

Chapter 5

Agent-Based Models of Supply Chains

B. Behdani, K.H. van Dam, and Z. Lukszo

Abstract Based on the modelling steps discussed in Part I, this chapter aims to present ways in which agent-based simulation models of supply chains can be developed and used to improve the performance of these systems in both normal and abnormal situations. An industrial supply chain with a network of several independent companies is a good example of a socio-technical system. The physical and social networks of the actors involved in their operation collectively form an interconnected, complex system in which the actors determine the development and operation of the physical network and, likewise, the physical network affects the behaviour of the actors. In this type of system, the many interactions taking place in the social and physical subsystems can result in the complex, dynamic behaviour of the supply chain as a whole. Accordingly, any attempt to improve the functioning of the supply chain requires a comprehensive understanding of this behaviour under different supply network configurations. Most of the current approaches to the modelling and simulation of supply chains do not capture the rich socio-technical dynamics present. The agent-based modelling approach, however, seems to be very promising as a means to address this complex behaviour. To demonstrate its applicability, we will present agent-based simulation models for two different industrial supply chains: an oil refinery and a multi-plant chemical enterprise. Using the models described in this chapter, the outcomes of the system under a broad range of possible agent behavioural rules and environmental events can be explored, and improved levels of system functioning can be identified.

5.1 Introduction

Supply chains present an interesting research subject from a socio-technical systems perspective. In these chains, complex production technologies interact with geographically distributed, intelligent, autonomous entities; each with its own dynamics, goals, desires and plans. There is a significant challenge in modelling such systems which function in dynamic, stochastic, socio-economic environments and

B. Behdani (✉)
Delft, the Netherlands
e-mail: b.behdani@tudelft.nl

show intra-organisational and inter-organisational complexity. The focus in this chapter is on *operational decision-making*, which has a relatively short time frame and in which decisions and actions are directly linked to the functioning of production units and the transportation of mass flows.

This chapter presents two case studies: one is a project designed to help an oil refinery in dealing with disturbances and ensuring smooth operation at minimal cost; the other is support for a multi-plant chemical enterprise, aimed at helping them to respond to different types of abnormal situations, including unexpected plant shut-down. The two systems lie in a similar domain (i.e. supply chains) but their individual characteristics and the specific problem formulations result in different models that are used in different ways. The models have a common starting point, but these diverge into different paths because they address different problems and thus have different conceptualisations and implementations. This chapter will demonstrate the link between the problem and domain characteristics of the two cases and the choices made during the development and use of the models. The cases illustrate the possibilities offered by agent-based models in the field of supply chain management. Furthermore, the description of the modelling choices for these case studies demonstrates how agent models of such complex systems can be developed in practice.

5.1.1 Supply Chains

A set of companies whose activities can together be characterised as the production, service provision, distribution or trade of raw materials, by-products or market goods is regarded as an “industrial network”. A supply chain is a good example of such a network. Along each *supply chain*, numerous links and nodes exist, grouping material suppliers, production facilities, distribution services, warehouses, service centres and customers into a single cohesive entity. The *actors* in the supply chain coordinate their activities and collaborate together to improve the efficiency of the whole system in general and their own enterprise in particular. The *physical* elements of the supply chain encompass processing equipment which may be located at different sites. As the complexity of this system increases, the likelihood of a breakdown at any point in the network increases as well. This can happen, for example, due to the heavy interdependence of the network components, lengthened supply and distribution lines or heightened competition on the market. As a consequence, different independent decision-makers in the supply chain might come into conflict; something which can lead to the instability of the system and place the network in an abnormal situation.

5.1.2 Abnormal Situation Management

Abnormal situations in supply chains encompass a range of events outside the “normal” operating modes, including human error, fires, shipment delays, (unplanned)

maintenance and equipment failure. Note that disturbances may be part of normal operation. When the normal operating mode is interrupted by a disturbance or series of disturbances, such that the network cannot function properly, this is called an *abnormal situation*. In such a case, the planned production targets may not be met unless swift action is taken to minimise the negative effects. Abnormal situations are rarely local; a disruption in one part of an industrial network can influence the performance of other components and the network as a whole. This raises the following questions:

1. To what extent does one part of a network influence a system's overall performance?
2. What are suitable strategies and control mechanisms for coping with and recovering from abnormal situations?

This chapter describes how agent-based simulation models of supply chains are developed and used to support decision-makers in the management of abnormal situations in supply chains and to help answer these critical questions.

5.1.3 Supply Chain Modelling

Agent-based modelling is not the only modelling paradigm used for *supply chain modelling* and simulation. Over the past two decades, many other types of analytical models have been presented to study the dynamic behaviour of supply chains during normal and abnormal situations. Most of these models are based on mathematical programming and operations research approaches. Overviews of some key publications are given by Tayur et al. (1999) and Min and Zhou (2010). Although these analytical models cover a broad range of areas in supply chain management, from functional aspects (e.g. inventory management) to inter-organisational issues (e.g. collaboration along the supply chain), most of them do not explicitly take into account the social aspects of the system; hence, the dynamics of the modelled system are dominated by the physical realities rather than the interactions between decision-makers.

