A Bandwidth Market for Traffic Engineering in Telecommunication Networks

J.J. Combrink

THESIS PRESENTED IN PARTIAL FULFILMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE
AT THE UNIVERSITY OF STELLENBOSCH

Prof. A.E. Krzesinski

April 2004

Declaration

I, the undersigned, hereby declare that the work contained in this thesis is my own original work and that I have not previously in its entirety or in part submitted it at any university for a degree.

Abstract

Traffic engineering determines the bandwidth allocation required to meet the traffic loads in a network. Similarly an economic market determines the resource allocation required to meet the demand for resources. The term *bandwidth market* denotes traffic engineering methods that use economic market methodology to determine the bandwidth allocation required to meet the traffic loads. A bandwidth market is an attractive traffic engineering method because of its distributed nature and ability to respond quickly to changes in network architecture or traffic loads.

Network terminology is frequently used to define bandwidth markets. Our approach is to use the concepts of microeconomics to define a bandwidth market. The result is that our bandwidth markets are similar to economic markets, which is advantageous for applying economic principles correctly.

This thesis presents the theoretical basis for two bandwidth markets. The first bandwidth market is a framework for building bandwidth markets. The second bandwidth market represents a society of cooperating individuals. The society distributes resources via a mechanism based on economic principles. An implementation of the bandwidth market is presented in the form of an optimisation algorithm, followed by its application to several test networks.

We show that, in the test networks examined, the optimisation algorithm reduces the network loss. For all test networks, the network loss achieved by the optimisation algorithm compares well with the network loss achieved by a centralised optimisation algorithm.

Opsomming

Verkeersingenieurswese bepaal die nodige bandwydtetoekenning om die verkeersvolume in 'n netwerk te dra. Op 'n soortgelyke wyse bepaal 'n ekonomiese mark die nodige hulpbrontoekenning om die aanvraag vir hulpbronne te bevredig. Die terme ***bandwydtemark*** stel verkeersingenieurswesetegnieke voor wat ekonomiese-mark metodes gebruik om die bandwydtetoekenning vir die verkeersvolume in 'n netwerk te bepaal. 'n Bandwydtemark is 'n aantreklike verkeersingenieurswesetegniek omdat dit verspreid van aard is en vinnig kan reageer op veranderinge in netwerk argitektuur en verkeersvolume.

Netwerkterminologie word gereeld gebruik om bandwydtemarkte te definieer. Ons benadering is om mikro-ekonomiese begrippe te gebruik om 'n bandwydtemark te definieer. Die resultaat is dat ons bandwydtemarkte soortgelyk aan ekonomiese markte is, wat voordelig is vir die korrekte toepassing van ekonomiese beginsels.

Hierdie tesis lê die teoretiese grondwerk vir twee bandwydtemarkte. Die eerste bandwydtemark is 'n raamwerk vir die ontwikkeling van bandwydtemarkte. Die tweede bandwydtemark stel 'n vereniging van samewerkende individue voor. Die vereniging versprei bandwydte deur middel van 'n meganisme wat gebasseer is op ekonomiese beginsels. 'n Implementasie van hierdie bandwydtemark word voorgestel in die vorm van 'n optimeringsalgoritme, gevolg deur die toepassing van die optimeringsalgoritme op 'n aantal toetsnetwerke.

Ons wys dat die bandwydtemark die netwerkverlies verminder in die toetsnetwerke. In terme van netwerkverlies vaar die bandwydtemark goed vergeleke met 'n gesentraliseerde optimeringsalgoritme.

Acknowledgement

This work was performed within the Telkom-Siemens Center of Excellence for ATM & Broadband Networks and their Applications and was supported by grants from the South African National Research Foundation, Telkom SA Limited and Siemens Telecommunications. I spent a sabbatical at the Teletraffic Research Centre at the University of Adelaide, Australia. My host there was Prof P.G. Taylor.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

In many telecommunication networks the network connectivity and the offered traffics change over time. If routes fail then traffic can be lost, and if traffics change then links may become overloaded. Traffic losses and link overloads are undesirable because they degrade the quality of the service experienced by the network users. In order to preserve quality of service, a network operator may apply *network engineering* and acquire additional bandwidth for overloaded links, or the network operator can apply *traffic engineering* and reroute the traffic over links that have spare bandwidth. This thesis focuses on traffic engineering and investigates methods for rerouting traffic.

It may be worthwhile for a network to respond to changes in traffic demand if the gain in revenue from a more accurate routing configuration outweighs the revenue loss incurred by additional signaling information on the network. Routing changes are mainly caused by changes in the traffic demand or changes in the underlying physical network. Typically a response mechanism, aware of local traffic and network configuration properties, will adjust the network resources allocated to the routes affected by the change. Major changes in the network are caused for example by link/node failures or traffic surges, which may require substantial changes in traffic routing and reallocation of network resources.

The *welfare* of a network can be evaluated in terms of the traffics and the network resources. A network operator can apply traffic engineering by employing a centralised optimisation algorithm that uses information about the traffics and the network resources to calculate a routing configuration that maximises the network welfare. Typically this information is acquired, a new routing configuration is calculated and the new routes are deployed. During the time between information acquisition and the deployment of the new routes, the traffics and resources of the network may change such that the newly calculated routing configuration no longer maximises the network welfare.

Alternatively, traffic engineering can be applied by a distributed optimisation method which deploys multiple agents in the network, each agent capable of adjusting the routing configuration.

The term agent implies an autonomous entity which acts on behalf of a client, ensuring the well-being of that client. Autonomous agents act without centralised control from a system coordinator, behaving entirely by local information and rules. The advantage of such a system over a centralised optimisation algorithm, is that an agent is able to ensure the well-being of a client without the delayed administration of a centralised entity. Moreover, the centralised execution of the optimisation algorithm represents a single point of failure whose malfunction would compromise the optimal operation of the network. A distributed system can be designed such that the actions of the agents not only ensure the well-being of the clients but also maximise the welfare of the network.

This thesis investigates the design of systems of distributed autonomous agents targeted to maximise the welfare of a network. A typical design of a system of distributed autonomous agents will include an objective for each agent and rules for collaboration amongst the agents. We investigate systems of distributed autonomous agents with agent objectives and rules of collaboration similar to those found in economics. Therefore we use economic terminology to define the systems of distributed autonomous agents and we apply economic principles in these *markets*.

Other authors have used economic principles to apply traffic engineering in telecommunication networks. The application of economic principles for traffic engineering is often called a *bandwidth market*. In [19] Kuwabara *et al.* define a market model that uses a pseudo price for controlling the allocation of link bandwidth. Gibney *et al.* also use a pseudo pricing approach in [14]. We investigate pseudo prices in economics (Chapter 5) before using a pseudo pricing approach for pricing bandwidth. In [32] Wellman describes a bandwidth market with a mechanism to calculate the *market clearing prices*. These market clearing prices yield an allocation of resources which maximises the welfare of the network.

The work presented in this thesis differs from most previous works, because it uses fundamental economic terminology to define the entities in a bandwidth market. For example, Wellman used economic terminology to define a bandwidth market in [32]. However, most previous works define bandwidth markets in network terms to which economic principles are applied. The advantage of defining the bandwidth market in the correct economic terminology is that it serves as foundation for examining the economic literature concerning the particular type of market.

An economy does not evolve because of the application of economic theory. The evolution of an economy is the result of a demand for resources and the efforts of multiple agents to transform and distribute resources in order to meet the demand. Economic theory is the scientific study of the evolution of an economy. Agents in an economy may however use economic theory to predict the evolution of the economy and thereby make optimum choices that maximise their well-being.

When designing a bandwidth market the designer must be careful not to apply economic theory to agents in an environment that does not resemble an economy. A common mistake is to name a network entity as an economic entity and ignore the differences in properties or behaviours that are premises for economic theory to render the correct result. In a bandwidth market where economic entities are defined incorrectly, the results of applying the theory of economics are limited. Other than presenting an optimisation algorithm through well known concepts, the usefulness of defining such a bandwidth market is questionable.

A distributed autonomous agent system is a distributed algorithm executed on a distributed processing platform[1]. We design two such systems of distributed autonomous agents and implement one of them. Anerousis and Lazar [3] survey a number of algorithms for creating virtual paths in networks and provide a taxonomy for the algorithms. According to their classification our distributed autonomous agent system is a data-driven, asynchronous, decentralised algorithm. The algorithm requires information about the network status, it is executed periodically, it is asynchronous to events in the network and it is a distributed multi agent system.

1.1 Game theory, markets and strategic equilibrium

Multi agent systems [15, 22, 30, 34] and equilibrium have been thoroughly studied in economics [11, 13, 25, 33]. This is not surprising, since events in real economic markets are determined by the strategy and behaviour of many independent agents (corporations and individuals). Their seemingly uncoordinated actions lead to an astonishingly stable economic market, in which most agents do well in achieving their goals such as maximising profit and maximising utility.

The agents in markets are interested in acquiring market resources. Competition for the same resources leads to conflict between the agents.

Game theory [9, 12, 16, 20, 23, 24] aims to analyse various problems of conflict between parties with diverse interests. Game theory has been applied to network routing problems and bandwidth allocation [1, 2, 17, 21]. Chapters 4 and 5 use game theory to analyse the markets which are designed to maximise the welfare of a network.

A market can be modeled as a game in which the players are the agents, the rules of the game are the market mechanisms and the strategy of an agent is the behaviour of the players in every possible situation. There are several formal representations of games. We use the *normal form* representation of a game. The definition of a game in normal form has four parts: players, strategies, payoff functions and additional rules. The first three may be symbolised by a triple $(\mathcal{N}, S, \mathbf{p})$ where $\mathcal{N}$ is a set of N players, $S = S_1 \times \cdots \times S_N$ is the strategy space which is the Cartesian product of the individual strategy sets S_i of the players $i \in \mathcal{N}$ and $\mathbf{p}$ is the vector of payoff functions. Given the chosen strategies $s_i \in S_i, i \in \mathcal{N}$ of all the players in the game, the payoff function $p_i(s_1, \ldots, s_N)$ of a player i is a numerical measure of how well the player does in the game. Additional rules [9, p.X] specifies the extent to which the players can communicate with one another, whether the players can or cannot enter into binding agreements, whether rewards obtained in the game may be shared and what information is available to the players in the game.

Games may differ in duration. In a *single period game* players choose their strategies, the game

[1] A distributed algorithm executed on a distributed processing platform is not necessarily a distributed autonomous agent system.

is played, the payoffs are evaluated and the game ends. A *multi period game* consists of T single period games, called the constituent games. The payoff of a multi period game is the sum $\sum_{i=1}^{T} \mathbf{p}(S)$ of the payoffs of the constituent games. The players may choose new strategies for each constituent game and this is done based on the strategies chosen in the previous constituent games. If $T \to \infty$ the multi period game is called a *super game.* Our market games are super games.

We assume that all players are *rational.* A rational player always[2] chooses a strategy that will maximise its payoff.

Cooperating players may follow a joint strategy which provides a mutually greater payoff than when the players do not cooperate. Games (and markets) are classified as *cooperative* or *non-cooperative.* Our market games are cooperative. A game is cooperative if the players are able to make binding, unbreakable agreements. Binding agreements are necessary for players to cooperate and they restrict the strategies of the cooperating players. In order to demonstrate this consider a simple two player super game where each player $i \in \{1,2\}$ can choose between strategy a_i or strategy b_i for each constituent game. The payoff function of each constituent game is given in Table 1.1 where the left hand value is the payoff of *Player 1* and the right hand value the payoff of Player 2.

	a_2		b_2	
a_1	0	0	9	−8
b_1	−8	9	8	8

Table 1.1: 2-player game.

Cooperating players may agree on the mutually best strategy pair (b_1, b_2). Games are non-cooperative if the players cannot make binding agreements. In this example if the players do not cooperate, Player 1 prefers the strategy pair (a_1, a_2) over (b_1, a_2) and (a_1, b_2) over (b_1, b_2). Thus regardless whether Player 2 chooses the strategy a_2 or b_2, the best strategy for Player 1 is strategy a_1. Since Player 1 and Player 2 are in the same situation the best strategy for Player 2 is also strategy a_2 and the players will choose the strategy pair (a_1, a_2) for each constituent game.

In any game an *equilibrium* strategy set is a set of player strategies by which the players earn a payoff such that no player has incentive to deviate from its strategy. The equilibrium discussed in this section is also referred to as a *strategic* equilibrium. In a multi agent system a strategic equilibrium is a strategy set that is persistently chosen by the system agents. The designer of the multi agent system is therefore interested in ensuring that all equilibria have favourable conditions (such as maximising the system welfare). Naturally it is important that the system designer knows the conditions of every possible strategic equilibrium.

[2]A reputation of being irrational can be advantageous to a player [18]. For example, in a super game, a player may choose strategies that do not maximise its payoff for some constituent games and let the other players think that it is irrational. This may change the choice of strategies of the other players which may enable the seemingly irrational player to take advantage and eventually earn a larger payoff for the super game. Note that such a player is rational in context of the super game.

A more general definition of a strategic equilibrium is: a combination of player strategies, agreed upon by cooperation or chosen by individual players, such that no player or group of players in coalition has an incentive to deviate from their equilibrium strategy.

One of the most well known strategic equilibriums is the *Nash equilibrium* [20, p.25]. A Nash equilibrium occurs when no player in the game can do better by changing its *own* strategy, thus abstaining from changes to its (conjectured) equilibrium strategy[3]. The Nash equilibrium applies only to non-cooperative games. Note that in the game with the payoff functions given in Table 1.1 the strategy pair (a_1, a_2) is a Nash equilibrium.

The Nash equilibrium cannot characterise a strategic equilibrium in cooperative games where cooperating players can agree to jointly change their strategies. That is to say, players that cannot do better by changing their own strategy can do better by agreeing with other players to jointly change their strategies. Cooperative strategic equilibrium occurs when no group of players is able to better their payoffs by jointly changing their strategies. Such a strategic equilibrium can be characterised as a *strong Nash equilibrium* [27, p.26]. A strong Nash equilibrium occurs when no subgroup (coalition) of players in the game can change their joint strategy such that each member in the subgroup has a better or at least the same payoff.

In most games the strong Nash equilibrium is too strong, meaning that the strong Nash equilibrium rarely exists. In choosing a cooperative equilibrium combination of player strategies it is more useful to define a *Pareto efficient* combination of player strategies.

Definition 1.1 *A combination of player strategies* **S** **dominates** *a combination of player strategies* **Z** *if each player earns at least as much payoff with the strategy set* **S** *as with strategy set* **Z** *and at least one of the players earns a strictly higher payoff with* **S** *than with* **Z***.*

For example consider the payoff vectors $(1, 1)$ and $(3, 4)$. An obvious interpretation is that the combination of player strategies resulting in a payoff vector $(3, 4)$ dominates the combination of player strategies resulting in a payoff vector $(1, 1)$. However for the payoff vectors $(3, 4)$ and $(4, 3)$, neither combination of player strategies dominates the other.

A Pareto efficient combination of player strategies is such that there is no other combination of player strategies by which some players do better without any player being worse off.

Definition 1.2 *A Pareto efficient combination of player strategies is a combination of player strategies that is not dominated by any other combination of player strategies.*

The Pareto efficient combination of player strategies proves to be more useful in many contexts.

The shaded interior region and the points on the enclosing boundary in Figure 1.1 illustrate the *payoff space* or the attainable payoffs for a two player game, where p_1 is the payoff of Player 1

[3]The formal definition is presented in Definition 4.4.

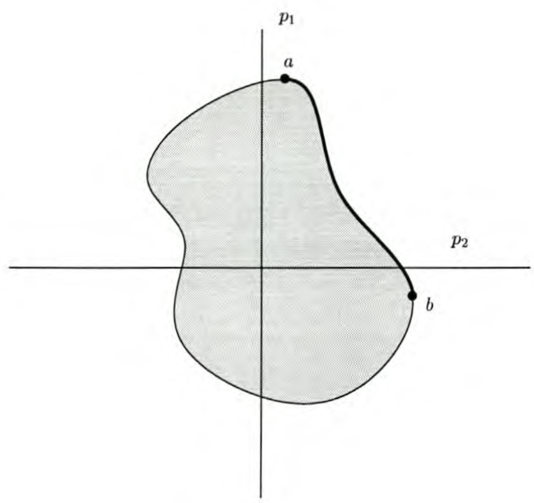

Figure 1.1: The Pareto frontier.

and p_2 the payoff of Player 2. According to Definition 1.1 the points on the marked boundary between a and b are not dominated by any point in the payoff space. The points on the marked boundary are *Pareto efficient payoff pairs* and the marked boundary is called the *Pareto frontier.* Each payoff pair in the payoff space is dominated by a Pareto efficient payoff pair on the Pareto frontier. The *Pareto inefficient payoff pairs* are the payoff pairs in the payoff space which are not on the Pareto frontier.

In most cases players in a coalition benefit more by choosing a Pareto efficient combination of player strategies rather than a Pareto inefficient combination of player strategies. In Chapter 4 we investigate a case where this is true. It is often difficult if not impossible for players to agree on which Pareto efficient combination of player strategies they should choose. There will always be a trade-off in choosing one Pareto efficient combination of player strategies over another Pareto efficient combination of player strategies, because changing from one Pareto efficient combination of player strategies to another Pareto efficient combination of player strategies will increase the payoff of one player at the cost of decreasing the payoff of another player.

In some games payoffs are *transferable*, which means that one player may transfer some of its payoff to another player. For instance if Players 1 and 2 choose a Pareto efficient point close to a then Player 1 may transfer an amount of payoff to Player 2 as a side-payment for the cooperation. In Chapter 4, we present a market model in which payoffs are transferable. Thus the cooperating players will prefer any Pareto efficient combination of player strategies that yields a higher aggregated payoff over an alternative that yields a lower aggregated payoff. Payoffs are not transferable in the cooperative market model presented in Chapter 5. The Pareto efficient combinations of player strategies which the players may agree on are therefore limited in the market model presented in Chapter 5.

Several theories [12, p.234-289] describe how players can *bargain* for a cooperative equilibrium. In our market models we assume the players have already agreed upon a combination of player strategies and we incorporate this agreed equilibrium combination of player strategies in the agent behaviours. It is a complex and interesting experiment to allow the players to bargain for a cooperative strategic equilibrium rather than provide them with a pre-determined cooperative equilibrium combination of player strategies. On the other hand it may prove to be too complex and resource intensive to be practical in real networks.

1.2 Problem statement

Using the concepts presented in the previous sections and some concepts of autonomous agent systems [30, Chapter 7], we state the problem investigated in this thesis.

Definition 1.3 *A* **mechanism design problem** *is a system of agents with a defined set of preferred* system state(s). *The system state determines the payoff of each individual agent and the goal of each agent is to maximise its payoff.*

Definition 1.4 *A* **mechanism** *for a mechanism design problem is a set of actions available to each agent in a system. The action of an agent can change the system state.*

A mechanism design problem and a mechanism together comprise a game.

Definition 1.5 *A* **solution mechanism** *for a mechanism design problem is a mechanism defining a game of which every strategic equilibrium necessarily leads to one of the preferred system state(s).*

This thesis defines two mechanism design problems and establishes a mechanism for each mechanism design problem. Although we do not prove that these mechanisms are mechanism design solutions for the mechanism design problems, the aim is to establish a mechanism design solution for each mechanism design problem.

In particular, the goal of this thesis is to define a mechanism design problem equivalent to the network resource allocation problem where resources are limited in quantity and distributed among agents. The mechanism design problem is defined such that it resembles an economic market. A mechanism that resembles the mechanisms used in real economic markets can then be designed. This mechanism may lead to an optimal resource allocation, which is the desired outcome defined by the mechanism designer.

In order to define a mechanism design problem that resembles an economic market we model the allocation of resources in a network as a market and identify the appropriate entities of the mechanism design problem in this market model.

1.3 Thesis layout

This thesis is organised as follows: Chapter 2 introduces the framework of the communication network. Chapter 3 presents the fundamental definitions used in the market models presented in this thesis.

Chapter 4 presents a general market model and identifies the entities of a mechanism design problem and a mechanism.

The following two chapters describe a second market model. Chapter 5 presents the definitions of the market model and identifies the entities of a mechanism design problem and mechanism. Chapter 6 presents an implementation of the market model presented in Chapter 5 along with a strategy for each agent and the results of testing an implementation of this market model. Chapter 7 presents our conclusions.

The notation used in this thesis is as follows. Bold capital letters denote matrices $\mathbf{X}$ of which $\mathbf{x}_i$ is the i-th row vector and x_{ij} the entry in the i-th row and j-th column. Let $[\mathbf{X}]_i = \mathbf{x}_i$ and $[\mathbf{X}]_{ij} = x_{ij}$ for any matrix $\mathbf{X}$. A vector is denoted by a lower case letter $\mathbf{x}$ and its i-th entry by x_i.

The vector $\mathbf{0}$ denotes the zero vector and $\mathbf{e}_i$ denotes the i-th unit vector. For a function $f(\mathbf{r}) = f(r_1, \ldots, r_n)$ we use the notation $D_i f(\mathbf{r}) = \partial f(r_1, \ldots, r_n)/\partial r_i$ to represent the partial derivative of the function f with respect to r_i. The sets $\mathbb{R}$, $\mathbb{R}_+$, $\mathbb{N}$, $\mathbb{N}_0$ and $\mathbb{Z}$ are the sets of real numbers, positive real numbers including zero, positive integers, positive integers including zero and integers respectively.

Chapter 2

The network framework

In this Chapter we define the network framework, which is a simplified definition of the real network [4] for which the market models in this thesis are designed.

2.1 Network

2.1.1 Physical link layer

The physical layer of a network (layer 1) is the physical fabric connecting communicating nodes in a communication network. The network physical layer can be static or dynamic. It is however not uncommon for a network to contain a hybrid of dynamic and static network components.

Static physical layer

The simplest physical network layer consists of physical links connecting communicating nodes. This is known as a static physical layer where the bandwidth between nodes is constant. Small networks like a local area network (LAN) or a wide area network (WAN) have a static physical layer.

Dynamic physical layer

Backbone networks, carrying large traffics from a variety of sources, have more advanced physical layers consist of switching network nodes connected by physical network links. The role of the switches is to produce single concatenated links from originating nodes to destination nodes via transit nodes. These switches enable dynamic changes to the bandwidth between nodes. Optical network technology is the most widely used and cost effective dynamic physical layer.

2.1.2 Optical physical network

Optical switches with optical fibre links are widely used in building dynamic physical layers. An optical switch is capable of switching ingress fibre links to egress fibre links, called Fibre-Switch Capable (FSC).

Different optical wavelengths can be multiplexed onto a single fibre using Wavelength Division Multiplexing (WDM). Each of these wavelengths can be "switched" to other wavelengths on different fibres (Lambda Switch-capability LSC), adding another dimension to switching and making the optical switch more complex. A cable of two multiplexed fibre optic links has a capacity of 120Gbps (1,451,520 × 64kbit/s channels). WDM has an economic advantage, since the previously installed FSC link capacities can be expanded at nominal cost. In 2001 Telkom [26] hosted Optical fibre links with WDM capable of carrying 240,000 channels, compared to an FSC capacity of only 30,000 channels.

There are two well known types of optical switches:

Optical Cross-connect (OXC) The signals traversing an OXC switch undergo optical/electrical/optical conversion (O/E/O). Error detection is done on the electrical signal and the LSC is accomplished by tunable lasers.

Photonic Cross-connect (PXC) Signals traversing a Photonic Cross-Connection (PXC) switch do not undergo electrical conversion. PXC switches are optical to optical devices, reducing the complexity of the switch, allowing a smaller footprint, and lowering power consumption.

Depending on the nature of the switch an OXC switch can switch around 252 to 1008 pairs of in-bound and out-bound ports at less than 1 microsecond per switching pair.

2.1.3 Logical link layer

In a dynamic physical layer, a logical link is a set of physical links between a node pair, cross connected by node switches (OXC/PXC). In the case of a static physical layer the logical links and physical links are the same entities.

The logical link layer is a meshed network topology on top of a sparse physical link topology, provisioned at different rates, such as T1, OC3, OC12, OC48, OC192, etc., depending on the requirement between nodes. As can be seen from Figure 2.1 the logical link layer (layer 2) is not necessarily fully connected.

