic Children: Understanding and Facilitating Predictability for Engagement in Learning
- 14 Negin Samaeemofrad (UL), Business Incubators: The Impact of Their Support
- 15 Onat Ege Adali (TU/e), Transformation of Value Propositions into Resource Re-Configurations through the Business Services Paradigm
- 16 Esam A. H. Ghaleb (UM), Bimodal emotion recognition from audio-visual cues
- 17 Dario Doti (UM), Human Behavior Understanding from motion and bodily cues using deep neural networks

142 Bibliography

- 18 Remi Wieten (UU), Bridging the Gap Between Informal Sense-Making Tools and Formal Systems - Facilitating the Construction of Bayesian Networks and Argumentation Frameworks
 - 19 Roberto Verdecchia (VU), Architectural Technical Debt: Identification and Management
 - 20 Masoud Mansoury (TU/e), Understanding and Mitigating Multi-Sided Exposure Bias in Recommender Systems
 - 21 Pedro Thiago Timbó Holanda (CWI), Progressive Indexes
 - 22 Sihang Qiu (TUD), Conversational Crowdsourcing
 - 23 Hugo Manuel Proença (LIACS), Robust rules for prediction and description
 - 24 Kaijie Zhu (TUE), On Efficient Temporal Subgraph Query Processing
 - 25 Eoin Martino Grua (VUA), The Future of E-Health is Mobile: Combining AI and Self-Adaptation to Create Adaptive E-Health Mobile Applications
 - 26 Benno Kruit (CWI & VUA), Reading the Grid: Extending Knowledge Bases from Human-readable Tables
 - 27 Jelte van Waterschoot (UT), Personalized and Personal Conversations: Designing Agents Who Want to Connect With You
 - 28 Christoph Selig (UL), Understanding the Heterogeneity of Corporate Entrepreneurship Programs
-
- 2022 01 Judith van Stegeren (UT), Flavor text generation for role-playing video games
 - 02 Paulo da Costa (TU/e), Data-driven Prognostics and Logistics Optimisation: A Deep Learning Journey
 - 03 Ali el Hassouni (VUA), A Model A Day Keeps The Doctor Away: Reinforcement Learning For Personalized Healthcare
 - 04 Ünal Aksu (UU), A Cross-Organizational Process Mining Framework
 - 05 Shiwei Liu (TU/e), Sparse Neural Network Training with In-Time Over-Parameterization
 - 06 Reza Refaei Afshar (TU/e), Machine Learning for Ad Publishers in Real Time Bidding
 - 07 Sambit Praharaaj (OU), Measuring the Unmeasurable? Towards Automatic Co-located Collaboration Analytics
 - 08 Maikel L. van Eck (TU/e), Process Mining for Smart Product Design
 - 09 Oana Andreea Inel (VUA), Understanding Events: A Diversity-driven Human-Machine Approach
 - 10 Felipe Moraes Gomes (TUD), Examining the Effectiveness of Collaborative Search Engines
 - 11 Mirjam de Haas (UT), Staying engaged in child-robot interaction, a quantitative approach to studying preschoolers' engagement with robots and tasks during second-language tutoring
 - 12 Guanyi Chen (UU), Computational Generation of Chinese Noun Phrases

- 13 Xander Wilcke (VUA), Machine Learning on Multimodal Knowledge Graphs: Opportunities, Challenges, and Methods for Learning on Real-World Heterogeneous and Spatially-Oriented Knowledge
- 14 Michiel Overeem (UU), Evolution of Low-Code Platforms
- 15 Jelmer Jan Koorn (UU), Work in Process: Unearthing Meaning using Process Mining
- 16 Pieter Gijsbers (TU/e), Systems for AutoML Research
- 17 Laura van der Lubbe (VUA), Empowering vulnerable people with serious games and gamification
- 18 Paris Mavroustakos Blom (TiU), Player Affect Modelling and Video Game Personalisation
- 19 Bilge Yigit Ozkan (UU), Cybersecurity Maturity Assessment and Standardisation
- 20 Fakhra Jabeen (VUA), Dark Side of the Digital Media - Computational Analysis of Negative Human Behaviors on Social Media
- 21 Seethu Mariyam Christopher (UM), Intelligent Toys for Physical and Cognitive Assessments
- 22 Alexandra Sierra Rativa (TiU), Virtual Character Design and its potential to foster Empathy, Immersion, and Collaboration Skills in Video Games and Virtual Reality Simulations
- 23 Ilir Kola (TUD), Enabling Social Situation Awareness in Support Agents
- 24 Samaneh Heidari (UU), Agents with Social Norms and Values - A framework for agent based social simulations with social norms and personal values
- 25 Anna L.D. Latour (LU), Optimal decision-making under constraints and uncertainty
- 26 Anne Dirkson (LU), Knowledge Discovery from Patient Forums: Gaining novel medical insights from patient experiences
- 27 Christos Athanasiadis (UM), Emotion-aware cross-modal domain adaptation in video sequences
- 28 Onuralp Ulusoy (UU), Privacy in Collaborative Systems
- 29 Jan Kolkmeier (UT), From Head Transform to Mind Transplant: Social Interactions in Mixed Reality
- 30 Dean De Leo (CWI), Analysis of Dynamic Graphs on Sparse Arrays
- 31 Konstantinos Traganos (TU/e), Tackling Complexity in Smart Manufacturing with Advanced Manufacturing Process Management
- 32 Cezara Pastrav (UU), Social simulation for socio-ecological systems
- 33 Brinn Hekkelman (CWI/TUD), Fair Mechanisms for Smart Grid Congestion Management

Curriculum Vitae



Brinn Hekkelman took an interest in mathematics early on, joining the ‘wiskundekring’ (‘math circle’) organised by Wim Berkelmans at the VU during the first years of his high school. After graduating he enrolled in a double bachelor program mathematics and theoretical physics at the University of Amsterdam. Continuing this line, with a focus on algebra,

he enrolled in a master mathematical physics also at the UvA. During the master he spent one year first studying at UNICAMP in Brazil and then doing a Tesla minor on bridging the gap between science and society at the UvA. This set him on a course away from pure theory to more applied domains. With a PhD in computer science on the topics of fairness and mechanism design in power systems at the CWI and TU Delft, he was able to work on a complex practical problem from a theoretical background. After completing the PhD research he started as a researcher at the Netherlands Bureau for Economic Policy Analysis, working with fairness and mechanism design in new contexts.

Electricity users who share a grid connection encounter congestion when they cause too much power to flow through the connection. To resolve congestion, the available grid capacity must be divided over the users.

What is a fair way to divide the capacity? What if the network has a more complex structure, with different users being affected by congestion at different places in the network? And what if users not only consume but may also produce electricity, enabling matching between them to mitigate congestion?

In this dissertation we adapt and define notions of fairness for these complex settings. We also design algorithmic mechanisms that compute corresponding congestion solutions. We prove that our mechanisms are fair and that they provide the right incentives for users.