The socio-technical nature of supply chain problems, however, is the motivation behind using the agent-based modelling approach as an alternative paradigm. The benefit of this approach is that it takes an actor-centric perspective instead of modelling only the outcomes of their activities. The actions of each actor (represented as an agent) and the interactions between them are explicitly represented in such models, and consequently the behaviour of the entire system emerges (van Dam 2009). Swaminathan et al. (1998) provided one of the first efforts for modelling supply chains using agent-based modelling: a flexible, agent-oriented framework which enables rapid development and customised, agent-based decision support tools for supply chain management by using a library of structural and control elements (e.g. supply chain roles and policies). Other examples of using agent-based models for supply chains are given by Julka et al. (2002a, 2002b); Kaihara (2003); Sirola et al.

(2003) and the pioneering work of Gjerdrum et al. (2000). Moyaux et al. (2006) present a review of the applications of an agent-based paradigm for supply chain management. More recent work on modelling supply chains using the agent-based modelling approach includes Srinivasan et al. (2006), Mele et al. (2007), van Dam et al. (2008) and Behdani et al. (2010a, 2010b).

5.1.4 Chapter Organisation

The rest of this chapter will follow the ten model building steps, described in Chap. 3 of this book, for two case studies. The first two steps are discussed for the two models in Sects. 5.2 and 5.3, respectively. Afterwards, steps 3 through 10 are presented integrally in Sects. 5.4 through 5.11. In each of the steps the development of both models will be addressed, focusing on those aspects in which the development is influenced by specific choices stemming from different problem statements, application domains and the different conceptualisations that flow from this. The chapter thus does not provide a complete picture of the development of each of the models; instead it highlights the specific choices that were made in two related models and demonstrates where the models converge or diverge. In Sect. 5.12 these findings are summarised and concluding remarks on the lessons learned on agent-based modelling and supply chain management are given.

5.2 Oil Refinery Supply Chain

In this section the first two modelling steps ('Problem formulation and actor identification' and 'System identification and decomposition') are followed in a case study of an oil refinery supply chain.

5.2.1 Step 1: Problem Formulation and Actor Identification

A hierarchy of decisions has to be made in managing the supply chain of an oil refinery: strategic (e.g. capacity investments, the addition of units, technology upgrades and supply chain reconfiguration), tactical (e.g. production planning, policy evaluation, disruption management) and operational (e.g. procurement, storage, scheduling, throughput level) (Pitty et al. 2008). Disturbances in the supply chain, including equipment failures and shipping delays, have consequences for the operation of a plant. What is the best response to a certain situation, and which changes in operation or emergency procurement are required and efficient? During nominal process conditions the operation is optimised by choosing which crudes to buy, how much crude is needed and from which supplier to order. The mode of operation is scheduled based on predicted demands, and the throughput for operation of

the refinery is set based on actual consumer demand. When an abnormal situation manifests itself, this normal approach is no longer adequate.

In the supply chain domain as considered in this case study, the following actors can be identified:

- Problem owner:** The supply chain manager is responsible for overseeing the operation of the entire supply chain. However, in some circumstances, the operator of a single node in the supply chain could be considered to be the problem owner.
- Stakeholder:** In addition to the supply chain manager, stakeholders include the managers of the various supply chain entities, shippers, consumers, local governments, national governments, investors, port authorities, etc.
- Modeller/analyst:** The authors of this chapter are together considered to be the modeller, whose role is to develop a model for the problem owner to help solve the problem as motivated above.
- Facilitator:** In this case the role of the facilitator is not applicable, because the model building process has been a more or less straight-forward exercise in a domain which has been well studied in literature. In future development and extensions of these models, a facilitator could be required to lead the discussions between stakeholders and modellers.
- Domain expert:** The domain expert for this case study includes literature sources in process systems engineering and computer-aided chemical engineering, as well as several experts in the academic field of supply chain research.

With this problem owner and these other actors in mind, the system is identified and broken down in the next section.

5.2.2 Step 2: System Identification and Decomposition

An *oil refinery supply chain* begins with the oil reservoirs, both onshore and offshore. Crude oil is tapped from these sites and then transported to various refineries around the world, mostly by pipelines or large ships (very large crude carriers, or VLCC). Transport times of crude are relatively long; it takes four to six weeks for a VLCC carrying crude oil from the Middle East to reach refineries in Asia, for example. The crudes are then processed in crude distillation units (CDU) and separated into fractions based on their boiling points. These fractions are processed further in different downstream refining units such as reformers, crackers and blending pools to arrive at the various products. A single crude mix may yield numerous products and their variants through a suitable alteration of processing conditions. Hence, refineries must adapt their operations to the different crude batches to meet the required product specifications.