On the edges of the logical links are packet routers. The role of a packet router is to transfer encapsulated data fragments, called packets, via the logical links to another packet router. A finite buffer is allocated to a packet router, for storing incoming packets while previously received packets

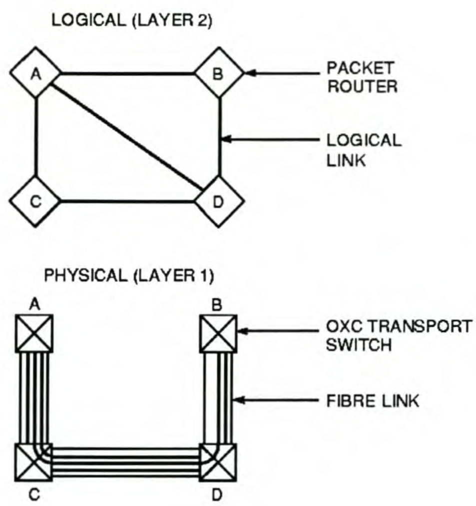

Figure 2.1: Physical and logical link layers.

are dispatched. Almost all packet routers are either Internet Protocol (IP) or Asynchronous Transfer Mode (ATM) routers.

2.1.4 Internet Protocol and Asynchronous Transfer Mode

IP routers use routing tables for forwarding packets over logical links. In the last decade of the 21 century, the version of IP (IPv4) was deemed insufficient for providing service to a considerable portion of communication formats (Voice, High Definition Video, etc.). One deficiency of IPv4 was its inability to provide a guaranteed quality of service. IPv4 packets are routed from ingress to egress router according to routing tables, which are populated and updated by path discovery protocols such as OSPF, normally minimising the number of logical links crossed between ingress and egress router. This is called destination-based connectionless routing. In IPv4 there are no mechanism for routing packets on a specified path, making traffic engineering difficult.

ATM [10] provides a connection-orientated service to protocols in higher abstraction layers. Traffics using an ATM network can be engineered to follow distinct paths through the network. Because of this property, most telephony telecommunication networks invested in ATM as a transport protocol.

The native protocol of most computers is IP and in computer networks the role of ATM is limited to transporting IP packets. Because of IP's dominance as the native protocol for computers, it will undoubtedly be the standard transport protocol for future integration of services on a shared communication network.

2.1.5 Multi-protocol Label Switching

The development of Multi-protocol Label Switching (MPLS) [8] commenced from a protocol called Tag Switching. The main function of the MPLS protocol is to enable the set up of explicit routes through a packet switched network and thus provide a connection-orientated structure to a connectionless network. The development of MPLS was motivated by a need to address

- scalability,
- decoupling of the routing scheme and the forwarding implementation,
- the complexity of mapping IP to ATM, and
- better price versus performance in routers.

The steep rise in demand for guaranteed quality traffic transport services motivated telecommunication companies to acquire a network infrastructure that could integrate and offer customer specified traffic transport services. Initially such a network infrastructure was not available and telecommunication companies invested huge amounts of money in network solutions that could offer limited services. Because of these large investments, the transition to a network solution that could integrate and offer a variety of customer specified traffic transport services must be able to interface with legacy networks. An attractive feature that MPLS offers to telecommunication companies is that it can interface with almost all legacy networks, translate the service requirements for the traffics of the legacy networks traversing the MPLS domain and maintain the quality of the services offered.

IP routers may be replaced by superior low cost MPLS routers called label switch routers (LSRs) and existing ATM routers will be converted to MPLS routers by reprogramming the control software. An LSR is responsible for identifying the next hop of each incoming MPLS packet and forwarding the packet accordingly. There are explicit paths (Label Switch Paths) over the links on which the MPLS packets are switched across the MPLS network.

Label Switch Paths

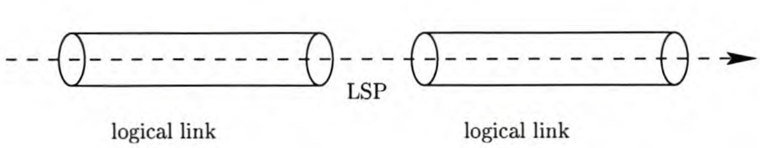

Figure 2.2: A label switch path along a logical link.

MPLS packets are forwarded along explicit paths called label switch paths (LSPs). LSPs are set up by configuring the LSRs, marking a path or a so called *flow*, through the MPLS network.

Each LSR has a table (Label Information Base) which is used to map a set of incoming labels to outgoing labels. Essentially these mapped pairs of labels identify flows through the LSR. These flows are further sub-classified into Forwarding Equivalence Classes (FEC). A FEC is a unique identifier for a flow through an MPLS network, flow granularity being a key feature in enabling MPLS to support multiple virtual networks; each with its own distinct flows. A packet submitted to an ingress LSR (origin node) requiring transportation to an egress LSR (destination node) via an explicit path has a label and a FEC assigned to it. Indicated by its label and FEC, a packet will be forwarded to the next LSR. At each LSR, the incoming label of an incoming MPLS packet is replaced by the mapped outgoing label. MPLS packets of an LSP will thus be routed across a series of logical links (Figure 2.2) by the core LSRs situated between the logical links. The traffic routing scheme is thus decoupled from the forwarding implementation in the LSR router. This is a major advance, since hardware or software routing optimisation agents can be changed without any changes in the forwarding implementation of the LSRs.

2.1.6 Virtual Network/ Virtual Path Connection Network

Constraint Routed Label Switch Paths (CRLSPs) are LSPs conforming to constraints or service quality preferences, for instance maximum delay or minimum bandwidth.

CRLSPs with common ingress/egress nodes and common constraints are grouped into a traffic trunk, called a Virtual Path Connection (VPC) in ATM networks. A traffic trunk is thus a collection of paths through a network between ingress and egress nodes, satisfying certain criteria. A service connection request to the traffic trunk is provided with equal *Quality of Service* (QoS), regardless of the CRLSP of the traffic trunk over which the connection is routed.

Similar to the grouping of CRLSPs into traffic trunks, traffic trunks are grouped into Virtual Networks (VNETs). These VNETs are fully connected virtual networks having certain characteristic properties (QoS). VNETs are dynamic and adaptive to changes in service demands, transport network alterations, and have high survivability during network failure. From the *network client's* perspective, the choice VNET is a reliable fully connected virtual network conforming to all the network client specifications and providing a sustained QoS. A separate virtual network for each service class exists on the physical network. When working with a specific virtual network, we call the traffic trunks the *origin-destination pairs* (OD-pairs) of that virtual network.

2.2 Quality of service

Network clients are the sources of connections (calls) to the network. Network clients wish to specify the quality of the data transmission service they are provided with. This specification is called the *Quality of Service* (QoS). We characterise and discuss briefly the major components of QoS.

Latency

Latency, also known as mean packet delivery time, is the elapsed time between the submission of data to an ingress node (origin node) and arrival at the egress node (destination node). This time delay is caused by *node delay*, which is incurred in examining and resubmitting data to the appropriate link transport medium and *link delay*, incurred in transporting data over the physical network link. Latency is of cardinal importance to real time services (interactive voice, interactive gaming).

Throughput speed

Measured in bit rate, throughput speed represents the amount of data throughput per unit time, which is largely dependent on the bandwidth of the network.

Jitter

Packets do not arrive at the egress node at regular time intervals. Some services, especially Constant Bit Rate (CBR) service classes, require data to arrive at the egress node at a sustained rate. A common metric for jitter over a time interval, is maximum difference in packet delivery time.

Transport reliability

Fallible physical transmission medium renders data transmission faulty to some extent. If detected, such transmission faults would require retransmission, or in case of real-time critical services, packet loss and degradation of service. The degree of unreliability is measured in packet loss probability or bit loss probability, referred to as packet or bit delivery grade of service (GoS).

Service connection blocking

Due to the conservation of QoS of the connections currently active on the network, additional service connections may be rejected. The number of connection requests rejected serves as measure of revenue loss. Under-provisioned networks will reject connection requests and lose revenue, in order to conserve the QoS of the active connections. From an QoS perspective, service connection blocking is part of the service Quality offered.

Security

Some network services require secure transmission. Security measures such as data encryption are applied at edge nodes.

Chapter 3

Fundamental definitions of the market

The definitions in this Chapter describe our *fundamental market model.* In later Chapters we use our fundamental market model as a foundation for designing more detailed market models.

3.1 Definitions

A market agent is assigned to every origin-destination (OD) pair, under the assumption that a connection is unbiased as to the LSP on which it is routed, except for the origin and destination of the LSP.

The primary interest of each agent is to manage network resources[1] for the transportation of traffic for that OD-pair, thus managing the capacity reserved for each LSP.

3.1.1 Payoff and utility

The term utility and payoff are both used in game theory and in the theory of economics. We present definitions of payoff and utility [16, p.3] using simple concepts.

Consider an individual with a choice of strategies $S = \{s_1, s_2, s_3, \ldots\}$. Each of these strategies describes an *action* or series of actions taken by the individual in every possible situation encountered. For example when a router is running out of buffer space the strategies may be: s_1 to drop a packet, s_2 to send a request to reduce the packet transmission rate or s_3 to drop a packet and send a request to reduce the packet transmission rate. Let $\mathcal{R} = \{\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \ldots\}$ be a set of vectors.

[1] The terms *resource* and *commodity* are used interchangeably.

Each vector quantifies an *outcome* of following a strategy. The outcome of a request to reduce the packet transmission rate may be to alleviate the shortage of buffer space.

Definition 3.1 *A* **payoff function** *of an individual is a function $p : S \to \mathcal{R}$ such that if the individual follows a strategy $s \in S$ then $p(s)$ is the outcome relevant to the individual. $p(s)$ is the* payoff *of the individual when following the strategy s.*

The outcomes of a strategy choice allow us to infer how well the individual does by choosing a particular strategy. For simple outcomes like revenue generation it is obvious that an individual does better by choosing strategies that yield higher revenue. However when outcomes are complex, it is less obvious how well an individual does with a strategy choice. For instance if an individual sends data over a network the relevant outcome is the quality of the data transmission (QoS). How well the individual does with a certain QoS depends on the preference of that individual.

A *utility function* quantifies the preference of an individual. The practical advantage of having a concise utility function rather than many individual preferences is significant [9, p.58]. The definition, rules and applications of utility functions are sometimes called *utility theory*.

Consider a set of k entities $E = \{e_1, \ldots, e_k\}$. Each of these entities may be anything of value to an individual: a resource, a bundle of resources or the payoff of following a strategy. For example e_1 may be a computer, e_2 a network connection and e_3 a resource bundle consisting of a computer and a network connection, or E may contain vectors describing QoS specifications.

Let $a, b, c \in E$. We define the *binary relations* $\succ, \prec, \sim$ on the entities in the set E such that

- $a \succ b$ if and only if entity a is preferred over entity b,
- $a \sim b$ if and only if the individual is indifferent to entity a or entity b,
- $a \succ b$ implies $b \prec a$.

The conditions

- either $a \succ b$, $a \prec b$ or $a \sim b$, and
- if $a \succ b$ and $b \succ c$ then $a \succ c$,

are met by the binary relations $\succ, \prec, \sim$ on the entities in set E. The binary relations $\succ, \prec, \sim$ are called *preference relations* on the entities in the set E.

Definition 3.2 *A* **utility function** *of an individual is a function $u : E \to \mathbb{R}_+$ such that $u(a) > u(b)$ if and only if $a \succ b$. $u(a)$ is the* utility *of the entity a.*

A utility function or payoff function may include other variables in its domain. An entity or strategy may then lead to different utilities or payoffs, depending on the values of these other variables.

A linear ordering of entities according to their utility will arrange the entities in order of an individual's preference. It is important to mention that the preference of an individual is observed whereupon a utility function is established which the individual seems to be maximising [9, p.56]. Strictly speaking it is wrong to say an individual is trying to maximise its utility; very likely the individual does not even know such a thing exists. Thus throughout this thesis when referring to an agent that maximises its utility we refer to an agent that acts as though it maximises its utility.

3.1.2 Commodities and utility

The *basic commodity* of the market is the bandwidth of the physical link between an OD-pair. This bandwidth can be reserved for use in different ways, for example the link bandwidth can carry direct traffic, or it can be used by transit routes traversing the link. Bandwidth can thus be assigned to routes, so that a route represents bandwidth assigned from the several physical links that make up the route. The route bandwidth is a *composite commodity* produced from a composition of basic commodities. Although commodities are discrete valued and allocated in unit quantities, in order to use derivatives we represent commodities as continuous variables and work with continuous functions.

Commodities are classified according to their OD pairs which in terms of MPLS are LSPs with identical originating and terminating nodes. For example commodities 1-3, 1-2-3 and 1-4-3 are commodities of class 1-3. Note that an OD-pair agent owns commodities of only one class.

Let N be the number of market agents and $\mathcal{N} = \{1, \ldots, N\}$ be a set of indices identifying the market agents. $\mathcal{N}$ is then also the set of indices identifying the commodity classes. The primary interest of an agent $i \in \mathcal{N}$ is to own an amount of commodities satisfying the *commodity requirement* of the network clients. An agent can thus be viewed as a *broker* who acts within the market in order to acquire commodities on behalf of the *consumers* (groups of network users).

An agent benefits by owning commodities. Let $\mathbb{R}^n_+$ denote the n-fold product of the set of positive real numbers, where $n \in \mathbb{N}$. An agent i owns a bundle $\mathbf{b}_i \in \mathbb{R}^n_+$ of commodities, where n is the number of commodities. The preference for owning a bundle of commodities of class i is given by the agent's utility function

$$u_i(\mathbf{b}_i, d_i)$$

where $u_i : \mathbb{R}^n_+ \times \mathbb{R}_+ \to \mathbb{R}$, $\mathbf{b}_i$ is a bundle of commodities and $d_i \in \mathbb{R}_+$ is the requirement for commodities of class i. Each utility function $u_i(\mathbf{b}_i, d_i)$ is bounded for bounded values of $\mathbf{b}_i$ and d_i.

Definition 3.3 *For a given commodity requirement d_i, $u_i(\mathbf{b}_i, d_i) > u_i(\mathbf{b}_i^*, d_i)$ if and only if an agent i prefers a bundle $\mathbf{b}_i$ of commodities over a bundle $\mathbf{b}_i^*$ of commodities.*

In our model it is necessary that the utility functions of all the agents be twice differentiable and monotone increasing so that $D_j u_i(\mathbf{b}_i, d_i) > 0$ (see Section 1.3) for $j \in \{1, \ldots, n\}$. Furthermore the utility function $u_i(\mathbf{b}_i, d_i)$ is *concave* on $\mathbf{b}_i \in \mathbb{R}^n_+$.

Definition 3.4 *A multi-variable twice differentiable function* $u(\mathbf{b}_i, d_i)$ *is concave on* $\mathbf{b}_i \in \mathbb{R}_+^n$ *if* $D_j^2 u(\mathbf{b}_i, d_i) \leq 0$ *for* $j \in \{1, \ldots, n\}$.

Definition 3.5 *A multi-variable twice differentiable function* $u(\mathbf{b}_i, d_i)$ *is convex on* $\mathbf{b}_i \in \mathbb{R}_+^n$ *if* $D_j^2 u(\mathbf{b}_i, d_i) \geq 0$ *for* $j \in \{1, \ldots, n\}$.

The characteristics of the utility function $u_i(\cdot)$ with respect to the commodity requirement d_i will be discussed when presenting the particular market model. For simplicity, we will on occasion assume that the commodity requirements are time independent, so that each commodity requirement d_i is constant.

There are two kinds of utility functions in economic theory. An *ordinal utility function* generates a ranking of commodity bundles in order of most preferred to least preferred while a *cardinal utility function* quantifies by how much one commodity bundle is preferred to another. When allocating commodity bundles to agents, it is necessary to know how much an agent prefers one commodity bundle to another. It is also necessary to know how much one agent prefers a commodity bundle than another agent. We therefore use cardinal utility functions.

Because we use cardinal utility functions, a utility defines a numerical score representing the *value* of a bundle of commodities [25, p74] to an agent. The value quantifies how well the commodity bundle $\mathbf{b}_i$ does in meeting the commodity requirement d_i.

A unique *class number* is assigned to every commodity class and a unique *commodity number* is assigned to each commodity in every class. For example commodities 1-3, 1-2-3 and 1-4-3 are numbered 1, 2 and 3 respectively and are of class 1-3 which is numbered 1. The bundle $\mathbf{b}_i$ therefore contains an amount b_{ij} of commodity number j of class i.

Throughout this thesis, when referring to the *utility of an agent* we refer to the utility of the bundle of commodities owned by that agent.

Definition 3.6 *Given a commodity requirement* d_i*, the utility* $u_i(\mathbf{b}_i, d_i)$ *is the rate at which agent* i *receives benefits for owning the commodity bundle* $\mathbf{b}_i$.

There is an important difference between models of markets in a real economy and the models we present. In a real economy commodities are created and destroyed (or consumed), whilst in our market models in general[2] commodities are indestructible and cannot be created, they merely change form. A consumer in a real economy that has utility for a certain commodity bundle may find that the utility of that commodity bundle diminishes due to the degraded quality or consumption (destruction) of that commodity bundle. In our market models the only factors that can change the utility of a commodity bundle are factors external to the commodity itself, such as the requirement for commodities in that commodity bundle. Commodities in the market models we investigate are therefore limited in quantity and are indestructible.

[2] Commodities may be be destroyed (link failure) or created (bandwidth added to network) only on rare occasions.

3.1.3 The production of commodities

An agent can transform commodities into other classes of commodities. The detail of the transformation depends on the capability of the agent to transform commodities, which is called the *technology* of the producing agent. The market uses a production function $P_k(\mathbf{X})$ to produce a class k commodity from a *production bundle* $\mathbf{X}$ of commodities. The domain of the production $P_k(\mathbf{X}) = \mathbf{y}_k$ is a non-zero matrix where $x_{ij} \geq 0$ is the amount of commodity number j of class i used in the production. The range of $P_k(\mathbf{X}) = \mathbf{y}_k$ is a vector $\mathbf{y}_k$ where $y_{kj} \geq 0$ is the amount of commodity number j of class k produced.

The use of a production function is demonstrated by an example production

$$P_3 \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$$

which uses one unit of commodity number 1 of class 1 and one unit of commodity number 1 of class 2 to produce one unit of commodity number 2 of class 3.

Assumption 3.1 *A production $P_k(\mathbf{X}) = \mathbf{y}_k$ is restricted to produce only a single commodity in the bundle $\mathbf{y}_k$.*

Corollary 3.1 *From Assumption 3.1 it follows that for a production $P_k(\mathbf{X}) = \mathbf{y}_k$ there exists only one element $y_{kj} > 0$ in the vector $\mathbf{y}_k$.*

Assumption 3.2 *For a production $P_k(\mathbf{X}) = \mathbf{y}_k$ with $y_{k\ell} > 0$ either $x_{ij} = y_{k\ell}$ or $x_{ij} = 0$ for all i and j.*

Assumption 3.2 is a reasonable assumption in a network where a route uses equal amounts of capacity from each link which it traverses.

Corollary 3.2 *According to Assumption 3.1 and Assumption 3.2 we can write a production producing commodity number ℓ of class k as $P_k(\mathbf{X}) = P_k(\lambda\mathbf{R}) = \lambda\mathbf{e}_\ell = \mathbf{y}_k$, where $\mathbf{R}$ is a matrix containing only 1's and 0's, and $\lambda > 0$ is the amount of commodity number ℓ of class k produced. We call $\mathbf{R}$ the* unit production bundle.

Our next assumption concerning productions is very important. In section 3.1.4 we see that this assumption is part of a policy which simplifies the market mechanisms.

Assumption 3.3 *The technology of a producing agent is restricted such that composite commodities can only be produced from basic commodities.*

Production bundles that are linearly independent are *distinct*:

Definition 3.7 *A production bundle* $\mathbf{X} \neq \mathbf{0}$ *is distinct from a production bundle* $\mathbf{Y}$ *if and only if* $\mathbf{X} \neq \beta\mathbf{Y}$ *for all* $\beta \in \mathbb{R}_+$.

Distinct production bundles have unit production bundles that are unequal:

Lemma 3.1 *For distinct production bundles* $\mathbf{X}$ *and* $\mathbf{Y}$ *and their respective unit production bundles* $\lambda_X\mathbf{R}_X = \mathbf{X}$ *and* $\lambda_Y\mathbf{R}_Y = \mathbf{Y}$ *it follows that* $\mathbf{R}_X \neq \mathbf{R}_Y$.
Proof. Because $\mathbf{X}$ is distinct from $\mathbf{Y}$, according to Definition 3.7 $\mathbf{X}$ is not expressible as $\mathbf{X} = \beta\mathbf{Y}$, in particular for $\beta = \lambda_X/\lambda_Y$ from which it follows that $\mathbf{R}_X \neq \mathbf{R}_Y$. ∎

Definition 3.8 *Each* distinct *production bundle* $\mathbf{X}$ *maps to a unique commodity number* ℓ_X *of class* k *with* $P_k(\alpha\mathbf{X}) = \alpha P_k(\mathbf{X}) = \alpha\mathbf{y}_k = \alpha\lambda_X\mathbf{e}_{\ell_X}$. *This mapping is unique in the sense that for any production bundle* $\mathbf{Y}$ *distinct from the production bundle* $\mathbf{X}$, $P_k(\mathbf{R}_X) \neq P_k(\mathbf{R}_Y)$, *where* $\mathbf{R}_X$ *and* $\mathbf{R}_Y$ *are unit production bundles with* $\lambda_X\mathbf{R}_X = \mathbf{X}$ *and* $\lambda_Y\mathbf{R}_Y = \mathbf{Y}$.

Definition 3.9 *Two commodities are distinct if and only if the commodities are produced from distinct production bundles.*

Theorem 3.1 *For any two distinct production bundles* $\mathbf{X}$ *and* $\mathbf{Y}$ *it follows that* $P_k(\mathbf{X}+\mathbf{Y}) \neq P_k(\mathbf{X}) + P_k(\mathbf{Y})$.
Proof. Assume that $P_k(\mathbf{X}+\mathbf{Y}) = P_k(\mathbf{X}) + P_k(\mathbf{Y})$ for at least one pair of distinct production bundles $\mathbf{X}$ and $\mathbf{Y}$. According to Corollary 3.2 we can write $P_k(\mathbf{X}+\mathbf{Y}) = \lambda_{X+Y}\mathbf{e}_{\ell_{X+Y}}$, $P_k(\mathbf{X}) = \lambda_X\mathbf{e}_{\ell_X}$ and $P_k(\mathbf{Y}) = \lambda_Y\mathbf{e}_{\ell_Y}$, with $\lambda_{X+Y}, \lambda_X, \lambda_Y > 0$. However $\mathbf{X}$ and $\mathbf{Y}$ are distinct production bundles thus according to Definition 3.8 $\mathbf{e}_{\ell_X} \neq \mathbf{e}_{\ell_Y}$, implying that $\alpha_1\mathbf{e}_{\ell_X} + \alpha_2\mathbf{e}_{\ell_Y} \neq \alpha_3\mathbf{e}_r$ for all $\alpha_1, \alpha_2, \alpha_3 > 0$ and all $r \in \mathcal{R}_k$ where $\mathcal{R}_k$ is the index set of the commodities of class k. This contradicts the assumption $P_k(\mathbf{X}+\mathbf{Y}) = \lambda_{X+Y}\mathbf{e}_{\ell_{X+Y}} = P_k(\mathbf{X}) + P_k(\mathbf{Y}) = \lambda_X\mathbf{e}_{\ell_X} + \lambda_Y\mathbf{e}_{\ell_Y}$. ∎

From Theorem 3.1 it follows that P_k is a *non-linear transformation*[3], because there exist production bundles $\mathbf{X}$ and $\mathbf{Y}$ such that $P_k(\alpha\mathbf{X}+\mathbf{Y}) \neq \alpha P_k(\mathbf{X}) + P_k(\mathbf{Y})$.

Definition 3.10 *The inverse of* $P_k(\cdot)$ *is called the decomposition function* $P_k^{-1}(\cdot)$. *If* $P_k^{-1}(\mathbf{b}) = \mathbf{X}$ *then* $\mathbf{X}$ *is called the* **decomposition**[4] *of* $\mathbf{b}$. *The production function characterises the decomposition function as follows: For every composite commodity number* j *of class* k *and* $\lambda > 0$ *let* $P_k^{-1}(\lambda\mathbf{e}_j) = \lambda\mathbf{R}_j$ *be a linear transformation where* $P_k(\lambda\mathbf{R}_j) = \lambda\mathbf{e}_j$ *is a mapping from* $\mathbf{R}_j$ *to a unique* $\mathbf{e}_j$ *for each* $\mathbf{R}_j$ *according to Definition 3.8.*

Note that $\mathbf{b}$ may contain more than one composite commodity whereby the decomposition of $\mathbf{b}$ may comprise several production bundles.

[3] P is a linear transformation if $P(\alpha\mathbf{X}+\mathbf{Y}) = \alpha P(\mathbf{X}) + P(\mathbf{Y})$ for all $\mathbf{X}$ and $\mathbf{Y}$ otherwise P is a non-linear transformation.

[4] *If* $\mathbf{b}$ *contains only a single composite commodity then the decomposition is the production bundle of this composite commodity.*

When using the production functions $P_k(\cdot)$ and $P_k^{-1}(\cdot)$ it is important to be aware of their (non) linearity. Consider for example the production bundles $\mathbf{X}_1$, $\mathbf{X}_2$, $\mathbf{X}_3$ and production function $P_k(\cdot)$ and decomposition function $P_k^{-1}(\cdot)$: because P_k is a non-linear transformation $P_k(\mathbf{X}_1)+P_k(\mathbf{X}_2)+P_k(\mathbf{X}_3) \neq P_k(\mathbf{X}_1+\mathbf{X}_2+\mathbf{X}_3)$, though $P_k^{-1}\left(P_k(\mathbf{X}_1)+P_k(\mathbf{X}_2)+P_k(\mathbf{X}_3)\right) = \mathbf{X}_1+\mathbf{X}_2+\mathbf{X}_3$, because P_k^{-1} is a linear transformation.

A commodity is called a basic commodity if it cannot be produced from a bundle of other commodities except a bundle containing the basic commodity itself:

Definition 3.11 *A commodity number γ contained in bundle $\mathbf{y}_k$ as $y_{k\gamma} > 0$ with $y_{kj} = 0$ for all $j \neq \gamma$ is a basic commodity of class k if and only if $P_k(\mathbf{X}) = \mathbf{y}_k$ implies that*

$$\mathbf{x}_i = \begin{cases} \mathbf{y}_k & i = k \\ \mathbf{0} & \textit{otherwise.} \end{cases}$$

A commodity is therefore a basic commodity if and only if any production producing the commodity only has the commodity itself in the production bundle.

Assumption 3.4 *We assume that there is only one commodity in each class $i \in \mathcal{N}$ satisfying the conditions defining a basic commodity in Definition 3.11.*

For the remainder of this thesis, where no confusion can arise we let the domain of the production $P_i(\mathbf{x}) = \mathbf{y}_i$ and the range of the decomposition $P_i^{-1}(\mathbf{y}_i) = \mathbf{x}$ be a vector $\mathbf{x}$ where x_k is the amount of basic commodity of class k used to produce a composite commodity of class i in the commodity bundle $\mathbf{y}_i$.

Without loss of generality, let commodity number 1 denote the commodity number of the basic commodity of class $i \in \mathcal{N}$ and b_i the amount of basic commodity of class i in the commodity bundle $\mathbf{b}_i$. The commodity number $j > 1$ of class i is a composite commodity and b_{ij} the amount of composite commodity number j of class i in the commodity bundle $\mathbf{b}_i$.

Our agents have incorporated production capability. An agent can therefore acquire a bundle of basic commodities with the intent of producing a composite commodity. Since each agent owns and produces composite commodities of one specific class, we consider them to be *specialists* in that class of commodity.

Definition 3.12 *The* **composite commodity allocation** *of a market is a matrix* $\mathbf{B}$ *representing the commodity bundles $\mathbf{b}_i$ of the agents $i \in \mathcal{N}$.*

Definition 3.13 *We define*

$$\mathcal{K}_{i,j} = \left\{k : \left[P_i^{-1}(\mathbf{e}_j)\right]_k = 1\right\}$$

as the index set of the basic commodities required to produce a composite commodity number j of class i. The set $\mathcal{K}_{i,j}$ contains the indices of the specialists from which basic commodities are required to produce a composite commodity j of class i.

3.1.4 Trading

Given constant commodity requirements and resources, it is desirable that the market reaches an equilibrium where no further trades of commodities take place. This equilibrium is regarded as a *market equilibrium* since it is not an equilibrium reached in the choosing of trading strategies but rather an equilibrium allocation of commodities.

Definition 3.14 *A* **market equilibrium** *occurs when none of the agents in a market is able to trade commodities.*

At a market equilibrium, each agent is either content with the utility of its commodity bundle or discontent with the utility of its commodity bundle but unable to trade commodities that will improve the utility of its commodity bundle. The market moves towards market equilibrium by agents trading resources.

Definition 3.15 *The* **state** *of a market is the information about everything in the market. The market state will include among others information about the agents in the market, the agent utility functions, the requirement for commodities, amount of basic commodities and the composite commodity allocation (see Definition 3.12).*

Definition 3.16 *A* **market process** *is a sequence of market states, representing the states of a market over time.*

An agent may make poor-trades which are poor by global standards, but which are reasonable by local standards. A poor-trade necessitates a future trade to reverse the losses incurred by the poor-trade. Such poor-trades cause *oscillations.* When trading resources in a market, it is desirable that trading proceeds in such a way as to prevent oscillations and so minimise communication between agents (overhead cost). An oscillating market has redundant trades which increase the overhead cost of trading and increase the time before equilibrium is reached. Avoiding oscillations inherently lessens overhead cost. This conforms to our aim of converging the market process (non-oscillating) towards market equilibrium.

The autonomy of the agents in a distributed market may inadvertently give rise to oscillations. The behaviour of one agent may indirectly stimulate the reaction of another agent, which in turn can influence the behaviour of the original agent. Such pairs, or perhaps groups of agents with recurrent behaviour can potentially cause recurrent trades of the same resource, thereby producing oscillations which increase overhead cost.

In order to avoid oscillations and reach equilibrium we enforce a policy which restricts the trading strategies of the agents in the market. The policy is applied by setting the following restrictions on trading.

Basic resource cache: A basic resource belongs only to the specialist of that basic resource. A basic resource is not distributed among many agents, which simplifies the market since the market agents only have to contact one agent (the specialist) to acquire a certain basic resource.

Resource production from basic resources: Only basic resources are traded, which simplifies the technology of production (only basic resources are used in production bundles).

An agent i initiates a trade whenever there is an opportunity to improve the utility of its bundle of commodities $\mathbf{b}_i$ or an opportunity to collaboratively improve the utilities of the bundles of commodities of a group of agents. The agent will then take one of the following actions:

- Select a production formula for a composite resource and contact the respective specialists to negotiate a trade of the basic resources required to produce the composite resource. This may take place in the form of *auctions* (section 3.1.7).
- Select a composite resource and contact the respective specialists to negotiate the decomposition and trade of the basic resources from which the composite resource was produced. This may also take place in the form of *auctions* (section 3.1.7).

The following subsections present the *market mechanisms* used in the trading process. The trading of commodities is a two-phase process: the selection of a commodity, ruled by the current commodity price, followed by the trade and production or decomposition of that commodity.

3.1.5 Market demand and supply

An agent owns commodities in order to satisfy the commodity requirements of its clients. If an agent owns more commodities than its clients require, then there is an excess amount of commodities and the agent may sell commodities to other agents. Alternatively if an agent owns fewer commodities than its clients require, then there is a deficit amount of commodities and the agent may buy commodities from other agents. Hence the clients' requirement for commodities causes their agents to either *demand* commodities from the market or *supply* commodities to the market.

Since only basic commodities are traded in the market, demand and supply only apply to basic commodities.

Definition 3.17 *The* **market demand** *and* **market supply** *of basic commodity of class k refers to the aggregated demand and supply of basic commodity of class k of the agents $\mathcal{N} - \{k\}$.*

Definition 3.18 *The* **specialist demand** *and* **specialist supply** *of basic commodity of class k refers to the demand and supply of basic commodity of class k of the specialist k.*

Due to the trading restrictions, from the perspective of any agent $i \in \mathcal{N} - \{k\}$ the specialist k is the only other supplier or demander of basic commodity of class k.

A *demand curve* is the unit price that an agent or group of agents is prepared to pay per quantity of commodity demanded and a *supply curve* is the unit price that an agent or group of agents is prepared to accept per quantity of commodity supplied. A demand curve describes the negative relation between the unit price of a commodity and the quantity of commodity that agents demand: an increased unit price leads to lower quantity of commodity demanded. Conversely a supply curve describes the positive relation between the unit price of a commodity and the quantity of commodity that agents supply: an increased unit price leads to increased quantity of commodity supplied. Due to the diminishing marginal utility[5] of commodities the demand curves are convex (Definition 3.5). Similarly due to the increased marginal cost of commodities supplied the supply curves are convex.

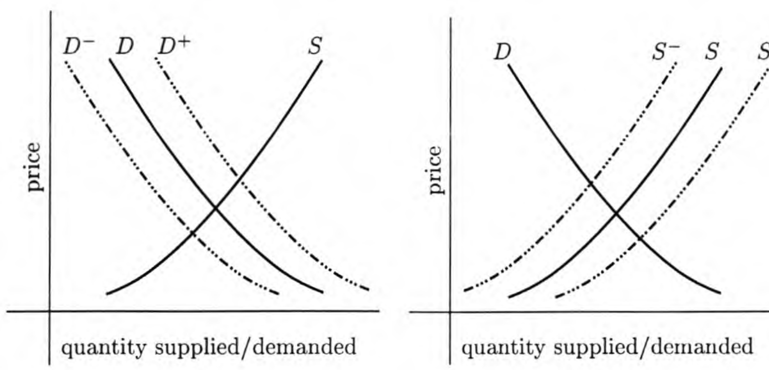

Figure 3.1: Demand and supply curves demonstrating changes in supply and demand.

Figure 3.1 illustrates the demand curve D and supply curve S. According to economic theory the following rules apply to supply and demand:

- If demand decreases then the quantity used at every price decreases, implying that the demand curve D shifts to the left onto D^-.
- If demand increases then the quantity used at every price increases, implying that the demand curve D shifts to the right onto D^+.
- If supply decreases then the quantity supplied at every price decreases, implying that the supply curve S shifts to the left onto S^-.
- If supply increases then the quantity supplied at every price increases, implying that the supply curve S shifts to the right onto S^+.

If the quantity of basic commodity supplied mismatches the quantity of basic commodity demanded, the agents will trade basic commodities in order to decrease the difference between the

[5] Marginal utility is the increase in utility per additional unit of commodity.

quantity of basic commodity supplied and the quantity of basic commodity demanded [6, p.111-116, p.165]. Later in this section we show how this occurs.

If the quantity of basic commodity supplied equals the quantity of basic commodity demanded, then a market equilibrium occurs. The intersection of the demand curve and the supply curve renders a market equilibrium. At this market equilibrium the (equilibrium) quantity of basic commodity supplied is the *market clearing quantity* and the (equilibrium) unit price is the *market clearing unit price*. The market clearing unit price is special because it is dually the highest unit price at which the quantity demanded is not less than the quantity supplied and the lowest unit price at which the quantity supplied is not less than the quantity demanded. It is therefore favourable for both the supplier and demander.

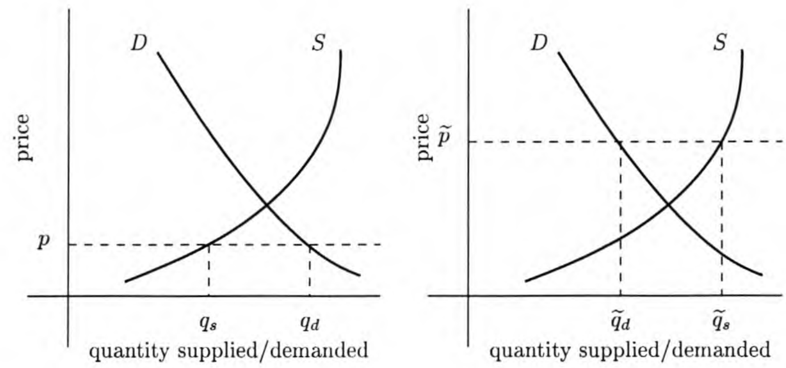

Figure 3.2: Demand and supply curves demonstrating market disequilibria.

Figure 3.2 depicts the supply curve S and demand curve D of a basic commodity. The graphs illustrate market disequilibria where the quantity supplied q_s is less than the quantity demanded q_d and the quantity supplied $\widetilde{q}_s$ is more than the quantity demanded $\widetilde{q}_d$ respectively.

Consider the case where $q_s < q_d$. Because the quantity demanded is higher than the quantity supplied the supplier will increase the quantity supplied and charge a higher unit price. Then because of the higher unit price the quantity demanded will decrease. Consider the case where $\widetilde{q}_s > \widetilde{q}_d$. Because the quantity supplied is higher than the quantity demanded the supplier will decrease the quantity supplied and charge a lower unit price. Then because of the lower unit price the quantity demanded will increase.

It is clear that the supplier and demander will end up in a market equilibrium, where the quantity supplied equals the quantity demanded.

In Chapter 5 demand and supply will play an important role in defining the unit price of a basic commodity.

3.1.6 The pricing of commodities

Each commodity has a price. How prices are calculated and used are specific to the particular market model. We define the calculation and use of prices in Chapter 4 and Chapter 5.

3.1.7 Auctioning

An *auction* is a mechanism for determining the distribution and the highest unit price of commodities where multiple agents with varying demands for commodities buy the commodities from single sellers. Likewise, an auction is also a mechanism for determining the distribution and the lowest unit price of commodities where multiple agents with varying supplies of commodities sell the commodities to single buyers.

The details of the auctioning mechanism will be discussed when presenting a particular market model.

3.2 Relating the market to an MPLS network

The previous sections defined a framework of a market and the objectives of the market agents. We now specify the relationship between the market and an MPLS network.

Note that the commodity classes have nothing to do with the service classes in MPLS. A commodity class is a unique number identifying a basic commodity (the bandwidth of a single link LSP) and composite commodities (the bandwidth of multi-link LSPs). A service class in MPLS identifies a virtual network layered onto the physical architecture of the IP network. The traffics on a virtual network are protected from traffics on other virtual networks, so that from the point of view of the user a virtual network seems to be a separate physical network.

In an MPLS network, bandwidth is not allocated to an LSP. LSPs are bundled together into aggregates, and these aggregates' capacities are protected from each other. This is done by fair queuing mechanisms at the label switching routers (LSRs). Bandwidth may however be *reserved* for LSPs.

Our market-based routing scheme is set within an MPLS network service class (virtual network), with protected link (direct route) bandwidth. In market terms, the virtual network link bandwidths represent the *original basic commodity bundle*. For the remainder of this thesis let $\mathbf{q} = (q_1, \ldots, q_N)$ denote the basic commodities in the original basic commodity bundle of the market.

Some virtual link bandwidth is reserved for each direct route LSP. The remaining bandwidth on each virtual link may be reserved for multi-link LSPs traversing the virtual link.

During the market process the following relations are true :

$$\sum_{r \in \mathcal{N}} [P_r^{-1}(\mathbf{b}_r)]_i = q_i$$

where $\mathcal{N}$ is the set of agents (OD pairs) and b_{rj} the bandwidth reserved for LSP j of OD pair r. In other words the decomposition of all the composite commodities in the market is equal to the original basic commodity bundle. In market terms, the traffic offered to an OD-pair $i \in \mathcal{N}$ is the commodity requirement d_i.

The task of the agents in the market is to manage the bandwidth reserved for connections on the LSPs. Given a configuration of capacities on LSPs, connections are routed via LSPs with enough available reserved capacity.

3.3 Market overview

A commodity class is assigned to each commodity in the market. By means of productions, agents may combine basic commodities of different classes to produce a composite commodity of another class. Each agent i in the market has a utility function $u_i(\cdot)$ which maps the bundle of commodities it owns to a utility value. Owning more commodities increases the utility value of the bundle of commodities owned by an agent. Some classes of commodities increase the utility of certain agents more than others.

3.3.1 The market variables and market equilibrium

We make a distinction between *endogenous* variables and *exogenous* variables of a market model [6, p.9].

Definition 3.19 *The endogenous variables of a market model are those variables that are only affected by the choices of individuals (agents or players).*

Definition 3.20 *The exogenous variables are the ultimate givens of the market model and their values are not affected or chosen by individuals in the market.*

The exogenous variables of the market model presented in this Chapter are the set of agents $\mathcal{N}$, the basic commodity bundle $\mathbf{q}^t$, the commodity requirements $\mathbf{d}^t$, the production functions $\mathbf{P}(\cdot)$ and the utility functions $\mathbf{u}(\cdot)$. The exogenous variables may change over time t.

The endogenous variables of the market model presented in this Chapter are the composite commodity allocation $\mathbf{B}^t$. The endogenous variables may also change over time t.

Agents in a market will trade if they have incentives to do so. The trades cause the endogenous variables to change and may lead to a market equilibrium. When there are no incentives for agents to trade, a market equilibrium occurs and none of the endogenous variables in the market changes during a market equilibrium. A market in equilibrium can only be brought into disequilibrium by a change in an exogenous variable [6, p.21].

3.3.2 The basic market models in economics

All models of economic markets are descendants of four *basic market models*. In this section we classify our fundamental market model as a descendant of one of these basic market models.

Table 3.1 presents the characteristics of these four basic market models. Here the agents represent the suppliers of the basic commodities. It is important to note that the table presents the normal characteristics of each of these market models. For example an oligopoly normally has differentiated products, however it may rarely have homogeneous products. The characteristics of monopolistic competition are so similar to that of an oligopoly that the economic theory of the behaviour of agents in a monopolistic competition is applicable to an oligopoly [12, p.45].

Characteristic	Perfect competition	Monopolistic competition	Oligopoly	Monopoly
Number of agents	many	many	few	one
Type of Product	homogeneous	differentiated	differentiated	unique
Restrictions on entry	none	some	significant	entry blocked
Price competition	fierce	considerable	some	none
Commodities example	non-durable (minerals)	non-durable (food)	durable (refrigerators)	rare and durable (diamonds)

Table 3.1: Characteristics of the four basic market models.

We note that:

1. The commodities of our fundamental market model are durable (indestructible). According to Table 3.1 this is a characteristic of an oligopoly or monopoly.
2. In our fundamental market model there are distinct classes of commodities and the marginal utility of composite commodities may differ, therefore commodities are differentiated which is a characteristic of monopolistic competition and an oligopoly.
3. Each class consists of several commodity types. The basic market model is therefore not a monopoly.
4. An agent can manufacture composite commodities. There are only a few agents supplying the basic commodities required for these productions. This is a characteristic of an oligopoly.

According to the arguments about we classify our fundamental market model presented in this Chapter as a descendant of the oligopoly basic market model in Table 3.1.

3.3.3 Market models

It is evident that there are opportunities for the agents to trade resources in order to arrive at a mutually better allocation of resources.

We consider two market models and their respective means of solving the resource allocation problem. Before describing the market models we present a few definitions.

Definition 3.21 Self-interest *is a behaviour of an agent whose goal is to increase the well-being of itself. Any increase or decrease of another agent's well-being is inconsequential.*

Definition 3.22 Altruism *is a behaviour of an agent (altruist) whose goal is to increase the well-being of one or more agents other than itself. The goal of the altruist need not include the increase of its own well-being.*

Definition 3.23 Cooperation *is a behaviour of two or more agents who coordinate their strategy choices in order to increase the well-being of the individual cooperating agents.*

In general the utility of an agent is a measure of the well-being of that agent. Note that a self-interested agent may (and will) cooperate in order to increase its own well-being, although the increase or decrease of the other agents' well-being is inconsequential.

- *The dual-oligopoly* is a market of *self-interested* agents that interact to find a better resource allocation. Each agent trades resources in order to increase its own utility. Agents may cooperate or compete by setting prices for resources in the oligopoly presented in Chapter 4.
- *The social community* is a market where a society of *self-interested* agents[6] *cooperate* in order find a better resource allocation. Rules are established by which all agents abide. These rules restrict the trading of resources and dictate the prices of resources. The reactions of the agents to the restrictions on trading and resource prices constitute a better resource allocation. Chapter 5 presents the social community.

[6]Although it may appear that an agent is altruistic in context of a single period game, an agent only cooperates in order to increase its own well-being. Each agent is therefore self-interested.

Chapter 4

The dual-oligopoly

This Chapter maps the resource allocation problem in a network to a market model that is a descendant of the oligopoly basic market model.

In section 4.4 the naming dual-oligopoly will become apparent. This Chapter defines the structure of a market model and does not specify a method for setting buying or selling prices, neither does it specify whether or how the agents in the market should cooperate. This structure of a market model may serve as basis for modeling a bandwidth market that conforms closely to market models in economic literature.

4.1 Wealth and utility

In the dual-oligopoly, the utility function $u_i(\mathbf{b}_i, d_i)$ is defined as the rate at which an agent i earns revenue by routing the traffic load d_i over the route bandwidths $\mathbf{b}_i$. Hence $\sum_{i \in \mathcal{N}} u_i(\mathbf{b}_i, d_i)$ is the network's rate of earning revenue.

We introduce an additional commodity called *wealth.* When we refer to commodities however, we refer to the the commodities in the market excluding the commodity wealth. The role of wealth in this model is similar to the role of money in real economic markets. According to [11, p.11], money facilitates trade in three ways: it serves as a medium of exchange; it allows value to be stored; and it allows the value of a commodity to be evaluated by the use of a common measure.

Each agent is *endowed* with an *initial* amount of wealth and a bundle of commodities. Commodities and wealth are traded whenever there is an opportunity to increase the utility of the agents participating in the trade.

There exists an important interplay between wealth and utility. In economic theory, money is usually modeled as a commodity with a positive marginal utility for all the agents in a market. Commodities and wealth in this market model have positive marginal utilities. Let m_i be the amount of wealth owned by agent i. The total wealth in the market is a constant $M = \sum_{i \in \mathcal{N}} m_i$. Let n be the maximum number of distinct commodities per commodity class. We extend the utility function $u_i(\cdot)$ of Chapter 3 and define a utility function $\widetilde{u}_i(\cdot)$ which includes wealth as a parameter

$$\widetilde{u}_i(\mathbf{b}_i, m_i, d_i) \quad = \quad u_i(\mathbf{b}_i, d_i) + m_i. \tag{4.1}$$

$\widetilde{u}_i(\mathbf{b}_i, m_i, d_i)$ is bounded for bounded values of $\mathbf{b}_i$, m_i and d_i. From Equation (4.1) and for all $j \in \{1, \ldots, n\}$

$$\begin{aligned} D_j \widetilde{u}_i(\mathbf{b}_i, m_i, d_i) &= D_j u_i(\mathbf{b}_i, d_i) \\ \partial \widetilde{u}_i(\mathbf{b}_i, m_i, d_i)/\partial m_i &= 1. \end{aligned}$$

The marginal utility $D_j \widetilde{u}_i(\mathbf{b}_i, m_i, d_i)$ of a commodity j of class i is thus independent of the wealth m_i and the marginal utility of wealth is equal to 1 for each agent $i \in \mathcal{N}$.

4.2 Money and wealth

Wealth and money are not synonymous. An agent in a real world market has utility for money because money has the potential to be exchanged for commodities. If money cannot be exchanged for commodities, the money will lose its value and the agent will have no utility for money. Money therefore loses all its value when no trades can take place in the market.

In the real world economy there is a constant exchange of commodities from creation (resources that are mined for example), transformation of commodities (manufacturing) up to consumption or destruction. There is therefore a constant flow of commodities through the market. The money in the market will always have some value – given that there are commodities to be traded for money.

Consider a market with only commodities that are indestructible and limited in quantity. There is always the risk that there will be no supply of commodities. This will happen if none of the owners of commodities wishes to sell. If there is no supply of a commodity, money will not be exchanged for that commodity. In such a case money is rendered valueless. The periods of no supply may be erratic and the agents owning money will receive no benefits for owning money during these periods.

Thus in a market with only indestructible commodities in limited quantity, commodities can only be exchanged for other indestructible commodities in limited quantity. This is true, only if we value commodities objectively. A rare diamond is both limited in quantity and virtually indestructible, but it has only aesthetic value. The monetary value of a diamond is the aesthetic value proclaimed by the highest bidder.

In Chapter 3 we stated that the commodities in our market models are indestructible and limited in quantity. The money used in our market therefore has the characteristics of an indestructible commodity in limited quantity in order to be exchangeable for commodities. Thus we use the term "wealth" rather than "money". Here, *wealth is the rate of earning money.* The initial wealth endowment is therefore a rate of earning money allocated to an agent. If for a certain time period, the agent keeps this wealth, he receives a benefit. This benefit is the utility of the wealth. An agent may exchange its wealth for commodities and by doing so exchanges its wealth benefits for the benefits of owning the commodities.

Wealth is indestructible and allocated to the market agents as a fixed total quantity M. If no trades can take place in the market, the agents still receive benefits for keeping their wealth.

4.3 Trading

Agents may trade commodities and wealth amongst each other, subject to the constraints on trading defined in Chapter 3. The concept of *buying a quantity of commodity* refers to the trading of an amount of wealth for a quantity of commodity, the concept of *selling a quantity of commodity* refers to the trading of a quantity of commodity for an amount of wealth. When referring to *trading a quantity of commodity* or *trading* in general we refer to either buying a quantity of commodity or selling a quantity of commodity.

The trading of commodities or setting of a price in the market is an instantaneous event at any time $t \in \mathbb{R}_+$. At any time t the goal of an agent i is to maximise the utility $\widetilde{u}_i(\mathbf{b}_i^t, m_i^t, d_i^t)$ of its commodity bundle $\mathbf{b}_i^t$, wealth m_i^t and commodity requirement d_i^t. A specialist sets buying and selling prices in order to increase its wealth and thereby its utility. Agents increase their utilities by buying basic commodities to produce composite commodities and decomposing composite commodities to sell basic commodities.

4.3.1 The specialist and prices

According to the basic market model of Chapter 3, the basic commodity of class k may only be bought from the specialist k or sold to specialist k. Each specialist k sets a buying and a selling unit price for the basic commodity of class k.

At time t the specialist k

- sells quantities of basic commodity of class k at a selling price p_k^t,
- and/or buys quantities of basic commodity of class k at a buying price $\tilde{p}_k^t$,
- or does nothing.

For any agent i the *buying price* c_{ij}^t and *selling price* $\tilde{c}_{ij}^t$ of a unit of composite commodity number j

of class i is then

$$c_{ij}^t = \sum_{k\in\mathcal{K}_{i,j}} p_k^t \tag{4.2}$$

$$\tilde{c}_{ij}^t = \sum_{k\in\mathcal{K}_{i,j}} \tilde{p}_k^t \tag{4.3}$$

respectively, where

$$\mathcal{K}_{i,j} = \left\{k : \left[P_i^{-1}(\mathbf{e}_j)\right]_k = 1\right\}$$

is the set containing the classes of the basic commodities required to produce a composite commodity number j of class i as defined in Definition 3.13.

4.3.2 The agents, productions and decompositions

At time t let $r_{ij}^t > 0$ denote the production of a quantity r_{ij}^t of composite commodity j of class i and $s_{ij}^t > 0$ denote the decomposition of a quantity s_{ij}^t of composite commodity number j of class i. Let $r_{ij}^t = 0$ or $s_{ij}^t = 0$ denote the production or decomposition of zero quantity of composite commodity number j of class i at time t.

An agent $i \in \mathcal{N}$ producing a quantity r_{ij}^t of a composite commodity number j of class i buys a quantity r_{ij}^t of basic commodity of class $k \in \mathcal{K}_{i,j}$ from each specialist $k \in \mathcal{K}_{i,j}$ at time t. The total amount of wealth paid to the specialists is $r_{ij}^t c_{ij}^t$, where c_{ij} is defined in Equation 4.2.

An agent i decomposing a quantity s_{ij}^t of a composite commodity number j of class i sells a quantity s_{ij}^t of basic commodity of class $k \in \mathcal{K}_{i,j}$ to each specialist $k \in \mathcal{K}_{i,j}$ at time t. The total amount of wealth received from the specialists is $s_{ij}^t \tilde{c}_{ij}^t$, where $\tilde{c}_{ij}$ is defined in Equation 4.3.

At time t the quantity of basic commodity owned by the specialist k increases (or decreases if negative) by $\sum_{j:k\in\mathcal{K}_{i,j}} (s_{ij}^t - r_{ij}^t)$ and the wealth increases by $\sum_{j:k\in\mathcal{K}_{i,j}} (r_{ij}^t p_k^t - s_{ij}^t \tilde{p}_k^t)$ as result of trading with the agent i. Hence, at time t the quantity of basic commodity owned by the specialist k increases (or decreases if negative) by

$$z_k^t = \sum_{i\in\mathcal{N}-\{k\}} \sum_{j:k\in\mathcal{K}_{i,j}} (s_{ij}^t - r_{ij}^t)$$

and the wealth increases by

$$\tilde{z}_k^t = \sum_{i\in\mathcal{N}-\{k\}} \sum_{j:k\in\mathcal{K}_{i,j}} (r_{ij}^t p_k^t - s_{ij}^t \tilde{p}_k^t)$$

as result of trading with all the agents $i \in \mathcal{N} - \{k\}$.

Thus given the prices $\mathbf{C}^t$ and $\tilde{\mathbf{C}}^t$, at time t each agent (and specialist) $i \in \mathcal{N}$ maximises

$$\widetilde{u}_i\left(\mathbf{b}_i^t + z_i^t \mathbf{e}_1 + \sum_j (r_{ij}^t - s_{ij}^t)\mathbf{e}_j,\ m_i^t + \tilde{z}_i^t - \sum_j r_{ij}^t c_{ij}^t + \sum_j s_{ij}^t \tilde{c}_{ij}^t,\ d_i^t\right)$$

by choosing to produce the quantity $\mathbf{r}_i$ or decompose the quantity $\mathbf{s}_i$ of composite commodities and setting the selling price p_i^t and buying price $\tilde{p}_i^t$ of the basic commodity of class i, subject to the constraints $m_i^t \geq 0, b_{ij}^t \geq 0$ for all i, j. Let $(\mathbf{r}_i^t, \mathbf{s}_i^t, p_i^t, \tilde{p}_i^t)$ denote the *action* of the agent i at time t.

4.4 The market structure

In this section we discuss the structure of the dual-oligopoly.

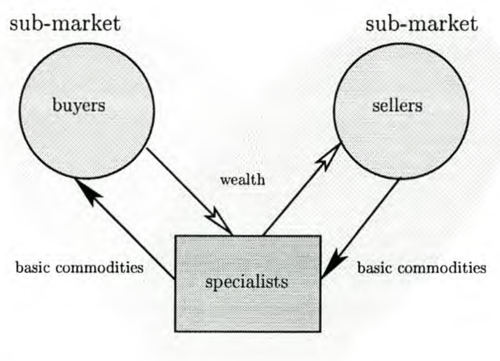

Figure 4.1: The sub-markets in the dual-oligopoly

A specialist in the dual-oligopoly can be viewed as a mediator between two *sub-markets*. Figure 4.1 illustrates the sub-markets in the dual-oligopoly. In the first sub-market the buyers are all the market agents to which the specialists sell quantities of basic commodities. A buyer agent in this sub-market has a choice of composite commodities to produce. However there are few specialists from which it can buy basic commodities. The demand for basic commodity of class k is the market demand (see Definition 3.17) for basic commodity of class k and the supply of basic commodity of class k is the specialist supply (see Definition 3.18) of basic commodity of class k.

In the second sub-market the sellers are all the market agents from which the specialists buy quantities of basic commodities. A seller agent in this sub-market has a choice of composite commodities to decompose. However there are few specialists to which it can sell basic commodities. The supply of basic commodity of class k is the market supply of basic commodity of class k and the demand of basic commodity of class k is the specialist demand of basic commodity of class k.

The specialists set the unit price for buying and selling in the respective sub-markets. These sub-markets are interdependent, each specialist buys basic commodities from the second sub-market to sell in the first sub-market. The agents in each sub-market consist of all the agents in the dual-oligopoly. These sub-markets are identical except for the direction of trading (selling or buying), thus we consider only the first sub-market in the following discussion.

The distinct characteristic [12, p.5] of an oligopoly is that the decisions of each of the few supplying agents in an oligopoly has an observable effect on any agent. The decisions of the agents in each of these sub-markets has an observable effect on the other agents. Each of these sub-markets is an oligopoly.

There is a separate market demand curve (see Section 3.1.5) for each specialist in the first sub-market. We demonstrate this by observing the effect of a change in the selling price p_k of basic commodity of class k.

Consider a specialist k that increases the selling price p_k of the basic commodity of class $k \in \mathcal{K}_{i,j}$. This will cause the price c_{ij}^t for producing a unit of composite commodity number j of class i to increase. The composite commodities that do not have the basic commodity of class k in their production bundles will thus become cheaper relative to composite commodities that do have the basic commodity of class k in their production bundles. Agents will buy the cheaper basic commodities in order to increase their utilities. It is clear that the market demand for basic commodities of class $i \in \mathcal{N} - \{k\}$ will increase and the market demand for basic commodities of class k will decrease.

Consider a specialist k that decreases the selling price p_k of the basic commodity of class $k \in \mathcal{K}_{ij}$. This will cause the price c_{ij}^t for producing a unit of composite commodity j of class i to decrease. The composite commodities that have the basic commodity of class k in their production bundles will thus become cheaper relative to composite commodities that do not have the basic commodity of class k in their production bundles. Agents will buy the cheaper basic commodities in order to increase their utilities. It is clear that the market demand for basic commodities of class $i \in \mathcal{N} - \{k\}$ will increase and the market demand for basic commodities of class k will decrease.

The market demand for the basic commodity of each specialist $k \in \mathcal{N}$ is therefore a function of all the selling prices $\mathbf{p}^t$. In particular

- the market demand for basic commodity of class k is negatively correlated with the selling price p_k^t set by the specialist k, and
- the market demand for basic commodity of class k is positively correlated with the selling prices p_i^t set by the specialists $i \in \mathcal{N} - \{k\}$.

4.5 Agent behaviour

Chamberlin's model of a monopolistic competition [12, p.56-57] explains the behaviour of agents with market demand curves that are dependent as explained in the previous section.

In Figure 4.2 the curve dd' represents the increased sales which a specialist may enjoy by lowering its selling price. Let the current price and quantity sold be p_1 and q_1. A specialist will lower the selling price of the basic commodity it sells to p_2 so that the market demand for this basic

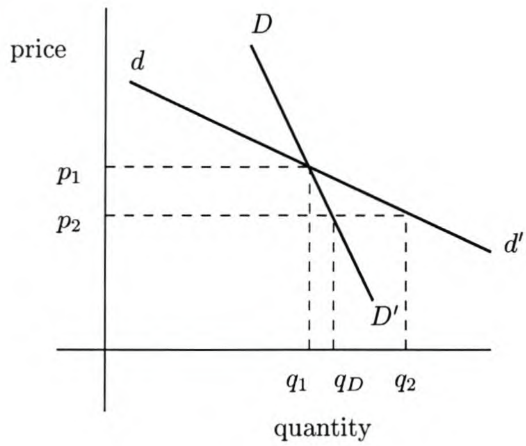

Figure 4.2: Chamberlin's modeling of demand.

commodity increases. The specialist does so under the assumption that the market demand for the basic commodities of other specialists will only decrease slightly. This decrease in market demands will be small enough not to provoke a decrease of the prices of the other basic commodities. The specialist will thus sell a larger quantity q_2 of basic commodity at a lower price p_2.

However each of the specialists in the market will behave in this way, which will result in a general decrease of selling prices. In Figure 4.2 the curve DD' shows the actual sales of the specialist as the general decrease of selling prices takes place. The specialist will therefore sell a quantity of basic commodity q_D only slightly larger than q_1 at a decreased price p_2.

It is clear that there is some price competition in each sub-market.

4.6 Social welfare

Social welfare refers to the joint utility of a group of agents. Let

$$\Theta_t = (\mathcal{N}, \widetilde{\mathbf{u}}(\cdot), \mathbf{B}^t, \mathbf{p}^t, \tilde{\mathbf{p}}^t, \mathbf{m}^t, \mathbf{d}^t, \mathbf{q}^t, M)$$

be a market state (see Definition 3.15) with $\mathcal{N}$ agents, agent utility functions $\widetilde{\mathbf{u}}(\cdot)$, a composite commodity allocation $\mathbf{B}^\mathbf{t}$, basic commodities selling prices $\mathbf{p}^t$ and buying prices $\tilde{\mathbf{p}}^t$, a wealth allocation $\mathbf{m}^t$, commodity requirements $\mathbf{d}^t$, basic commodities $\mathbf{q}^t$ and total wealth M. In this model, the social welfare is defined as the sum of the utilities of all the agents in the market

$$W(\Theta_t) = \sum_{i \in \mathcal{N}} \widetilde{u}_i(\mathbf{b}_i^t, m_i^t, d_i^t). \tag{4.4}$$

Social welfare functions are classified as Pareto or non-Pareto according to the *Pareto value judgement* [13, p.63]. The Pareto value judgement derives from the concept of Pareto efficiency. In order to define the Pareto value judgement we present a definition of Pareto efficiency.

4.6.1 Pareto efficiency

Definition 1.1 presented an informal definition of domination between combinations of player strategies. The following is a formal definition of domination.

Definition 4.1 *Let Γ be an N-player game in normal form with players $\mathcal{N} = \{1, \ldots, N\}$, player strategy sets $S_1, S_2, \ldots, S_N$ and payoff functions $p_1(\cdot), p_2(\cdot), \ldots, p_N(\cdot)$. Let $\succ$ be a binary relation on the set $S = S_1 \times \cdots \times S_N$ such that for $s \in S$ and $s^* \in S$, $s \succ s^*$ implies $p_j(s) \geq p_j(s^*)$ for all $j \in \mathcal{N}$ and $p_j(s) > p_j(s^*)$ for at least one $j \in \mathcal{N}$. The binary relation $\succ$ is non-commutative and transitive and we say that s^* is* **dominated** *by s if $s \succ s^*$.*

The following is a formal definition of a Pareto efficient combination of player strategies (see the informal Definition 1.2).

Definition 4.2 *Let Γ be an N-player game in normal form with players $\mathcal{N} = \{1, \ldots, N\}$, player strategy sets $S_1, S_2, \ldots, S_N$ and payoff functions $p_1(\cdot), p_2(\cdot), \ldots, p_N(\cdot)$. An N-tuple $s^* \in S$ where $S = S_1 \times \cdots \times S_N$ is a Pareto efficient N-tuple if there is no N-tuple $s \in S$ such that $s \succ s^*$.*

4.6.2 The Pareto value judgement

The Pareto value judgement states that welfare is increased if one or more agents are made better off and no agent becomes worse off in terms of utility.

Definition 4.3 *For Γ an N-player game in normal form with players $\mathcal{N} = \{1, \ldots, N\}$, player strategy sets $S_1, S_2, \ldots, S_N$ with $S = S_1 \times \cdots \times S_N$ and payoff functions $p_1(\cdot), p_2(\cdot), \ldots, p_N(\cdot)$, let $W(\cdot)$ be a welfare function $W : S \to \mathbb{R}_+$. Then the Pareto value judgement is the statement: For any pair of N-tuples $s_x, s_y \in S$, if $s_x \succ s_y$ then $W(s_x) > W(s_y)$.*

A welfare function is classified as a *Pareto welfare function* if the Pareto value judgement holds for that welfare function, otherwise it is a *non-Pareto welfare function.*

Figure 4.3 depicts *welfare indifference curves* of a Pareto and a non-Pareto welfare function in a society of two agents with utilities u_1 and u_2 respectively. All the points on a welfare indifference curve have the same social welfare value. The welfare indifference curves of the Pareto welfare function are SWP_1 and SWP_2 and the welfare indifference curves of the non-Pareto welfare function are SWR_1 and SWR_2. The relations $p_2 > p_1$ and $r_2 > r_1$ apply to the social welfare values on the welfare indifference curves in Figure 4.3.

Note that for the non-Pareto welfare function when moving from T to T', thus keeping the utility of agent 2 constant while increasing the utility of agent 1, initially increases the social welfare but eventually leads to a decrease in social welfare. This is not the case with the Pareto welfare function. Similarly for the non-Pareto welfare function when moving from P to P', thus keeping

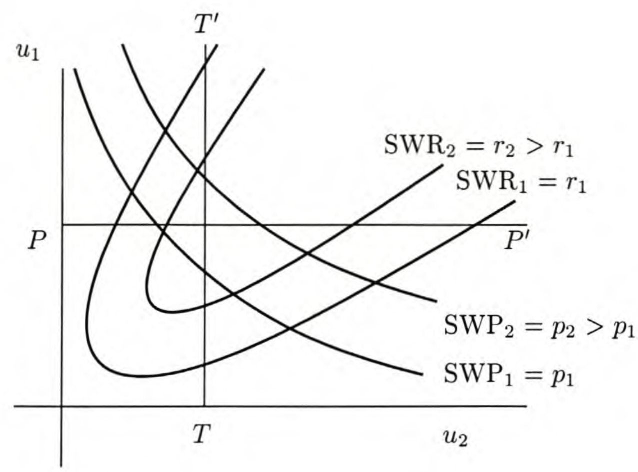

Figure 4.3: Welfare indifference curves of a Pareto and a non-Pareto social welfare function.

the utility of agent 1 constant while increasing the utility of agent 2, initially increases the social welfare but eventually leads to a decrease in social welfare.

In a society where social welfare is defined by a Pareto welfare function, increasing the utilities of some agents while keeping the utilities of the others constant cannot reduce the social welfare.

It is relatively easy to see a positive association between the individual's choice and that which is best for society when observing a society with a Pareto social welfare function. In contrast, if social welfare is defined by a non-Pareto welfare function it is not clear whether the individual's choice reflects that which is best for society. In some sense a Pareto welfare function is inferior to a non-Pareto welfare function. If certain groups of agents become better off and leave others behind, it creates divisions between the groups in a society. Such differences may be regarded as undesirable[1]. A Pareto social welfare function does not allow that possibility.

The social welfare function $W(\cdot)$ of the dual-oligopoly is classified as a Pareto welfare function. Chapter 5 presents a market model of a society with a non-Pareto welfare function.

4.6.3 Equilibrium and social welfare

For this section we assume that the commodity requirement $\mathbf{d}^t = \mathbf{d}$ and basic commodity allocation $\mathbf{q}^t = \mathbf{d}$ are constants with respect to time t.

A market equilibrium (Definition 3.14) occurs whenever no agent in the market chooses or is able to trade commodities. We say that the market is in a strategic equilibrium at a market state Θ_t if the actions of the agents in the market renders the basic commodity prices unchanged and no trades take place in the market. It is clear that if a strategic equilibrium occurs at a market

[1] In a network it may be undesirable to have seriously degraded service for traffics between some O-D pairs, while traffics in the rest of the network receive top-grade service.

state Θ_t then a market equilibrium occurs.

If the agents in the market are able to collaborate then a subgroup of agents may jointly choose their actions in order to increase the utility of the subgroup. The strategic equilibrium is then a strong Nash equilibrium [27, p.26]. However if cooperation is impossible a strategic equilibrium occurs when no agent has incentive to unilaterally deviate from its (conjectured) equilibrium strategy.

Definition 4.4 *Let Γ be an N-player game in normal form with players $\mathcal{N} = \{1, \ldots, N\}$, player strategy sets $S_1, S_2, \ldots, S_N$ and payoff functions $p_1(\cdot), p_2(\cdot), \ldots, p_N(\cdot)$. An N-tuple $(s_1^*, \ldots, s_N^*)$ with $s_i^* \in S_i$ is a Nash equilibrium N-tuple if for every j and for every $s_j \in S_j$*

$$p_j(s_1^*, \ldots, s_j, \ldots, s_N^*) \leq p_j(s_1^*, \ldots, s_j^*, \ldots, s_N^*).$$

A Nash equilibrium is not necessarily unique nor does its occurrence imply that a market equilibrium which maximises social welfare occurs.

4.7 Mechanism design

The dual-oligopoly is a game constituted by a mechanism design problem and a mechanism design solution as presented in Section 1.2. The mechanism presented in the Chapter is only a framework for building a detailed market mechanism.

For a given commodity requirement $\mathbf{d}^t = \mathbf{d}$, a quantity of basic commodity $\mathbf{q}^t := \mathbf{q}$ and an amount of wealth M, the mechanism design problem of the dual-oligopoly is $\mathrm{O} = (\mathcal{N}, \Im, U, \widetilde{W})$ where

- $\mathcal{N}$ is the set of agents,
- the set of outcomes is

  $$\Im = \left\{ (\mathcal{N}, \tilde{\mathbf{u}}(\cdot), \mathbf{B}, \mathbf{p}, \tilde{\mathbf{p}}, \mathbf{m}, \mathbf{d}, \mathbf{q}, M) \in \widetilde{\Im} : \sum_{r \in \mathcal{N}} [P_r^{-1}(\mathbf{b}_r)]_i = q_i, \ \sum_{i \in \mathcal{N}} m_i = M \right\},$$

  where

  $$\widetilde{\Im} = \left\{ (\mathcal{N}, \tilde{\mathbf{u}}(\cdot), \mathbf{B}, \mathbf{p}, \tilde{\mathbf{p}}, \mathbf{m}, \mathbf{d}, \mathbf{q}, M) : \mathbf{B} \in \mathbb{N}_0^R, \ \mathbf{p} \in \mathbb{R}^N, \ \tilde{\mathbf{p}} \in \mathbb{R}^N, \ \mathbf{m} \in \mathbb{R}_+^n \right\},$$

  with $N = |\mathcal{N}|$ and R the total number of distinct commodities in the market,
- $U = U_1 \times \cdots \times U_N$ where U_i is the set of utility functions for each agent $i \in \mathcal{N}$. Each $\tilde{u}_i \in U_i$ satisfies the requirements for utility functions presented in Section 4.1, and
- $\widetilde{W} : U \to \mathbb{P}(\Im)$ is a function mapping the agent utilities to subsets of outcomes, namely those desired by the mechanism designer. In particular

  $$\widetilde{W}(\mathcal{N}, \tilde{\mathbf{u}}(\cdot), \mathbf{d}, \mathbf{q}, M) = \{\Theta \in \Im : W(\cdot) \text{ has a local maximum at } \Theta \text{ with respect to } \mathbf{B} \text{ and } \mathbf{m}\}.$$

Given the mechanism design problem $\text{O} = (\mathcal{N}, \Im, U, \widetilde{W})$, we define a mechanism $(\mathcal{A}, \mu)$ for O where

- $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_N$ where $\mathcal{A}_i \subseteq \mathbb{N}_0^n \times \mathbb{N}_0^n \times \mathbb{R}_+ \times \mathbb{R}_+$ is the set of actions available to an agent $i \in \mathcal{N}$ and n the maximum number of distinct commodities of class i. An action $(\mathbf{r}_i, \mathbf{s}_i, p_i, \tilde{p}_i) \in \mathcal{A}_i$ represents buying and producing an amount r_{ij} and decomposing and selling an amount s_{ij} of composite commodity number j of class i. p_i and $\tilde{p}_i$ are the selling and buying unit prices of basic commodity of class i set by the specialist i. $\mathbf{A} \in \mathcal{A}$ is then an action profile, and
- $\mu : \mathcal{A} \times \Im \to \Im$ maps each action profile and outcome pair $(\mathbf{A}^t, \Theta_t) \in \mathcal{A} \times \Im$ onto an outcome $\Theta_{t+\epsilon} \in \Im$, where $\Theta_{t+\epsilon}$ is the first market state where an agent trades commodities or changes a basic commodity price after time t.

The mechanism design problem and the mechanism together constitute a super game. Note that the action profile $\mathbf{A}^t \in \mathcal{A}$ represents a sequence of market trades and setting of prices directing the transition from the market state Θ_t to the market state $\Theta_{t+\epsilon}$.

The dual-oligopoly is a super game Γ consisting of the single period games $\Gamma_t = (\mathcal{N}, \mathcal{A}, (\widetilde{u}_i)_{i\in\mathcal{N}})\,, t \in \mathbb{R}_+$ where $\mathbf{A}^t \in \mathcal{A}$ are the strategy choices of the players, $(\mathcal{N}, \widetilde{\mathbf{u}}(\cdot), \mathbf{B}^{t+\epsilon}, \mathbf{p}^{t+\epsilon}, \widetilde{\mathbf{p}}^{t+\epsilon}, \mathbf{m}^{t+\epsilon}, \mathbf{d}, \mathbf{q}, M) = \Theta_{t+\epsilon} = \mu(\mathbf{A}^t, \Theta_t)$ is the outcome and $\widetilde{u}_i(\mathbf{b}_i^{t+\epsilon}, m_i^{t+\epsilon}, d_i), i \in \mathcal{N}$ the player payoffs of the single period game Γ_t. The super game payoff of a player $i \in \mathcal{N}$ is $\int_0^\infty \widetilde{u}_i(\mathbf{b}_i^t, m_i^t, d_i)dt$.

Market equilibrium occurs at time t when an action profile $\mathbf{A}^t \in \mathcal{A}$ is a strategic equilibrium of the constituent game Γ_t.

4.8 Conclusion

The dual-oligopoly has a market structure that conforms closely to some market models in economic literature. We define a market mechanism for the dual-oligopoly that places few restrictions on the communications and actions of the market agents. A market mechanism and market agent strategy for the dual-oligopoly may render an *autonomous multi-agent system* that does well in maximising network welfare.

Chapter 5

The social community

In this Chapter we build a market model using the definitions in Chapter 3. The market model is named the *social community* because it resembles a society of individuals. The market agents are members of the social community and are bound by the rules of the social community. These rules define the mechanism for distributing resources in the market and essentially protect the welfare of the individual agents while enhancing the welfare of the group. In order to simplify the market model we assume that the agents are unable to break any of the rules of the social community.

The utility of an agent is a measure of how well the agent does with a commodity bundle in meeting the demand for commodities. The *goal of the social community* is to maximise the individual agent utilities. The agents realise the goal of the social community by trading commodities.

A pseudo price is assigned to each basic commodity. These prices are calculated periodically and communicated to the market agents. An agent chooses the composite commodities to produce or decompose, based on the prices of the basic commodities.

We develop a notion of *fairness* for each agent. The notion of fairness dictates the amount of commodities that agents trade and ensures that agents obtain commodities in proportion to their commodity requirements. An agent with a large requirement for commodities will obtain a larger amount of commodities than an agent with a small requirement for commodities. An agent can evaluate this measure of fairness by using information about the utilities and commodity bundles of the agents with which it trades.

The pseudo price of a basic commodity acts as a penalty for agents using that basic commodity to produce composite commodities. These penalties control the use of the basic commodities in productions.

5.1 The agent utility

Let $\mathcal{R}_i$ denote the index set of the commodities of class i. For the social community, the utility function $u_i(\mathbf{b}_i, d_i)$ is defined as the probability that a call is carried by the OD-pair whose routes are managed by the agent i. The supremum utility of an agent is therefore 1 and for any $d_i \in \mathbb{R}_+$, $\sup_{\mathbf{b}_i}\{u_i(\mathbf{b}_i, d_i)\} = 1$. Ordinarily, for a utility close to this supremum utility, the marginal utility $D_j u_i(\mathbf{b}_i, d_i)$, $j \in \mathcal{R}_i$ of the commodity bundle $\mathbf{b}_i$ becomes insignificant.

Definition 5.1 *For a given commodity bundle $\mathbf{b}_i$ an agent i prefers a lower commodity requirement d_i over a higher commodity requirement, hence the marginal utility $\partial u_i(\mathbf{b}_i, d_i)/\partial d_i$ of the commodity requirement d_i is negative.*

Furthermore we assume that the network clients using commodities of class i are indifferent to the particular commodity number of class i which they use:

Assumption 5.1 *For each agent $i \in \mathcal{N}$ in the social community the marginal utility of a commodity number j of class i is equal to the marginal utility of a commodity number k of class i. Thus $D_j u_i(\mathbf{b}_i, d_i) = D_k u_i(\mathbf{b}_i, d_i)$ for any $\mathbf{b}_i$, d_i, $i \in \mathcal{N}$ and $j, k \in \mathcal{R}_i$ where $\mathcal{R}_i$ denotes the index set of the commodities of class i.*

5.2 Public resources and shadow prices

We introduce public resources and *shadow prices* by presenting an example from the world economy concerning the pricing of a clean environment. A clean environment is a valuable *public resource*; not only in a qualitative sense, but society is prepared to pay a price in order to keep the environment clean. For instance, first world countries pay a price to dispose of their nuclear waste in the territories of developing countries. Both first world and developing countries value a clean environment. However in this example of nuclear waste, a developing country has more urgent problems, of which keeping a clean environment is the lesser important[1]. In terms of demand and supply some agents in society demand the right to use (pollute) the environment and society supplies the right to use the environment. It is difficult to define the supply of a clean environment, because a clean environment is not the property of a single agent in society. In economic terms a clean environment is defined as a *public resource.* The unique characteristic of a public resource is that it may be consumed by one agent without hindering the ability of the resource to be consumed by other agents [13, p.222].

Definition 5.2 *A public resource may be used by one agent without hindering the ability of the resource to be used by other agents.*

[1]Nuclear waste is a long term threat to a clean environment. Developing countries are more concerned with solving the short term problems of society. It is therefore attractive for developing countries to allow the regulated disposal of nuclear waste in their territories.

The pollution by one agent does not hinder the ability of another agent to pollute. It is therefore virtually impossible to quantify the supply of a clean environment. The supply and demand of a clean environment can thus not meet in a market place.

It is important to be able to quantify the supply of a clean environment, because there is an increasing *social cost* [2] for every additional agent polluting the environment. The concept of *shadow prices* provides a solution to the problem of quantifying supply. Shadow prices are prices calculated as if there were a market supply of a clean environment. The essential idea of shadow pricing is to assign pseudo prices to a public resource in such a way that its consumption by one agent does hinder the ability of the public resource to be consumed by others[3].

There exist various techniques to simulate such a market and evaluate the price of a clean environment. Each of these techniques aims to estimate the marginal social cost of the public resource consumed. The marginal social cost of the public resource is then used as a supply curve. Some of these techniques of calculating shadow prices include the cost of removing the pollution and what society is prepared to pay for a clean environment.

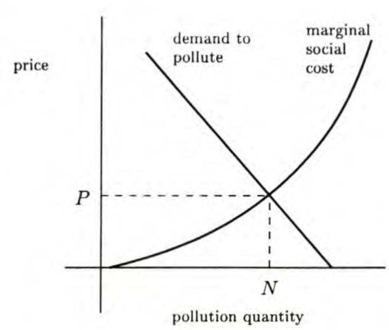

Figure 5.1: Pricing a clean environment.

Figure 5.1 depicts a demand and supply curve representing the demand to pollute and the marginal social cost of pollution. The intersection of the supply curve and demand curve renders the market clearing pollution quantity N and market clearing price P where the marginal cost of pollution equates the marginal cost polluters are prepared to pay. The shadow prices therefore act as a penalty for agents using the resources subjected to shadow pricing.

In the nuclear waste example, the shadow price for disposing nuclear waste in developing countries is substantially lower, since first world societies are prepared (and able) to pay much more for a clean environment[4] than if the waste were disposed in their own territories.

[2] Increasing social cost is detrimental to the well-being of all the agents in society.

[3] In order for shadow pricing to be successful in reducing the consumption of a public resource, the demand of that public resource must be elastic. In other words the agents in the society must be able to choose the amount of that public resource they consume.

[4] Pollution is threatening the cleanness of the global environment. Developing countries blame first world countries for using the global environment without proper compensation to their societies. First world countries are pressured to agree on environmental shadow prices [31] that will reduce pollution, to the detriment of their economies.

5.3 The price of commodities

Similar to shadow prices for public resources we define a price for basic commodities. This price is based on the amount of available basic commodities and the requirement for commodities.

Let $\Theta = (\mathcal{N}, \mathbf{u}(\cdot), \mathbf{B}, \mathbf{d}, \mathbf{q})$ be a *market state* with a set of market agents $\mathcal{N}$, a composite commodity allocation $\mathbf{B}$, commodity requirements $\mathbf{d}$, basic commodities $\mathbf{q}$ and utility functions $\mathbf{u}(\cdot)$ (see Section 3.1.2). In this section we define a price s_i for each basic commodity of class $i \in \mathcal{N}$ at the market state Θ. Each composite commodity number $j \in \mathcal{R}_i$ of class $i \in \mathcal{N}$ is priced according to the prices for basic commodities.

5.3.1 The price of basic commodities

Definition 5.3 *The quantity Q_i of basic commodity of class i supplied for the production of composite commodities of class $j \neq i, j \in \mathcal{N}$ is $Q_i = \sum_{j\neq i, j\in\mathcal{N}} \left[P_j^{-1}(\mathbf{b}_j)\right]_i$.*

According to Definition 5.3 and the constraints $\sum_{j\in\mathcal{N}}[P_j^{-1}(\mathbf{b}_j)]_i = q_i$, for all $i \in \mathcal{N}$, $Q_i = q_i - b_i$ is the quantity of basic commodity of class i supplied to productions and b_i the quantity of basic commodity not available to supply to productions.

The price s_i of supplying basic commodity of class i to productions is defined as

$$s_i = 1 - u_i\left(\left(q_i - Q_i\right)\mathbf{e}_1, d_i\right)$$

which is also the supply curve with s_i the unit price of basic commodity of class i and Q_i the quantity of basic commodity of class i supplied. By evaluating the derivatives

$$\begin{aligned} \partial s_i/\partial Q_i &= D_1 u_i((q_i - Q_i)\,\mathbf{e}_1, d_i) = D_1 u_i(b_i \mathbf{e}_1, d_i) > 0 \\ \partial^2 s_i/\partial Q_i^2 &= -D_1 D_1 u_i((q_i - Q_i)\,\mathbf{e}_1, d_i) = -D_1 D_1 u_i(b_i \mathbf{e}_1, d_i) \geq 0 \end{aligned}$$

we determine the slope and concavity of the supply curve: the supply curve is monotone increasing and convex (Definition 3.5). The supply curve is well defined in a sense that the price of supplying basic commodity increases as the quantity of basic commodity supplied increases.

One expects the supply of basic commodity of class i to decrease if the the client requirement d_i for commodities of class i increases. The derivative

$$\partial s_i/\partial d_i \quad = \partial\left(1 - u_i\left(\left(q_i - Q_i\right)\mathbf{e}_1, d_i\right)\right)/\partial d_i = -\partial u_i(b_i \mathbf{e}_1, d_i)/\partial d_i > 0$$

indicates that an increased requirement d_i increases the unit price at every quantity supplied or equivalently the quantity supplied decreases at every unit price, hence the supply decreases. The opposite is also true, a decrease in the requirement d_i decreases the unit price at every quantity supplied or equivalently the quantity supplied increases at every unit price, hence the supply increases.

The price s_i for the basic commodity of class i characterises a supply curve. In network terms a lower price s_i results in more direct route bandwidth q_i to be used by multi-link routes and *vice versa.*

5.3.2 The cost of a composite commodity

The cost of a composite commodity is a price that represents a penalty for producing a unit of that composite commodity.

Let $\mathbf{C}$ denote a matrix of unit prices of the composite commodities where c_{ij} is the unit price of composite commodity number $j > 1$ of class i. As previously stated, the price of the basic commodity of class i is

$$s_i = 1 - u_i(b_i\mathbf{e}_1, d_i).$$

We define the unit price of a composite commodity as the sum of the prices of the basic commodities in the unit production bundle. The unit price of composite commodity number $j > 1$ of class i is

$$c_{ij} = \sum_{k \in \mathcal{K}_{i,j}} s_k \tag{5.5}$$

with $\mathcal{K}_{i,j}$ as in Definition 3.13. Note that the price $c_{i1} = 0$ for all $i \in \mathcal{N}$.

5.4 The total cost of composite commodities and the utilisation of network resources

The total cost of composite commodities is

$$\begin{aligned}
\mathcal{C}(\Theta) &= \sum_{i \in \mathcal{N}} \mathbf{c}_i \cdot \mathbf{b}_i \\
&= \sum_{i \in \mathcal{N}} \sum_j c_{ij} b_{ij} \\
&= \sum_{i \in \mathcal{N}} \sum_j \left(\sum_{k \in \mathcal{K}_{i,j}} s_k \right) b_{ij} \\
&= \sum_{i \in \mathcal{N}} \sum_j \left(\sum_{k \in \mathcal{K}_{i,j}} (1 - u_k(b_k\mathbf{e}_1, d_k)) \right) b_{ij}.
\end{aligned}$$

An example of a (physical or virtual) network with resources that are fully utilised, is a fully connected network with direct links that have exactly the required bandwidth to carry all the traffics of the network users. Traffics however are variable in volume, origins and destinations and it is impossible to maintain a network with fully utilised resources. Networks are therefore provided with extra bandwidth to meet the uncertainty of traffic demands. If we assume the network can always meet the QoS criteria such as maximum delay, a measure of efficient routing in a network is the utilisation of resources of the network. The lower the utilisation of network resources for carrying a certain volume of traffic, the more efficient the routing scheme.

Traffic that is routed across multiple links increases the utilisation of the network resources. It is therefore more efficient to route traffics over direct links and short routes. The total cost of composite commodities increases whenever composite commodities are produced from basic commodities (a synonym for routing over multiple links) and decreases whenever composite commodities are decomposed (a synonym for routing over single links). In essence the total cost of composite commodities is a measure of efficient routing, where the composite commodity bundle $\mathbf{B}$ is the routing configuration. In market terms the total cost of composite commodity is a measure of an efficient composite commodity allocation (see Definition 3.12). Chapter 6 depicts the total cost of composite commodities in an implementation of the social community.

5.5 Social welfare

This section defines *social welfare*, which is a measure of how well the members of the social community perform with the composite commodity allocation $\mathbf{B}$ at a market state Θ. All the agents in the social community cooperate to increase social welfare.

The utility of each agent contributes to the welfare of the social community. An agent with a higher requirement for commodities contributes more to the social welfare than an agent with the same utility and a lower requirement for commodities. Let the social welfare contribution of an agent i be $u_i(\mathbf{b}_i, d_i)I_i$ where $I_i = d_i / \sum_{j \in \mathcal{N}} d_j$ represents the requirement for commodities of class i expressed as a fraction of the total requirement for commodities. The weighted average agent utility is

$$A(\Theta) = \sum_{i \in \mathcal{N}} u_i(\mathbf{b}_i, d_i)I_i. \tag{5.6}$$

In network terms $A(\Theta)$ is the fraction of the total network traffic carried.

We develop an *objective function* which quantifies the welfare of the agents in the market for a given allocation of resources. Optimising this objective function is analogous to maximising social welfare. It is an intricate process to design mechanisms and agent strategies that optimise the objective function, since there is no global control over the behaviour of the agents and each agent has only limited information about the total agent system. We say that each agent acts on local information only.

Definition 5.4 *The local information available to an agent i is its production function $P_i(\cdot)$, its utility function $u_i(\cdot)$, its composite commodity bundle $\mathbf{b}_i$ and the basic commodity price s_i. The basic commodity prices $s_j, j \in \mathcal{N}$ are calculated periodically and communicated to the agents. The basic commodity prices are therefore classified as local information available to an agent. When an agent initiates a trade with another agent, the utility functions and commodity bundles of the trading participants are also classified as local information available to each trading participant.*

Because of the restricted nature of the agents, the agents themselves cannot evaluate the objective function. The objective function can therefore not serve as an aid for the agents to reach an optimum resource allocation. However we will use the objective function as a performance measure of the social community.

The objective function is called the *social welfare penalty function*. Kuwabara *et al.* [19] minimised the variance of the agent utilities as the social welfare of an agent system. We use the variance of the agent utilities $V(\Theta)$ and the weighted average agent utility $A(\Theta)$ in Equation (5.6) to define the social welfare penalty function

$$W(\Theta) \quad = \quad \alpha_1 \sqrt{V(\Theta)} + \alpha_2 \left(1 - A(\Theta)\right), \tag{5.7}$$

where the constants $\alpha_1, \alpha_2 > 0$ represent the weights of the respective terms contributing to the social welfare penalty,

$$V(\Theta) = \sum_{i \in \mathcal{N}} \left(u_i(\mathbf{b}_i, d_i) - \mu(\Theta)\right)^2 / (N-1)$$

is the variance of the agent utilities and

$$\mu(\Theta) = \sum_{j \in \mathcal{N}} u_j(\mathbf{b}_j, d_j)/N.$$

Note that $0 \leq A(\Theta) \leq 1$.

Increasing the utility of one agent and keeping the utility of the other agents constant may decrease the social welfare. The social welfare function of the social community is therefore a non-Pareto social welfare function (see Definition 4.3).

5.6 The optimisation method

The fraction $A(\Theta)$ of the total network traffic carried can be increased by producing a composite commodity number j of class i if

$$R_{ij} = d_i D_j u_i(\mathbf{b}_i, d_i) - \sum_{k \in \mathcal{K}_{i,j}} d_k D_1 u_k(\mathbf{b}_k, d_k) > 0.$$

Likewise $A(\Theta)$ can be increased by decomposing the composite commodity number j of class i if $R_{ij} < 0$. The problem with this method is that it may be computationally expensive to calculate the optimal amount of composite commodity to produce or decompose. An optimisation algorithm could select the composite commodity number j of class i with the largest value R_{ij} and produce a small amount of composite commodity number j of class i. This process is repeated until $R_{ij} < \epsilon$ for some value $\epsilon > 0$. This method uses a large number of productions and decompositions of small amounts of commodity in an attempt to maximise $A(\Theta)$.

The social community determines that a commodity bundle $\mathbf{b}_i, i \in \mathcal{N}$ contains an excess amount of commodities relative to a commodity bundle $\mathbf{b}_j, j \in \mathcal{N}$ if and only if $u_i(\mathbf{b}_i, d_i) > u_j(\mathbf{b}_j, d_j)$.

Hence an agent i has excess commodities relative to an agent j if and only if agent i does better than agent j in meeting its commodity requirement. The optimisation method used by the social community transfers (produces or decomposes) commodities from agents with an excess amount of commodities to agents with a shortage of commodities. The notion of fairness in the social community ensures that commodities are not transferred from an agent with a shortage of commodities to an agent with an excess amount of commodities.

Although we did not find a mathematical proof that this method optimises $W(\Theta)$, we tested this method in Chapter 6 which applies an implementation of the social community to several test networks. The results compare well with a conventional optimisation method.

5.7 The rules of the social community

In the social community the pricing of basic commodities is a mechanism used by an agent to select the composite commodities to produce or decompose. The rules of the social community stipulate how these prices are calculated (see Section 5.3). Chapter 6 explains how the agents use the prices.

In addition to the pricing mechanism, the rules of the social community restrict the amount of composite commodities an agent produces or decomposes. A trade that satisfies these restrictions on the amount of composite commodity produced or decomposed is called a *fair trade*. Agents may only take part in fair trades.

5.8 Fair trades

This section presents the conditions of a *fair trade*. Note that these conditions restrict the amounts of composite commodities to produce and decompose and do not dictate which composite commodity to produce or decompose.

The notion of a fair trade is derived from a notion of fairness called *maxmin fairness* [28, 29]. Maxmin fairness states that the utility of one agent is not increased at the cost of decreasing the utility of another agent that has a lower or equal utility. Maxmin fairness is attractive because it promotes smaller differences between the agent utilities and the agents therefore do equally well in meeting their commodity requirements. An agent with a large requirement for commodities will acquire a larger amount of commodities than an agent with a smaller requirement for commodities.

Maxmin fairness only considers pairs of agents. When producing or decomposing composite commodities we must consider more than two agents. We extend maxmin fairness to multiple agents:

- If the utility of a single agent is increased at the cost of decreasing the utility of multiple agents, the utility of the single agent must be less than or equal to the utilities of the other agents.

- If the utilities of multiple agents are increased at the cost of decreasing the utility of a single agent, the utility of the single agent must be greater or equal to the utility of at least one of the other agents.

Definition 5.5 *A maxmin fair trade is a production or decomposition that conforms to one of the following:*

1. For a production $P_i(\mathbf{x}) = \mathbf{y}$ with $\mathbf{x}$ the production bundle of the composite commodity number j of class i in the bundle $\mathbf{y} \neq \mathbf{0}$, then $u_k(\mathbf{b}_k - x_k\mathbf{e}_1, d_k) \geq u_i(\mathbf{b}_i + \mathbf{y}, d_i)$ for all $k \in \mathcal{K}_{i,j}$.
2. For a decomposition $P_i^{-1}(\mathbf{y}) = \mathbf{x}$ with $\mathbf{y} \neq \mathbf{0}$ the bundle containing the composite commodity number j of class i to be decomposed and $\mathbf{x}$ the decomposition of the composite commodity number j, then $u_k(\mathbf{b}_k + x_k\mathbf{e}_1, d_k) \leq u_i(\mathbf{b}_i - \mathbf{y}, d_i)$ for at least one $k \in \mathcal{K}_{i,j}$.

Productions may lead to a condition of *dual-stability*. Dual-stability occurs when composite commodities are produced from the basic commodities of agents who themselves posses insufficient quantities of commodities. These agents will then produce composite commodities to replace their basic commodities, which leads to an equilibrium market state where all the agents use composite commodities. To produce a unit of composite commodity requires more than one unit of basic commodity hence the amount of commodities $\sum_j b_{ij}$ of each agent i is reduced. Furthermore the additional utility for a unit of composite commodity is the same as the additional utility for a unit of basic commodity (see Assumption 5.1), hence the allocation of commodities is inefficient.

The condition of dual-stability can be avoided by restricting agents to either supply basic commodities for productions or to produce composite commodities. This restriction is called the *donor-recipient* relationship.

Assumption 5.2 *The donor-recipient relationship applies to the social community. Hence for an agent $i \in \mathcal{N}$ that supplies basic commodity ($Q_i > 0$), the amounts of composite commodities $b_{ij} = 0$ for all $j > 1$.*

When an agent i executes a production of a composite commodity number j of class i each agent $k \in \mathcal{K}_{i,j}$ becomes a supplier ($Q_k > 0$) of basic commodity and from Assumption 5.2 $b_{kr} = 0$ for all composite commodities $r > 1$. Similarly when an agent i executes a decomposition of a composite commodity number j of class i each agent $k \in \mathcal{K}_{i,j}$ is a supplier ($Q_k > 0$) of basic commodity and from Assumption 5.2 $b_{kr} = 0$ for all composite commodities $r > 1$. The commodity bundles $\mathbf{b}_k$ of the agents $k \in \mathcal{K}_{i,j}$ in Definition 5.5 thus reduce to $b_k\mathbf{e}_1$ and we define a fair trade as:

Definition 5.6 *A fair trade is a production or decomposition that conforms to one of the following:*

1. For a production $P_i(\mathbf{x}) = \mathbf{y}$ with $\mathbf{x}$ the production bundle of the composite commodity number j of class i in the bundle $\mathbf{y} \neq \mathbf{0}$, then $u_k\left((b_k - x_k)\,\mathbf{e}_1, d_k\right) \geq u_i(\mathbf{b}_i + \mathbf{y}, d_i)$ for all $k \in \mathcal{K}_{i,j}$.

2. For a decomposition $P_i^{-1}(\mathbf{y}) = \mathbf{x}$ with $\mathbf{y} \neq \mathbf{0}$ the bundle containing the composite commodity number j of class i to be decomposed and $\mathbf{x}$ the decomposition of the composite commodity number j, then $u_k\left((b_k + x_k)\,\mathbf{e}_1, d_k\right) \leq u_i(\mathbf{b}_i - \mathbf{y}, d_i)$ at least one $k \in \mathcal{K}_{i,j}$.

Let

$$\Im_{\mathbf{q}} = \left\{ \mathbf{B} \in \mathbb{N}_0^R : \sum_{j \in \mathcal{N}} [P_j^{-1}(\mathbf{b}_j)]_i = q_i, \text{ for all } i \in \mathcal{N} \right\}$$

denote the set of all possible *discrete-valued* composite commodity allocations where R is the total number of distinct commodities in the market. Let $\mathcal{T}_{\mathbf{B}} \subseteq \Im_{\mathbf{q}}$ be the set of all possible discrete-valued composite commodity allocations resulting from a trade starting with the composite commodity allocation $\mathbf{B}$. Let $\mathcal{T}_{\mathbf{B}}^{\mathcal{Q}} \subseteq \mathcal{T}_{\mathbf{B}}$ be the set of all possible discrete-valued composite commodity allocations resulting from a **fair** trade starting with the composite commodity allocation $\mathbf{B}$.

Definition 5.7 *A composite commodity allocation* $\mathbf{B}$ *is fair and the market said to be in a fair market equilibrium if and only if no fair trades can take place (*$\mathcal{T}_{\mathbf{B}}^{\mathcal{Q}} = \emptyset$*).*

Definition 5.8 *A fair market is a market in which only fair trades take place.*

Chapter 6 shows that starting from a market state Θ_0 the market process $(\Theta_t)_{t\in\mathbb{N}_0}$ converges to a fair market equilibrium in a finite number of fair trades. There is only a finite number of market states, because the amount of commodities is fixed and quantized. Hence by proving that there are no recurrent market states in a fair market, we prove that the market process converges to a fair market equilibrium in a finite number of fair trades.

Proving that there are no recurrent market states in a fair market is equivalent to proving that there are no cycles in the directed graph $(\mathcal{V}, \mathcal{E}) = \left(\Im_{\mathbf{q}},\ \left\{(\mathbf{B}, \mathbf{A}) : \mathbf{B} \in \Im_{\mathbf{q}}, \mathbf{A} \in \mathcal{T}_{\mathbf{B}}^{\mathcal{Q}}\right\}\right)$, where $\mathcal{V}$ is the set of vertexes (the market states) and $\mathcal{E}$ the set of edges (the fair trades) such that $(\mathbf{B}, \mathbf{A}) \in \mathcal{E}$. It may be possible to prove that there are no recurrent market states in a fair market because we are able to specify the set $\mathcal{T}_{\mathbf{B}}^{\mathcal{Q}}$ for any $B \in \Im_{\mathbf{q}}$. Attempts were made to prove this. However we abandoned this course of action and rather showed in Chapter 6 that it is reasonable to suppose that there are no recurrent market states in a fair market.

5.9 Mechanism design

The social community is a game constituted by a mechanism design problem and a mechanism as presented in Section 1.2.

For a given commodity requirement $\mathbf{d}$ and a quantity of basic commodity $\mathbf{q}$ the mechanism design problem of the social community is $\mathcal{O} = (\mathcal{N}, \Im, U, \widetilde{W})$ where

- $\mathcal{N}$ is the set of agents,
- the set of outcomes is
$$\Im = \left\{ (\mathcal{N}, \mathbf{u}(\cdot), \mathbf{B}, \mathbf{d}, \mathbf{q}) : \mathbf{B} \in \mathbb{N}_0^R, \sum_{r \in \mathcal{N}} [P_r^{-1}(\mathbf{b}_r)]_i = q_i \right\}$$
with $N = |\mathcal{N}|$ and R is the total number of distinct commodities in the market,
- $U = U_1 \times \cdots \times U_N$ where U_i is the set of utility functions for each agent $i \in \mathcal{N}$. Each $u_i \in U_i$ satisfies the requirements for utility functions presented in Section 5.1, and
- $\widetilde{W} : U \to \mathbb{P}(\Im)$ is the function mapping the agent utilities to subsets of outcomes, namely those desired by the mechanism designer. In particular
$$\widetilde{W}(\mathcal{N}, \mathbf{u}(\cdot), \mathbf{d}, \mathbf{q}) \subseteq \{\Theta \in \Im : W(\cdot) \text{ has a local minimum at } \Theta \text{ with respect to } \mathbf{B}\}.$$

Given the mechanism design problem $\mathcal{O} = (\mathcal{N}, \Im, U, \widetilde{W})$, we define a mechanism $(\mathcal{A}, \mu)$ for $\mathcal{O}$ where

- $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_N$ where $\mathcal{A}_i \subseteq \mathcal{R}_i \times \mathcal{R}_i$ is the set of actions available to an agent $i \in \mathcal{N}$ with $\mathcal{R}_i$ the index set of the composite commodities of class i. An action $(j, k) \in \mathcal{A}_i$ represents an attempt[5] to execute a fair trade and production of the composite commodity number j of class i and an attempt to execute a fair trade and decomposition of the composite commodity number k of class i. $\mathbf{A} \in \mathcal{A}$ is then an action profile, and
- $\mu : \mathcal{A} \times \Im \to \Im$ maps each action profile and outcome pair $(\mathbf{A}^t, \Theta_t) \in \mathcal{A} \times \Im$ onto an outcome $\Theta_{t+1} \in \Im$. The function μ determines which fair trades take place and the quantity traded.

The mechanism design problem and mechanism together constitute a super game. Note that the action profile $\mathbf{A}^t \in \mathcal{A}$ represents a sequence of fair trades directing the transition from the market state Θ_t to the market state Θ_{t+1}. The auctioning mechanism in Chapter 6 is an implementation of the function μ of the mechanism $(\mathcal{A}, \mu)$. At any market state Θ_t this auctioning mechanism allows multiple fair trades (or actions in the action profile $\mathbf{A}^t \in \mathcal{A}$) to take place simultaneously while excluding any trades that are not fair. The auctioning mechanism also determines the quantities of basic commodities traded.

The social community is a super game Γ consisting of the single period games $\Gamma_t = (\mathcal{N}, \mathcal{A}, (u_i)_{i \in \mathcal{N}})$ with $t \in \mathbb{N}_0$ where $\mathbf{A}^t \in \mathcal{A}$ are the strategy choices of the players, $(\mathcal{N}, \mathbf{u}(\cdot), \mathbf{B}^{t+1}, \mathbf{d}, \mathbf{q}) = \Theta_{t+1} = \mu(\mathbf{A}^t, \Theta_t)$ is the outcome and $u_i(\mathbf{b}_i^{t+1}, d_i), i \in \mathcal{N}$ are the player utilities after the single period game Γ_t ended.

[5]The market state or the actions of other agents may render a fair trade impossible.

Definition 5.9 *In a super game the strategy of a player defines its actions in every possible constituent game.*

The strategy set of a player i in the super game Γ is

$$\mathcal{S}_i = \left\{ \left(\mathbf{a}_i^{\Theta} \in \mathcal{A}_i \right)_{\Theta \in \Im} \right\}$$

where $\mathbf{a}_i^{\Theta} \in \mathcal{A}_i$ is the action of the player i in the constituent game characterised by the outcome (market state) $\Theta \in \Im$. The agent strategy in Chapter 6 is a super game strategy $\mathbf{S}_i^* \in \mathcal{S}_i$ defined for each player $i \in \mathcal{N}$.

Let $\mathbf{S}_i \in \mathcal{S}_i$ be the super game strategy of each player $i \in \mathcal{N}$. Starting at a market state Θ_0 the strategies $\mathbf{S}_i \in \mathcal{S}_i, i \in \mathcal{N}$ determine the market state transitions in the market process $(\Theta_t)_{t \in \mathbb{N}_0}$. The super game payoff of each player $i \in \mathcal{N}$ is then

$$p_i((\mathbf{S}_j)_{j \in \mathcal{N}}) = \sum_{t \in \mathbb{N}} u_i(\mathbf{b}_i^t, d_i).$$

A fair market equilibrium (see Definition 5.7) occurs at a market state Θ_t where there exists no action profile $\mathbf{A}^{t_1} \in \mathcal{A}$ with $t_1 > t$ such that $\Theta_{t_1} \neq \Theta_t$.

As result of a fair trade in a constituent game, an agent may decrease its utility in order to increase the utility of another agent. In context of a constituent game an agent in the social community may appear to be an altruist (see Definition 3.22), however the agent only abides by the rules of the social community (see Section 5.7) and is self-interested (see Definition 3.21).

The next Chapter presents a strategy for agents in the social community and an auctioning mechanism that allows multiple fair trades to take place simultaneously while excluding trades that are not fair.

Chapter 6

An implementation of the social community

This Chapter presents a market mechanism that is an implementation of fair trading, and a strategy for an agent using this market mechanism. We present test results of the market mechanism and the agent strategy employed in the social community. The test results show that the state of the market converges to a fair market equilibrium that minimises the market welfare penalty function.

6.1 The market mechanism and agent strategies

This section presents a market mechanism and a strategy for using the mechanism. The mechanism consists of a *production part* and a *decomposition part* and defines the behaviour of an agent playing the role of a specialist that trades a basic commodity. The strategy defines the behaviour of an agent playing the role of a producer that produces or decomposes composite commodities.

We prove in Appendix A that the mechanism conforms to the fair trade criteria (see Definition 5.6) and that by using the strategy, an agent will always take part in at least one fair trade if it is possible for that agent to take part in any fair trade.

Consider an agent i using the basic commodities of the classes $k \in \mathcal{K}_{i,h}$ to produce a composite commodity number h of class i. For this production to be a fair trade the amount y_i of composite commodity produced is chosen such that

$$u_i(\mathbf{b}_i + y_i\mathbf{e}_h, d_i) \leq u_k((b_k - y_i)\mathbf{e}_1, d_k) \tag{6.8}$$

for all $k \in \mathcal{K}_{i,h}$ (see Definition 5.6). The production part of the market mechanism calculates such values of y_i for all agents $i \in \mathcal{N}$ simultaneously.

Consider an agent i decomposing the composite commodity number h of class i and disposing of the basic commodities of the classes $k \in \mathcal{K}_{i,h}$. For this decomposition to be a fair trade the amount y_i of composite commodity decomposed is chosen such that

$$u_i(\mathbf{b}_i - y_i\mathbf{e}_h, d_i) \geq u_k((b_k + y_i)\mathbf{e}_1, d_k) \tag{6.9}$$

for at least one $k \in \mathcal{K}_{i,h}$ (see Definition 5.6). The decomposition part of the market mechanism calculates such values of y_i for all agents $i \in \mathcal{N}$ simultaneously.

6.1.1 Production part

An agent that attempts to produce a composite commodity places bids for basic commodities at the respective auctioneers. Let $\mathcal{B}_k$ denote the index set of the agents that place bids for the basic commodity of class k. Similar to the WALRAS algorithm [32], a bid is a message to the auctioneer where the message contains information about the utility function, the commodity requirement and the current commodity bundle of the bidder. The bid also contains information specifying the composite commodity number h that the bidder i intends to produce by using the basic commodity of class k. The auctioneers are bound by the rules of the social community and use the information in the bids to ensure that only fair trades take place. An auctioneer $k \in \mathcal{N}$ allocates basic commodity to the bidders $i \in \mathcal{B}_k$. The allocation of y_i units of basic commodity to a bidder i means that the bidder is entitled to use up to y_i units of basic commodity in the production of composite commodity number h of class i.

Mechanism: The auctioneer

The task of the auctioneer is to allocate basic commodities to the bidders. The auctioneer uses the information contained in the bids to determine the amount of basic commodity to be allocated to each bidder. Adhering to the fair trade criteria, the maximum amount of basic commodities is allocated to each bidder. In the social community the auctioning of commodities is a mechanism to distribute commodities.

The auction is classified as a *multi-unit auction* because more than one unit of a commodity is auctioned at the auction. It is also a *private-value auction* since each bidder has its own valuation of the commodity and each bidder is ignorant of the value that the other bidders place on the commodity. A bidder may only place one bid at each auction.

Let the variable y_i denote the total amount of basic commodity allocated to the bidder i. When the auction starts, no basic commodities are allocated to the bidders and $y_i = 0$ for all $i \in \mathcal{B}_k$. During the auction the auctioneer may allocate basic commodities and thus increase the amounts $y_i, i \in \mathcal{B}_k$.

The auction consists of a sequence of *phases* in which basic commodities are allocated to the bidders. At each phase the bidder i with the lowest utility of all the bidders is identified. One unit of basic commodity is allocated to the bidder i and the phase ends. At each phase, the total amount of basic commodity y_i allocated to the bidder i therefore increases by one unit.

Because one unit of basic commodity is allocated to a bidder i at the end of each phase, a new bidder utility must be calculated. According to the information contained in the bid, the bidder i intends to produce the composite commodity number h of class i. Hence the bidder i will use an amount y_i of basic commodity of class k to produce an amount y_i of composite commodity number h. The bidder utility is therefore $u_i(\mathbf{b}_i + y_i\mathbf{e}_h, d_i)$ for $y_i > 0$. At the end of each phase the utility of the bidder i is set to $u_i(\mathbf{b}_i + y_i\mathbf{e}_h, d_i)$.

The auction terminates when further allocation of commodity will result in the utility of a bidder i : $y_i > 0$ exceeding the auctioneer's utility

$$u_k((b_k - \sum_{j\in\mathcal{B}_k} y_j)\mathbf{e}_1, d_k)$$

of the basic commodity of class k. The auction therefore terminates at the phase where another commodity allocation would falsify one of the relations

$$u_i(\mathbf{b}_i + y_i\mathbf{e}_h, d_i) \leq u_k((b_k - \sum_{j\in\mathcal{B}_k} y_j)\mathbf{e}_1, d_k) \tag{6.10}$$

for any $i \in \mathcal{B}_k$ such that $y_i > 0$. Note that the auction termination condition in Relations (6.10) implies that the Relations (6.8) are satisfied.

When the auction terminates the auctioneer sends a message to each bidder $i \in \mathcal{B}_k$. Each message contains information about the total amount y_i of basic commodity allocated to the bidder i by the auctioneer k.

Strategy: The bidder

The task of the bidder is to choose which composite commodity to produce. The agent strategy is defined such that an agent avoids producing a composite commodity that requires the use of an expensive basic commodity. Consider a price vector $\mathbf{a} = (5, 5, 1, 9)$, denoting the basic commodity prices in the production bundle of some composite commodity. An agent prefers to produce a composite commodity with a price vector $\mathbf{b} = (5, 5, 5, 5)$ rather than a price vector $\mathbf{a}$ as explained below.

We define a *price index* θ_{ij} for a composite commodity number j of class i which is a function of the price vector of the composite commodity j. The price index of a composite commodity number j of class i is

$$\theta_{ij} = \max_{k\in\mathcal{K}_{i,j}} s_k \tag{6.11}$$

where s_k is the price of the basic commodity of class k. The price index of a composite commodity is the price of the most expensive basic commodity in the production bundle of that composite commodity. An agent i places bids for the production of the composite commodity with the lowest price index θ_{ij}.

Let y_{ki} be the total amount of basic commodity allocated to the bidding agent i by each auctioneer k. The agent i uses an amount y_{ki} or less of each basic commodity $k \in \mathcal{K}_{i,j}$ in the production of a composite commodity. In order to produce an amount λ of a composite commodity number j of class i, an agent i needs λ amounts of basic commodities of each class $k \in \mathcal{K}_{i,j}$ (see Assumption 3.2). The agent i maximises utility by producing the maximum amount

$$\lambda = \min_{k:i \in \mathcal{B}_k} \{y_{ki}\} \tag{6.12}$$

of composite commodity.

6.1.2 Decomposition part

A decomposition and trade is the reverse action of a trade and production. There exist separate auctions for decomposition trades and production trades.

An agent that attempts to decompose a composite commodity places bids to *dispose* of basic commodities to the respective auctioneers. Let $\mathcal{B}_k$ denote the index set of the agents that place bids for the basic commodity of class k. A bid is a message to the auctioneer where the message contains information about the utility function, the commodity requirement and the current commodity bundle of the bidder. The bid also contains information specifying the composite commodity number h that the bidder i intends to decompose. The auctioneers acquire all the basic commodities disposed by the bidders. An auctioneer $k \in \mathcal{N}$ determines the amount of basic commodity *deallocated* from the bidders $i \in \mathcal{B}_k$. The deallocation of y_i units of basic commodity from a bidder i means that the bidder is entitled to dispose of y_i or more units of basic commodity.

Mechanism: The auctioneer

The task of the auctioneer is to deallocate basic commodities from the bidders. The auctioneer uses the information in the bids to determine the amount of basic commodity to be deallocated from the bidders. Adhering to the fair trade criteria, the maximum amount of basic commodities is deallocated from each bidder.

The auction is classified as a *multi-unit auction*, because more than one unit of a commodity is auctioned at the auction. It is also a *private-value auction*, since each bidder has its own valuation of the commodity and each bidder is ignorant of the value that the other bidders place on the commodity. Each bidder may place only one bid at an auction.

Let the variable y_i denote the total amount of basic commodity deallocated from the bidder i. When the auction starts, no basic commodities are deallocated from the bidders and $y_i = 0$ for all $i \in \mathcal{B}_k$. During the auction the auctioneer may deallocate basic commodities and thus increase the amounts $y_i, i \in \mathcal{B}_k$.

The auction consists of a sequence of *phases* in which basic commodities are deallocated from the bidders. At each phase the bidder i with the highest utility of all the bidders is identified. One unit of basic commodity is deallocated from the bidder i and the phase ends. At each phase, the total amount of basic commodity y_i deallocated from the bidder i therefore increases by one unit.

Because one unit of basic commodity is deallocated from a bidder i at the end of each phase, a new bidder utility must be calculated. According to the information contained in the bid, the bidder i intends to decompose the composite commodity number h of class i. Hence the bidder i will decompose an amount y_i of composite commodity number h and dispose of an amount y_i of basic commodity of class k. The bidder utility is therefore $u_i(\mathbf{b}_i - y_i\mathbf{e}_h, d_i)$ for $y_i > 0$. At the end of each phase the utility of the bidder i is set to $u_i(\mathbf{b}_i - y_i\mathbf{e}_h, d_i)$.

The auction terminates when further deallocation of commodity will result in the utility of a bidder $i : y_i > 0$ to be less than the auctioneer's utility

$$u_k((b_k + \sum_{j \in \mathcal{B}_k} y_j)\mathbf{e}_1, d_k)$$

of the basic commodity of class k. The auction therefore terminates at the phase where another commodity deallocation would falsify one of the relations

$$u_i(\mathbf{b}_i - y_i\mathbf{e}_h, d_i) \geq u_k((b_k + \sum_{j \in \mathcal{B}_k} y_j)\mathbf{e}_1, d_k) \tag{6.13}$$

for any $i \in \mathcal{B}_k$ such that $y_i > 0$. Note that the auction termination condition in Relations (6.13) implies that the Relations (6.9) are satisfied.

When the auction terminates the auctioneer sends a message to each bidder $i \in \mathcal{B}_k$. Each message contains information about the total amount y_i of basic commodity deallocated from the bidder i by the auctioneer k.

Strategy: The bidder

The task of the bidder is to choose which composite commodity to decompose. The agent strategy is defined such that an agent disposes of a basic commodity with a high price. Consider a price vector $\mathbf{a} = (5, 5, 5, 5)$, denoting the basic commodity prices in the production bundle of some composite commodity. An agent prefers to decompose a composite commodity with a price vector $\mathbf{b} = (5, 5, 1, 9)$ rather than a price vector $\mathbf{a}$ as explained below.

An agent i places bids for the decomposition of the composite commodity with the highest price index $\theta_{ij} = \max\limits_{k \in \mathcal{K}_{i,j}} s_k$.

Let y_{ki} be the total amount of basic commodity deallocated from the bidding agent i by each auctioneer k. The agent i disposes an amount y_{ki} or more of basic commodity of class $k \in \mathcal{K}_{i,j}$. When decomposing an amount λ of composite commodity number j of class i, an agent i disposes of amounts λ of basic commodities of each class $k \in \mathcal{K}_{i,j}$ (see Assumption 3.2). The agent i maximises utility by decomposing the minimum amount

$$\lambda = \max_{k: i \in \mathcal{B}_k} \{y_{ki}\} \tag{6.14}$$

of composite commodity.

6.2 The computational complexity of the social community

The computational complexity of an auction is $O(C)$, where C is the amount of commodities auctioned. The computational complexity of an agent strategy is $O(n)$ where n is the number of routes defined for that OD-pair. An auction process does not communicate with other auction processes. The auctions are concurrent processes and since the auction process serialises the resource allocation, they may be executed independently.

6.3 Testing the social community

We present test results of the social community using the mechanism and agent strategy defined in Section 6.1

We assume that calls (connections) arrive individually at the instants of a Poisson stream and that the call holding times are exponentially distributed. Let ρ_i denote the load offered to OD-pair i.

A connection uses one unit of capacity on each link over which it is routed. The measure of utility[1] is the GoS of an OD-pair, hence

$$u_i(\mathbf{b}_i, \rho_i) = 1 - B\left(\sum\nolimits_{j=1}^{n} b_{ij},\ \rho_i\right) \tag{6.15}$$

where b_{ij} is the capacity assigned to the LSP j connecting OD-pair i, n is the number of LSPs connecting OD-pair i and $B(\cdot)$ is the Erlang B loss probability defined by

$$B(c, \rho) = \frac{\rho^c / c!}{\sum\limits_{k=0}^{c} \rho^k / k!}.$$

The supremo utility is 1 and ρ_i characterises the requirement for commodities of class i (at agent i). The utility function $u_i(\mathbf{b}_i, \rho_i)$ is such that high blocking results in a low utility and *vice versa*. An agent with a low utility for its commodity bundle has incentive to acquire more commodities.

Figure 6.1 depicts the utility of an OD-pair i with $\rho_i = 30$.

6.3.1 Agent behaviour and equilibrium

The experiments in this subsection are all of an 8-node bi-directional network with pre-configured LSPs calculated by XFG [5]. The OD-pair call arrival intensity matrix is given in Table 6.1.

The bi-directional link capacities for the network model are given in Table 6.2 and depicted in Figure 6.2. As a first approach to test the social community we only test networks with bi-directional links.

[1] We consider only a single call class, and we note that the Erlang loss function $B(c, \rho)$ is convex.

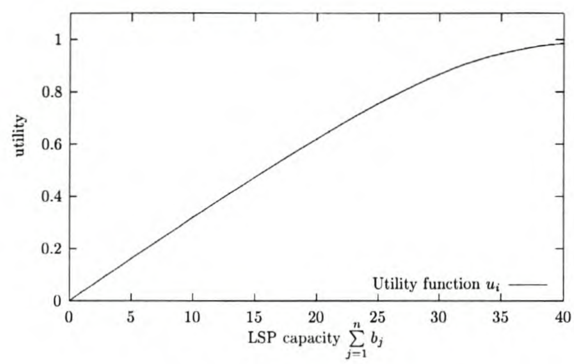

Figure 6.1: The utility of an OD-pair.

Node	1	2	3	4	5	6	7	8
1	-	13	15	2	20	10	4	6
2		-	49	6	64	29	11	17
3			-	7	76	34	13	21
4				-	9	4	2	2
5					-	45	17	27
6						-	8	12
7							-	4
8								-

Table 6.1: The arrival intensities to the 8-node network.

The market process (see Definition 3.16) consists of a series of steps. Agents may be invoked simultaneously or in any order at each step of the market process. The actions of an agent may change the state of the market resulting in a changed market state at the subsequent step. Step 0 refers to the starting conditions of an experiment.

There are 28 agents in the 8-node network. The agents were invoked in a random order once every step. The prices of the basic commodities were re-calculated at each step. Auctions took place every second agent invocation, thus every second step.

For each experiment, each direct route LSP was initially configured to have a capacity equal to that of the corresponding physical link. The multi-link LSPs were initially configured to have no capacity. The multi-link LSPs therefore represent the composite commodities that may be produced and the direct route LSPs represent the basic commodities. In market terms we view the LSP configuration as a market with all-basic commodities and a given production technology. All the trades in the social community are fair, thus a fair market equilibrium is reached whenever no trades can take place.

Node	1	2	3	4	5	6	7	8
1	-	305						128
2		-	178					
3			-	239				93
4				-	157	95		
5					-		125	
6						-	100	
7							-	151
8								-

Table 6.2: The physical link capacities of 8-node network.

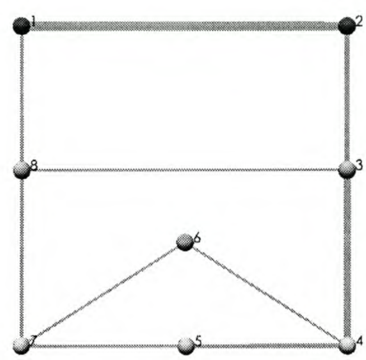

Figure 6.2: The 8-node network.

Experiment A

Figure 6.3 presents the bidding, auctioning, production and decomposition events of Experiment A. The commodity auctioned for production is the total amount of basic commodity allocated at an auction. The commodity auctioned for decomposition is the total amount of basic commodity deallocated at an auction. The commodity produced is the total amount of composite commodity produced by some agent. The commodity decomposed is the total amount of composite commodity decomposed by some agent.

Figure 6.4 presents graphs of the trace data as Experiment A progressed. The total cost of the composite commodities represents the efficiency of the composite commodity allocation (see Section 5.4).

The terms $\sqrt{V(\Theta)}$ and $1 - A(\Theta)$ of the social welfare penalty function $W(\cdot)$ (see Equation (5.7)) are represented by the variance of the OD-pair blocking and network loss probability respectively. These terms are called the measures of social welfare. Let $I_i = \rho_i / \sum_{j \in \mathcal{N}} \rho_j$ be the load offered to OD-pair $i \in \mathcal{N}$ expressed as a fraction of the total load offered to the network. $u_i(\mathbf{b}_i, \rho_i)$ is the probability that a connection offered to an OD-pair i is carried (see Equation (6.15)), hence $I_i u_i(\mathbf{b}_i, \rho_i)$ is the load carried by OD-pair i expressed as a fraction of the total load offered to the network. It follows that the network loss probability (the total load lost expressed as a fraction of the total load offered to the network) is $W(\Theta) = 1 - \sum_{i \in \mathcal{N}} I_i u_i(\mathbf{b}_i, \rho_i)$, where Θ is the state of the network.

In Experiment A at step 0 all the commodities in the market are basic commodities. Figure 6.3(a) shows that composite commodities are produced between step 0 and step 16 and no composite commodities are decomposed. Figure 6.4(a) shows that the total cost of the composite commodities increases as more composite commodities are produced. The production of composite commodities decreases the supply of basic commodities. The decreased supply of basic commodities leads to an increased basic commodity price and an increased cost of producing the composite commodities.

Figure 6.3(b) shows that after step 16 although some agents continue to place bids and commodities are auctioned, no trade takes place. Thus a market equilibrium is reached at step 16. Figure 6.4(b) shows that the network loss probability (GoS) settles at just above 0.05.

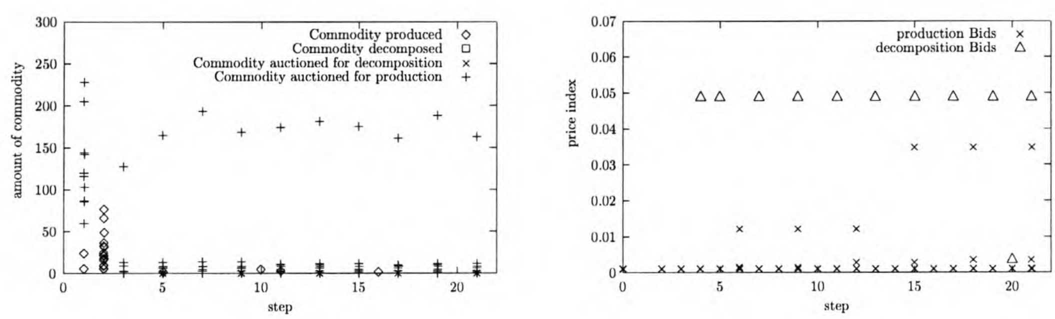

(a) The productions and decompositions.

(b) The bids.

Figure 6.3: The productions, decompositions and bids of Experiment A.

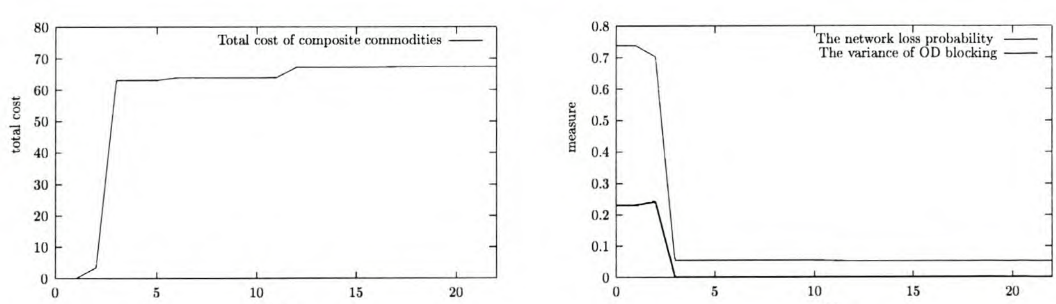

(a) The total cost of the composite commodities.

(b) The measures of social welfare.

Figure 6.4: Trace data of Experiment A.

Experiment B

Experiment B is similar to Experiment A except that the requirement for the commodity of class 4–5 is increased by 40 Erlangs. We chose to increase the requirement for the commodity of class 4–5 because there are no multi-link routes defined between nodes 4 and 5. The technology of the market is therefore such that no composite commodities of class 4–5 can be produced. Composite commodities produced from the basic commodity of class 4–5 must therefore be decomposed to increase the utility of agent 4–5.

The bandwidth of the physical link between nodes 4 and 5 represents 10% of the total amount of bandwidth in the network and the increased commodity requirement of 40 Erlangs represents a 7.6% increase in the total commodity requirement.

In Experiment B at step 20 the requirement for commodities of class 4–5 is increased. Figure 6.5(c) shows that the increased requirement for commodities of class 4–5 increases the network loss probability at step 20.

The increased requirement for commodities of class 4–5 increases the price of the basic commodity of class 4–5. Figure 6.5(b) shows that at step 20 the total cost of the composite commodities increases as result of the increase in the price of the basic commodity of class 4–5. Some composite commodities containing the basic commodity of class 4–5 in their production bundles are decomposed because of the increased price of the basic commodity of class 4–5. Figure 6.5(a) shows that these composite commodities are decomposed between step 24 and step 28.

Figure 6.5(a) shows that some productions take place after step 28 as result of the preceding decompositions. A market equilibrium is reached at step 40.

Figure 6.5(b) shows that although trading improved the measures of social welfare, the equilibrium measures of social welfare are somewhat worse than before the requirement of commodities of class 4–5 was increased.

Experiment C

Experiment C is similar to Experiment A except that at steps $400k$, $k \in \{1, \ldots, 1500\}$ the requirement for commodities of a randomly chosen class is increased by 40 Erlangs. At steps $400k + 200$, $k \in \{1, \ldots, 1500\}$ the requirement for commodities of the randomly chosen class is restored to its original value. The market exogenous variables thus change at steps $200k + 200$, $k \in \{1, \ldots, 1500\}$.

Figure 6.6 presents trace data from Experiment C. A market equilibrium was reached after each change in the market exogenous variables. Note that each subsequent market equilibrium is not necessarily the same as any of the previous market equilibria.

Experiment C spans $600,000$ steps (only 5000 steps are shown in Figure 6.6). Figure 6.6(b) shows

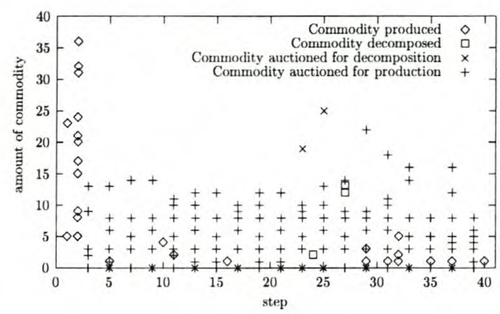

(a) The productions and decompositions.

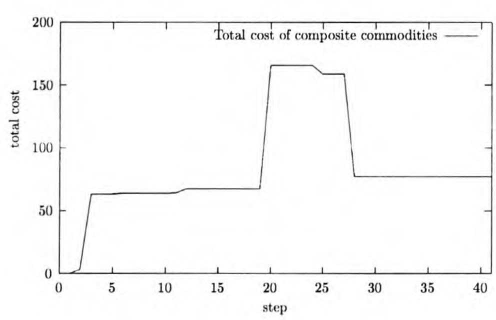

(b) The total cost of the composite commodities.

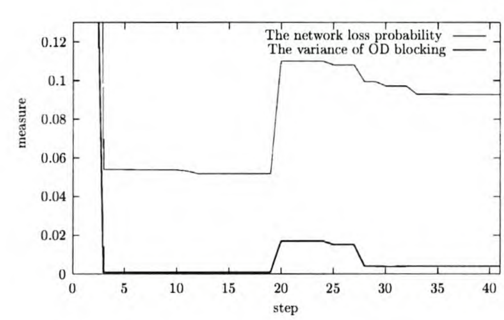

(c) The measures of social welfare.

Figure 6.5: The productions and decompositions and the trace data of Experiment B.

that trading improves the measures of social welfare after each increase in the requirement for commodities. The average number of steps to reach a market equilibrium after an increase in requirement for commodities is 28.07 and the standard deviation is 19.21. The average number of steps to reach a market equilibrium after traffic restoration of the requirement for commodities is 25.32 and the standard deviation is 18.19. The market was in disequilibrium for less than 107 steps after each of the 2998 instances where the market exogenous variables changed.

It is thus reasonable to assume that a fair market equilibrium will be reached in a finite number of fair trades.

6.3.2 XFG and the social community

We evaluate the social community's effectiveness in maximising social welfare by comparing networks configured by XFG [5] with networks configured by the social community. XFG is a centralised optimisation algorithm that calculates routes and route bandwidths for a given network topology and traffic demands such that the bandwidth configuration maximises the network's rate

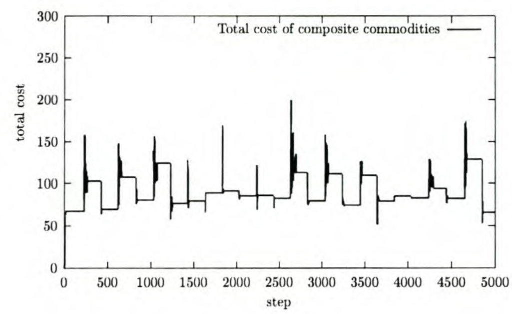

(a) The total cost of the composite commodities.

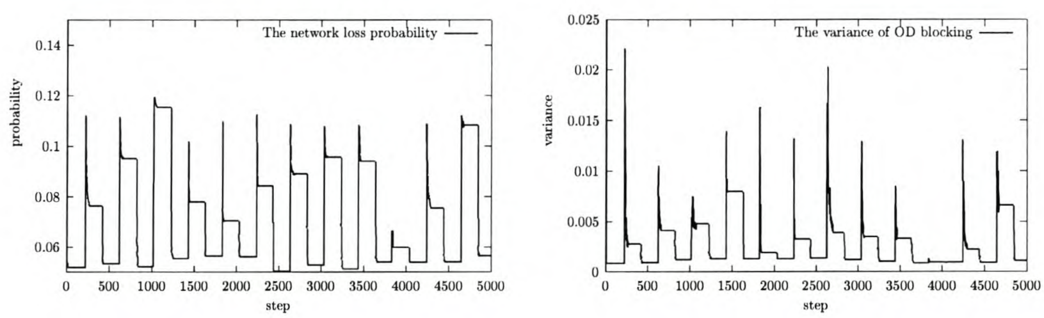

(b) The measures of social welfare.

Figure 6.6: The trace data of Experiments C.

of earning revenue $\sum_{i\in\mathcal{N}} \rho_i u_i(\mathbf{b}_i, \rho_i)$.

Experiments D_1 and D_2

Figure 6.7 shows the trace data of the Experiments D_1 and D_2 of a small network whose details are listed in Table 6.3. Experiment D_1 starts with routes and route bandwidths as configured by XFG. Experiment D_2 starts with routes configured by XFG and all the link capacities assigned to the direct routes.

The result of Experiment D_1 is that

- the social community makes insignificant changes the measures of social welfare as shown in Figure 6.7(d).

The results of Experiment D_2 are that

- Figure 6.7(a) shows that initially a large number of productions take place that improve the measures of social welfare substantially (see Figure 6.7(d)) followed by a series of productions and decompositions that improve the measures of social welfare only slightly.

- Figure 6.7(b) shows that the cost of the equilibrium allocation of composite commodities in Experiment D_1 and D_2 are similar.
- Figure 6.7(c) shows a close similarity between the equilibrium quantities of commodities assigned to each commodity class by XFG and by the social community.
- Figure 6.7(d) shows that the social community improves the measures of social welfare such that the equilibrium measures of social welfare in Experiments D_2 and D_1 are similar.

There is a limited number of possible bandwidth allocations in a small network. It is therefore not surprising that the bandwidth allocations calculated by XFG and by the social community are similar for a small network.

Experiments E_1, E_2, F_1 and F_2

Figure 6.10 shows the trace data of the Experiments E_1 and E_2 of the medium network in Figure 6.8 whose details are listed in Table 6.3. Experiment E_1 starts with routes and route bandwidths as configured by XFG, whereas Experiment E_2 starts with routes configured by XFG and all the link capacities assigned to the direct routes.

The result of Experiment E_1 is that

- the social community makes insignificant changes the measures of social welfare as shown in Figure 6.10(b).

The results of Experiment E_2 are that

- Figure 6.10(a) shows that initially a large number of productions take place that improve the measures of social welfare substantially (see Figure 6.10(d)) followed by a series of productions and decompositions that further improve the measures of social welfare.
- Figure 6.10(b) shows that the cost of the equilibrium allocation of composite commodities in Experiment E_1 and E_2 are similar.
- Figure 6.10(c) shows a close similarity between the equilibrium quantities of commodities assigned to each commodity class by XFG and by the social community.
- Figure 6.10(d) shows that the social community improves the measures of social welfare such that the equilibrium measures of social welfare in Experiments E_2 and E_1 are similar.

Figure 6.11 shows the trace data of the Experiments F_1 and F_2 of the large network in Figure 6.9 whose details are listed in Table 6.3. Experiment F_1 starts with routes and route bandwidths as configured by XFG, whereas Experiment F_2 starts with routes configured by XFG and all the link capacities assigned to the direct routes. The results of Experiments F_1 and F_2 are qualitatively similar to the results of Experiment E_1 and E_2.

Table 6.4 summarises the results of Experiments D_2, E_2, and F_2 and the bandwidth configured on the test networks by XFG. The "social community" is abbreviated as "SOC".

	small network	medium network	large network
network nodes	8	20	50
physical links	10	51	101
market agents	28	190	1225
distinct commodities	61	485	2729

Table 6.3: Network details.

equilibrium value	algorithm	small network	medium network	large network
network loss probability	XFG	0.0227	0.00001	0.00018
	SOC	0.0334	0.00006	0.00018
variance of OD blocking	XFG	0.0005	6.1056×10^{-10}	1.3187×10^{-7}
	SOC	0.0007	1.6921×10^{-8}	1.0361×10^{-7}
total cost of composite commodities	XFG	46.1025	378.921	638.784
	SOC	49.4818	402.210	674.332
number of distinct commodities used	XFG	61	485	2729
	SOC	51	330	1903
total number of trades	SOC	45	493	3715
average number of trades per agent	SOC	1.6	2.6	3.0

Table 6.4: Summary of results.

In summary, XFG achieved a slightly lower network loss probability than the social community, although the variance of the OD-pair blocking probabilities is similar. XFG calculated a more efficient bandwidth configuration according to the total cost of the composite commodities. The number of composite commodities (multi-link routes) used by the social community was less than the number of routes used by XFG. The average number of trades per agent increases as the number of nodes in the network increases.

6.4 Results

The social community improves social (network) welfare in the test networks.

For all the test networks the social community calculated an improved bandwidth configuration comparable with the optimal bandwidth configuration calculated by a centralised optimisation algorithm called XFG. The number of routes used by the bandwidth configuration calculated by the social community is less than the number of routes used by the bandwidth configuration calculated by XFG. Table 6.4 presents a summary of these results.

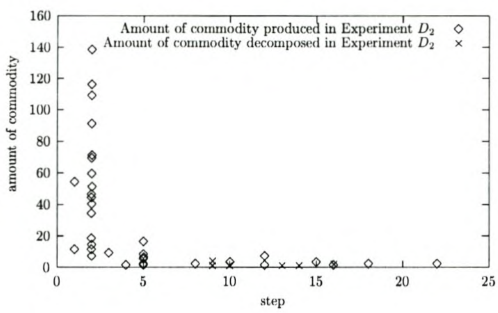

(a) The productions and decompositions.

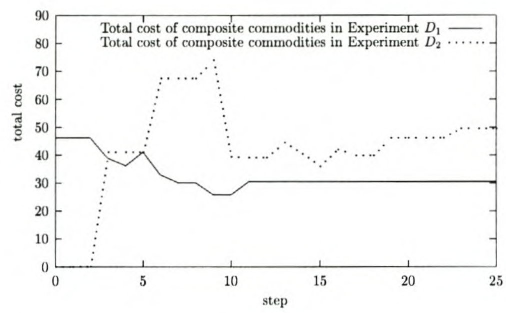

(b) The total cost of the composite commodities.

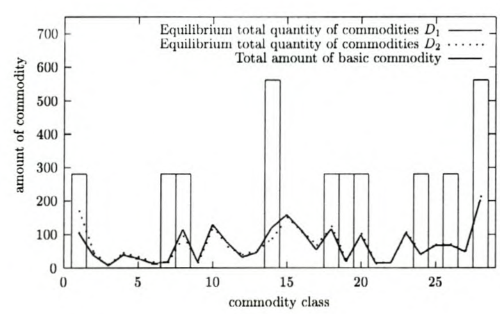

(c) The equilibrium quantity of commodities.

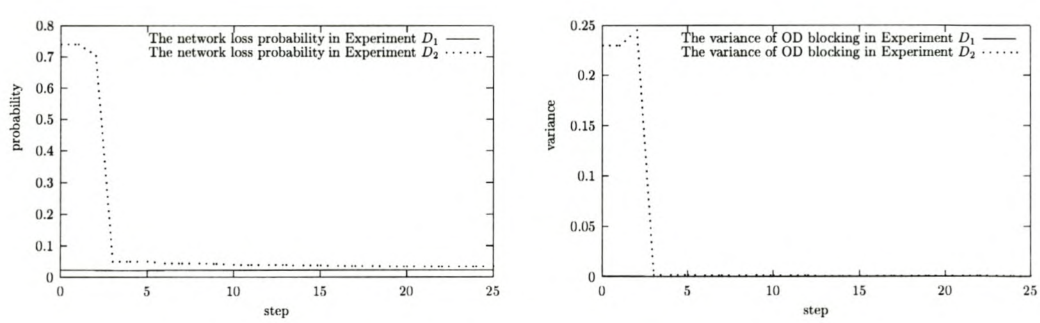

(d) The measures of social welfare.

Figure 6.7: The trace data of the experiments with a small network.

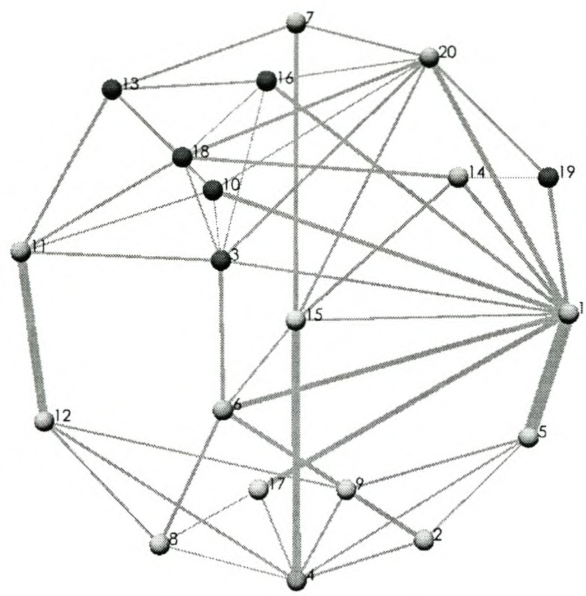

Figure 6.8: The medium network: 20 nodes.

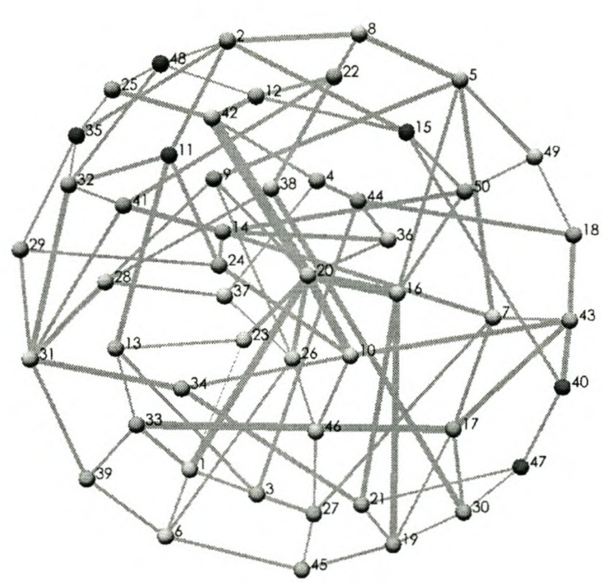

Figure 6.9: The large network: 50 nodes.

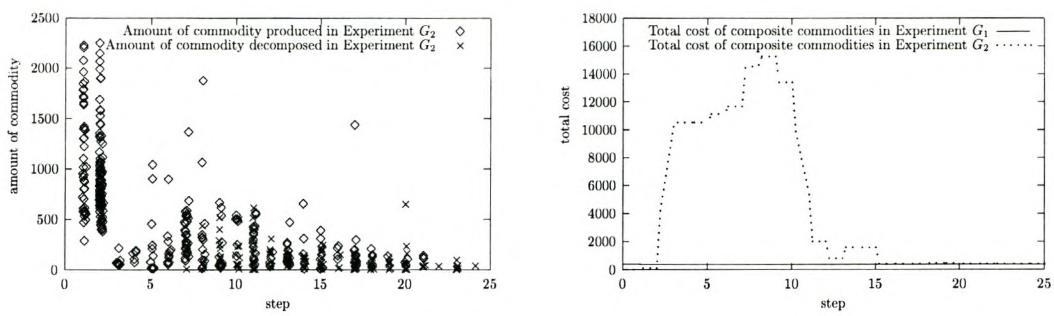

(a) The productions and decompositions. (b) The total cost of the composite commodities.

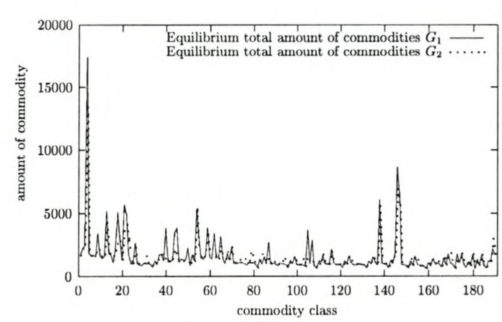

(c) The equilibrium quantity of commodities.

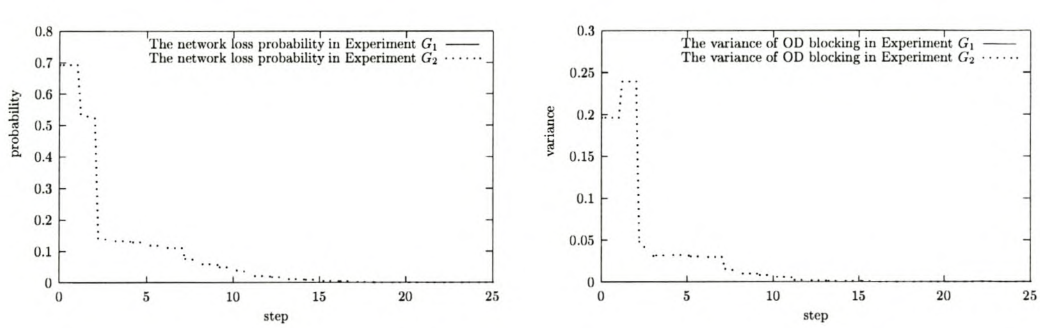

(d) The measures of social welfare.

Figure 6.10: The trace data of the experiments with a medium network.

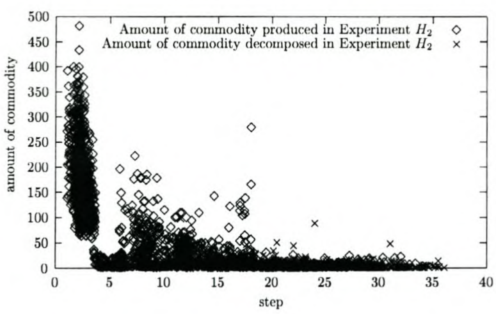

(a) The productions and decompositions.

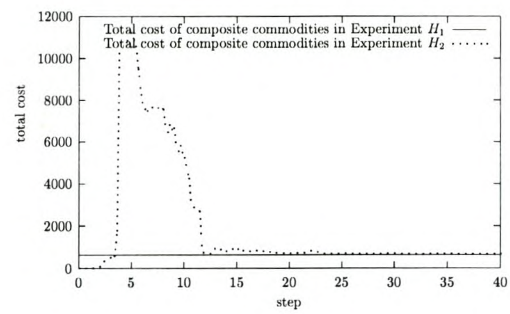

(b) The total cost of the composite commodities.

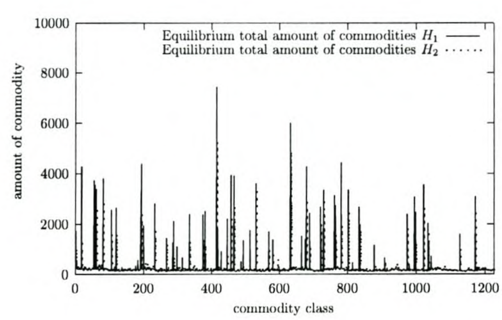

(c) The equilibrium quantity of commodities.

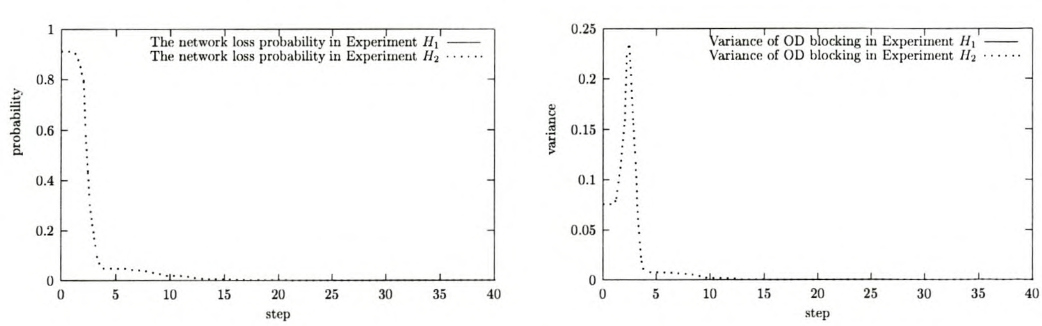

(d) The measures of social welfare.

Figure 6.11: The trace data of the experiments with a large network.

Chapter 7

Conclusion

This thesis focuses on traffic engineering and methods of routing traffic. Our aim is to map the problem of routing traffic onto the microeconomic problem of resource allocation in a market. We call this market a *bandwidth market*. As a prelude to this mapping we present concepts and entities of game theory, markets and networks.

We formulate the fundamental definitions of our bandwidth markets adhering to assumptions which simplify the market mechanisms and the communication between agents in the market. According to these fundamental definitions our bandwidth markets are classified as oligopoly basic market models.

We construct a bandwidth market called the *dual-oligopoly*. The dual-oligopoly conforms more closely to some market models in economic literature than many bandwidth markets encountered in the literature. A market mechanism and agent behaviour in the dual-oligopoly may lead to an optimum allocation of resources similar to that achieved in real economic markets.

We construct a second bandwidth market called the *social community*. We define a price for bandwidth in accordance with the economic theory of supply. We define rules and mechanisms in the social community that enforce cooperation between the market agents.

Finally we present an implementation of the social community and test this implementation. The tests conclude that the social community is a distributed traffic engineering method that reconfigures traffic routing to improve the network loss rate.

Appendix A

A proof of fairness and cooperation

Lemmas A.1 and A.3 prove that both parts of the auctioning mechanism presented in Section 6.1 allow multiple fair trades to take place simultaneously while excluding trades that are not fair. The mechanism thus conforms to the fair trade criteria in Definition 5.6.

Lemmas A.2 and A.4 prove that if no fair trades of the composite commodity chosen by the agent strategy are possible, then the agent is unable to do any fair trades of composite commodities. Thus by using the agent strategy presented in Section 6.1, the agents never abstain from fair trading and thus always cooperate with the other agents in the social community.

Let $\mathcal{R}_i$ denote the index set of the composite commodities of class i, and let

$$\mathcal{K}_{i,j} = \{k : [P_i^{-1}(\mathbf{e}_j)]_k = 1\}$$

denote the set containing the classes of the basic commodities required to produce a composite commodity number j of class i as defined in Definition 3.13.

Lemma A.1 *Consider an agent δ that places bids at the auctioneers $k \in \mathcal{K}_{\delta,c}$ for the production $P_\delta(\mathbf{x}_\delta) = \lambda \mathbf{e}_c$ producing an amount λ of composite commodity number c of class δ from the production bundle $\mathbf{x}_\delta$. Let $\mathcal{B}_k$ be the index set of the bidders that place bids for basic commodity at the auctioneer k. Note that the agent δ is in the bidder sets $\mathcal{B}_k$, where $k \in \mathcal{K}_{\delta,c}$.*

If each auctioneer $k \in \mathcal{K}_{\delta,c}$ allocates an amount y_{ki} of basic commodity of class k to each bidder $i \in \mathcal{B}_k$ such that the auction termination condition (6.10) is true and the agent δ produces an amount (see Equation (6.12)) $\lambda = \min_{k:\delta\in\mathcal{B}_k} \{y_{k\delta}\}$ of composite commodity number c of class δ then the production $P_\delta(\mathbf{x}_\delta) = \lambda \mathbf{e}_c$ is a fair trade.

Proof. From the auction termination condition (6.10), where an auctioneer k allocates an amount y_{ki} of basic commodity of class k to each agent $i \in \mathcal{B}_k$, we deduce that

$$u_\delta(\mathbf{b}_\delta + y_{k\delta}\mathbf{e}_c, d_\delta) \leq u_k((b_k - \sum\nolimits_{i \in \mathcal{B}_k} y_{ki})\mathbf{e}_1, d_k) \tag{A.1}$$

for each auctioneer $k : \delta \in \mathcal{B}_k$ and $y_{k\delta} > 0$. From Equation (6.12) the bidder δ produces an amount

$$\lambda = \min_{i:\delta \in \mathcal{B}_i} \{y_{i\delta}\}$$

of composite commodity number c of class k and thus uses an amount

$$x_{\delta k} = \lambda = \min_{i:\delta \in \mathcal{B}_i} \{y_{i\delta}\} \leq y_{k\delta} \tag{A.2}$$

of each basic commodity of class $k : \delta \in \mathcal{B}_k$.

Because utility functions are monotone increasing with respect to commodity amounts and by using the Relations (A.1) and $x_{\delta k} = \lambda \leq y_{k\delta}$ for all $k : \delta \in \mathcal{B}_k$ from Relation (A.2) we write

$$\begin{aligned} u_\delta(\mathbf{b}_\delta + \lambda\mathbf{e}_c, d_\delta) &\leq u_\delta(\mathbf{b}_\delta + y_{k\delta}\mathbf{e}_c, d_\delta) && \text{(A.3)} \\ &\leq u_k((b_k - \sum\nolimits_{i \in \mathcal{B}_k} y_{ki})\mathbf{e}_1, d_k) && \text{(A.4)} \\ &\leq u_k((b_k - \sum\nolimits_{i \in \mathcal{B}_k, i \neq \delta} y_{ki} - x_{\delta k})\mathbf{e}_1, d_k) && \text{(A.5)} \\ &\leq u_k\left((b_k - x_{\delta k})\,\mathbf{e}_1, d_k\right) && \text{(A.6)} \end{aligned}$$

for all $k \in \mathcal{K}_{\delta,c}$ for the production $P_\delta(\mathbf{x}_\delta) = \lambda\mathbf{e}_c$, where Relation (A.2) was used in steps (A.3) and (A.5) and Relation (A.1) was used in step (A.4).

By comparing Relation (A.6) with Definition 5.6 it follows that the production $P_\delta(\mathbf{x}_\delta) = \lambda\mathbf{e}_c$, producing an amount λ of the composite commodity number c of class δ from the production bundle $\mathbf{x}_\delta$ is a fair trade. ∎

Lemma A.2 *For an agent δ, if no fair trade (production) of the composite commodity with the lowest price index is possible then no fair trade (production) of any composite commodity of class δ is possible.*

Proof. Let $\mathbf{x}_\delta$ be the unit production bundle of the composite commodity number c of class δ. Consider the case where an agent $\delta \in \mathcal{N}$ cannot do a fair trade to produce any amount $\lambda > 0$ of the composite commodity c of class δ with $\theta_{\delta c} \leq \theta_{\delta r}$ for all $r \in \mathcal{R}_\delta$. Thus the production $P_\delta(\lambda\mathbf{x}_\delta) = \lambda\mathbf{e}_c$ of any amount $\lambda > 0$ of the composite commodity number c of class δ is not a fair trade and from Definition 5.6

$$u_\delta(\mathbf{b}_\delta + \lambda\mathbf{e}_c, d_\delta) > u_k\left(b_k\mathbf{e}_1 - \lambda\mathbf{e}_1, d_k\right)$$

for at least one $k \in \mathcal{K}_{\delta,c}$ and for $\lambda > 0$. Then

$$u_\delta(\mathbf{b}_\delta, d_\delta) \geq u_k\left(b_k\mathbf{e}_1, d_k\right) \tag{A.7}$$

for at least one $k \in \mathcal{K}_{\delta,c}$, because the utility functions are monotone increasing.

For the composite commodity number c of class δ with $\theta_{\delta c} \leq \theta_{\delta r}$ for all $r \in \mathcal{R}_\delta$ we write (see Definition 6.11)

$$\max_{k \in \mathcal{K}_{\delta,c}} \{s_k\} = \theta_{\delta c} \leq \theta_{\delta r} = \max_{k \in \mathcal{K}_{\delta,r}} \{s_k\}$$

which implies that

$$\max_{k \in \mathcal{K}_{\delta,c}} \{1 - u_k(b_k \mathbf{e}_1, d_k)\} \leq \max_{k \in \mathcal{K}_{\delta,r}} \{1 - u_k(b_k \mathbf{e}_1, d_k)\}$$

and

$$\min_{k \in \mathcal{K}_{\delta,c}} \{u_k(b_k \mathbf{e}_1, d_k)\} \geq \min_{k \in \mathcal{K}_{\delta,r}} \{u_k(b_k \mathbf{e}_1, d_k)\} \tag{A.8}$$

for all $r \in \mathcal{R}_\delta$.

From Relation (A.7) and (A.8)

$$\begin{aligned} u_\delta(\mathbf{b}_\delta, d_\delta) &\geq u_k(b_k \mathbf{e}_1, d_k) & \text{(A.9)} \\ &\geq \min_{k \in \mathcal{K}_{\delta,c}} \{u_k(b_k \mathbf{e}_1, d_k)\} \\ &\geq \min_{k \in \mathcal{K}_{\delta,r}} \{u_k(b_k \mathbf{e}_1, d_k)\} & \text{(A.10)} \end{aligned}$$

for all composite commodities $r \in R_\delta$ of class δ, where Relation (A.7) was used in step (A.9) and Relation (A.8) was used in step (A.10). Therefore for every composite commodity $r \in \mathcal{R}_\delta$ of class δ and for at least one $k \in \mathcal{K}_{\delta,r}$

$$u_k((b_k - \lambda_r)\mathbf{e}_1, d_k) \quad < \quad u_\delta(\mathbf{b}_\delta + \lambda_r \mathbf{e}_r, d_\delta)$$

with $\lambda_r > 0$ and from Definition 5.6 no fair trade (production) of any composite commodity of class δ is possible. ■

Lemma A.3 *Consider an agent δ that places bids at the auctioneers $k \in \mathcal{K}_{\delta c}$ to dispose basic commodities after the decomposition $P_\delta^{-1}(\lambda \mathbf{e}_c) = \mathbf{x}_\delta$ of an amount λ of composite commodity number c of class δ to form a production bundle $\mathbf{x}_\delta$. Let $\mathcal{B}_k$ be the index set of the bidders that place bids for basic commodity at the auctioneer k. Note that the agent δ is in the bidder sets $\mathcal{B}_k$, where $k \in \mathcal{K}_{\delta,c}$.*

If each auctioneer $k \in \mathcal{K}_{\delta,c}$ deallocates an amount y_{ki} of basic commodity of class k from each bidder $i \in \mathcal{B}_k$ such that the auction termination condition (6.13) is true and the agent δ decomposes an amount (see Equation (6.14)) $\lambda = \max_{k:\delta \in \mathcal{B}_k} \{y_{k\delta}\}$ of composite commodity number c of class δ then the decomposition $P_\delta^{-1}(\lambda \mathbf{e}_c) = \mathbf{x}_\delta$ is a fair trade.

Proof. From the auction termination condition (6.13), where an auctioneer k deallocates an amount y_{ki} of basic commodity of class k from each agent $i \in \mathcal{B}_k$, we deduce that

$$u_\delta(\mathbf{b}_\delta - y_{k\delta} \mathbf{e}_c, d_\delta) \geq u_k((b_k + \sum_{i \in \mathcal{B}_k} y_{ki}) \mathbf{e}_1, d_k) \tag{A.11}$$

for each auctioneer $k : \delta \in \mathcal{B}_k$ and $y_{k\delta} > 0$. From Equation (6.14) the bidder δ decomposes the amount

$$\lambda = \max_{i:\delta \in \mathcal{B}_i} \{y_{i\delta}\}$$

of composite commodity number c of class δ and thus disposes of an amount

$$x_{\delta k} = \lambda = \max_{i:\delta \in \mathcal{B}_i} \{y_{i\delta}\} = y_{k\delta} \tag{A.12}$$

of at least one basic commodity of class $k : \delta \in \mathcal{B}_k$.

Because utility functions are monotone increasing with respect to commodity amounts and by using the Relations (A.11) and $x_{\delta k} = \lambda = y_{k\delta}$ for at least one $k : \delta \in \mathcal{B}_k$ from Equation (A.12) we write

$$\begin{aligned} u_\delta(\mathbf{b}_\delta - \lambda \mathbf{e}_c, d_\delta) &= u_\delta(\mathbf{b}_\delta - y_{k\delta}\mathbf{e}_c, d_\delta) \\ &\geq u_k((b_k + \textstyle\sum_{i \in \mathcal{B}_k} y_{ki})\mathbf{e}_1, d_k) \\ &= u_k((b_k + \textstyle\sum_{i \in \mathcal{B}_k, i \neq \delta} y_{ki} + x_{\delta k})\mathbf{e}_1, d_k) \\ &\geq u_k\left((b_k + x_{\delta k})\,\mathbf{e}_1, d_k\right) \end{aligned} \tag{A.13}$$

for at least one basic commodity of class $k \in \mathcal{K}_{\delta,c}$ for the decomposition $P_\delta^{-1}(\lambda \mathbf{e}_c) = \mathbf{x}_\delta$.

By comparing Relation (A.13) with Definition 5.6 it follows that the decomposition $P_\delta^{-1}(\lambda \mathbf{e}_c) = \mathbf{x}_\delta$, decomposing an amount λ of the composite commodity number c of class δ to form the production bundle $\mathbf{x}_\delta$ is a fair trade. ■

Lemma A.4 *For an agent δ, if no fair decomposition of the composite commodity with the highest price index is possible then no fair decomposition of any composite commodity of class δ is possible.*
Proof. Let $\mathbf{x}_\delta$ be a unit production bundle of the composite commodity number c of class δ. Consider the case where an agent $\delta \in \mathcal{N}$ cannot do a fair trade to decompose any amount $\lambda > 0$ of composite commodity number c of class δ with $\theta_{\delta c} \geq \theta_{\delta r}$ for all $r \in \mathcal{R}_\delta$. Thus the decomposition $P_\delta^{-1}(\lambda \mathbf{e}_c) = \lambda \mathbf{x}_\delta$ of any amount $\lambda > 0$ of the composite commodity c is not a fair trade and from Definition 5.6

$$u_\delta(\mathbf{b}_\delta - \lambda \mathbf{e}_c, d_\delta) < u_k\left(b_k \mathbf{e}_1 + \lambda \mathbf{e}_1, d_k\right)$$

for all $k \in \mathcal{K}_{\delta,c}$ and for $\lambda > 0$. Then

$$u_\delta(\mathbf{b}_\delta, d_\delta) \leq u_k\left(b_k \mathbf{e}_1, d_k\right) \tag{A.14}$$

for all $k \in \mathcal{K}_{\delta,c}$, because the utility functions are monotone increasing.

For the composite commodity number c of class δ with $\theta_{\delta c} \leq \theta_{\delta r}$ for all $r \in \mathcal{R}_\delta$ we write (see Definition 6.11)

$$\max_{k \in \mathcal{K}_{\delta,c}} \{s_k\} = \theta_{\delta c} \geq \theta_{\delta r} = \max_{k \in \mathcal{K}_{\delta,r}} \{s_k\}$$

which implies that

$$\max_{k \in \mathcal{K}_{\delta,c}} \{1 - u_k(b_k \mathbf{e}_1, d_k)\} \geq \max_{k \in \mathcal{K}_{\delta,r}} \{1 - u_k(b_k \mathbf{e}_1, d_k)\}$$

and

$$\min_{k \in \mathcal{K}_{\delta,c}} \{u_k(b_k \mathbf{e}_1, d_k)\} \leq \min_{k \in \mathcal{K}_{\delta,r}} \{u_k(b_k \mathbf{e}_1, d_k)\} \tag{A.15}$$

for all $r \in \mathcal{R}_\delta$.

From Relation (A.14) $u_\delta(\mathbf{b}_\delta, d_\delta) \leq u_k\,(b_k \mathbf{e}_1, d_k)$ for all $k \in \mathcal{K}_{\delta,c}$, in particular

$$u_\delta(\mathbf{b}_\delta, d_\delta) \quad \leq \quad \min_{k \in \mathcal{K}_{\delta,c}} \{u_k\,(b_k \mathbf{e}_1, d_k)\}\,.$$

Finally by using relation (A.15)

$$\begin{aligned} u_\delta(\mathbf{b}_\delta, d_\delta) \quad &\leq \quad \min_{k \in \mathcal{K}_{\delta,c}} \{u_k\,(b_k \mathbf{e}_1, d_k)\} \\ &\leq \quad \min_{k \in \mathcal{K}_{\delta,r}} \{u_k(b_k \mathbf{e}_1, d_k)\} \\ &\leq \quad u_k(b_k \mathbf{e}_1, d_k) \end{aligned}$$

for every composite commodity number $r \in \mathcal{R}_\delta$ of class δ and for all $k \in \mathcal{K}_{\delta,r}$. Therefore for every composite commodity number $r \in \mathcal{R}_\delta$ of class δ and for all $k \in \mathcal{K}_{\delta,r}$

$$u_k((b_k + \lambda_r)\mathbf{e}_1, d_k) \quad > \quad u_\delta(\mathbf{b}_\delta - \lambda_r \mathbf{e}_r, d_\delta)$$

with $\lambda_r > 0$ and from Definition 5.6 no fair trade (decomposition) of any composite commodity of class δ is possible. ■

Bibliography

[1] E. Altman, R. El Azouzi, and O. Pourtallier. Avoiding paradoxes in routing games. In *Proceedings of the Seventeenth International Teletraffic Congress*, pages 24–28, Brazil, September 2001.

[2] E. Altman, J. Galtier, and C. Touati. Utility based fair bandwidth allocation. *International Conference on Networks, Parallel and Distributed Processing and Applications*, October 2002.

[3] N. Anerousis and A.A. Lazar. Virtual path control for ATM networks with call level quality of service guarantees. *IEEE ACM Transactions on Networking 6(2)*, pages 222–236, April 1998.

[4] G.R. Ash. Traffic engineering & QoS methods for IP-, ATM-, & TDM-based multiservice networks. *IETF Draft http://www.ietf.org/ietf/draft-ietf-tewg-qos-routing-01.txt*, 2001.

[5] S.A. Berezner, J.M. de Kock, A.E. Krzesinski, and P.G. Taylor. Local reconfiguration of ATM virtual path connection networks. In *Proceedings of the 5th International Conference on Broadband Communications, Hong Kong*, pages 621–630, November 1999.

[6] L.A. Boland. *Methodology for a New Microeconomics: the Critical Foundations.* Allen and Unwin, Inc., USA, first edition, 1987.

[7] J.J. Combrink. Co-operative monopolistic market-based routing in MPLS networks. *Southern African Telecommunication Networks and Applications Conference*, September 2002.

[8] B. Davie and Y. Rekhter. *MPLS Technology and Applications.* Morgan Kaufmann Publishers, USA, first edition, 2000.

[9] M.D. Davis. *Game theory - A Nontechnical Introduction.* Basic Books, Inc., Publishers, first edition, 1970.

[10] M. de Prycker. *Asynchronous Transfer Mode Solution for Broadband ISDN.* Ellis Horwood Limited, England, first edition, 1991.

[11] R.C. Epping. *A Beginner's Guide to the World Economy.* Vintage Original, first edition, 1992.

[12] J.W. Friedman. *Oligopoly and the Theory of Games.* North-Holland Publishing Company, first edition, 1977.

[13] K.D. George and J. Shorey. *The Allocation of Resources: Theory and Policy.* George Allen and Unwin (Publishers) Ltd., first edition, 1978.

[14] M. A. Gibney and N.R. Jennings. Dynamic resource allocation by market-based routing in telecommunications networks. *2nd International Workshop on Multi-Agent Systems and Telecommunications, Paris, France*, pages 102–117, 1998.

[15] A.L.G. Hayzelden and J. Bigham. Software agents in communications network management: An overview. *Knowledge Engineering Review*, August 1998.

[16] S. Huiasu and M. Malitza. *Coalition and Connection in Games.* Pergamon Press Ltd, England, first edition, 1980.

[17] B. Jennings and Å. Arvidsson. Co-operating market/ant based multi-agent systems for inteligent network load control. *IATA*, 1699:62–75, 1999.

[18] D.M. Kreps and R. Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.

[19] K. Kuwabara et al. An equilibrium market-based approach for distributed resource allocation and its applications to communication network control. In Scott Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, chapter 3, pages 53–73. World Scientific, USA, 1996.

[20] P. Morris. *Introduction to Game Theory.* Springer-Verlang New York, Inc., New York, USA, first edition, 1994.

[21] C.H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual Symposium on the Theory of Computation*, pages 749–753, Hersonissos, Crete, Greece, July 2001.

[22] H. V. D. Parnunak. "Go to the ant": Engineering principles from natural multi-agent systems. *Annals of Operations Research 75 (Special Issue on Artificial Intelligence and Management Science)*, pages 69–101, 1997.

[23] L.A. Petrosjan and N.A. Zenkevich. *Game Theory.* World Scientific Publishing Company. Pte. Ltd., Singapore, first edition, 1996.

[24] J. Rosenmüller. *The Theory of Games and Markets.* North-Holland Publishing Company, Elsevier North-Holland, first edition, 1981.

[25] D.L. Rubinfeld R.S. Pindyck. *Microeconomics.* Prentice-Hall, Inc., fifth edition, 2001.

[26] Telkom RSA. Annual report. *http://www.telkom.co.za/annualreports/annual_report2001/AnnualReport2001.pdf*, 2001.

[27] T. Sandholm and V.R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1-2):99–137, 1997.

[28] S. Sarkar and L. Tassiulas. Fair allocation of utilities in multirate multicast networks. In *Proceedings of the 37th Annual Conference on Communication, Control and Computing*, October 1999.

[29] S. Sarkar and L. Tassiulas. Scheduling for quality of service based service differentiation. In *Proceedings of the 17th International Teletraffic Conference*, 2002.

[30] Y. Shoham. *Introduction to Multi-Agent Systems.* (In preparation for publication), Stanford University, first edition, 2003.

[31] D.P. van Soest, J.A. List, and T. Jeppesen. A shadow price approach to measuring the cost of increased environmental stringency. *Systems Organising and Management, Seminar Series C*, Spring 2002.

[32] M.P. Wellman. A market-orientated programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

[33] M.P. Wellman. Market-oriented programming: Some early lessons. In Scott Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, chapter 4. World Scientific, USA, 1996.

[34] M. Wooldridge and N.R. Jennings. Pitfalls of agent-oriented development. In *Proceedings of the 2nd International Conference on Autonomous Agents (Agents'98)*, pages 385–291, USA, September 1998.