The refinery occupies a pivotal position in the supply chain, as its functional departments initiate and control the interactions with the external entities, these being oil suppliers, third party logistics providers, shippers, jetty operators and customers. The other stakeholders as identified in Sect. 5.2.1, such as the national governments, investors and port authorities, are considered to be outside the scope of the model, given the problem statement. They may still play an important role in the interpretation of model results and they may (indirectly) be affected by the decisions made by the problem owner, but they do not have to be included as actors in the *model* of the system. The same holds true for the supply chain manager, who does not need to be included in the model itself. Even though the supply chain manager is a member of the refinery company, which is included as an actor in the model, his tasks and behaviour are not captured in the model because the model is not designed to offer *online* decision support. The supply chain manager is the *user* of the model and can choose and validate decisions while experimenting with different actions.

The operation of the refinery supply chain requires various decisions in every cycle: what mix of products to make, which crudes to purchase and in what quantities, which mix to process and in what processing mode, etc. Different actors are responsible for the different decisions (Julka et al. 2002b). These actors and their interactions are shown in Fig. 5.1. The entities, shown as blocks in Fig. 5.1, communicate with each other through information flows (dashed arrows) in order to control the material flows (solid arrows).

The refinery's physical units (shaded blocks) may be further sub-divided into storage units such as crude and product tanks and processing units such as the CDU, reformer, cracker and blending pools. The functioning of these units and other supply chain activities is overseen by the functional departments: the storage department and the operations department. The actors and physical systems are shown in Table 5.1. Note that some actors, for example the logistics department, do not own or control a physical system, but they do have their own specific tasks and communicate with the other actors. All actors in the model are briefly described below:

Refinery company	is the overall institution which owns the refinery. It can be broken down into various departments. Relevant properties include the company's fixed and liquid assets.
Operations department	is in charge of production and operates the oil refinery. It selects a mode of operation (i.e. a recipe) and the throughput of the refinery based on sales. Properties include the characteristics of the refinery units.
Storage department	is responsible for managing the storage tanks for crudes and end products and monitors levels to ensure that limits are not violated. The storage department also alerts the operations department if there is not enough crude in storage during operation.
Sales department	is in contact with consumers about the sale of end products. This department also makes an estimate of

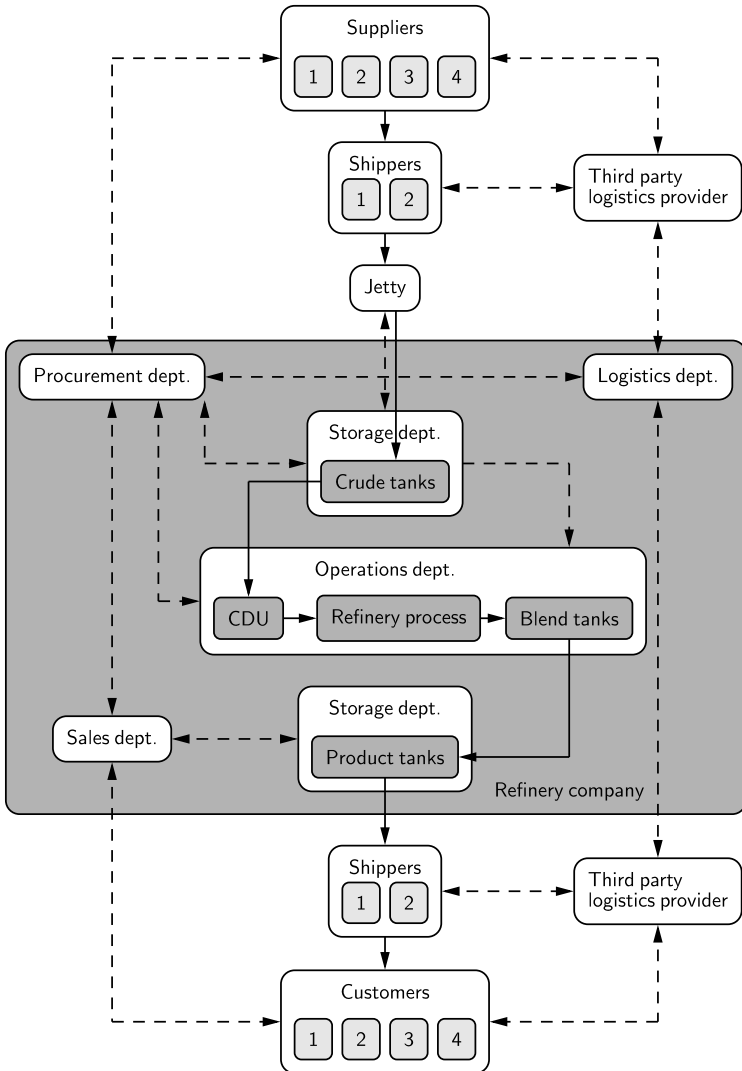


Fig. 5.1 Schematic depiction of an oil refinery supply chain (adapted from Pitty et al. 2008). Arrows with a solid line represent material flows; arrows with a dashed line represent information flows

Table 5.1 Agents and Physical Nodes and their relationships in the refinery supply chain case

Agent	Relationship	Physical Node
Refinery company	owns	Refinery units (incl. CDU)
Refinery company	owns	Raw materials storage tanks
Refinery company	owns	End product storage tanks
Operations dept. (Refinery)	controls	Refinery units (incl. CDU)
Storage dept. (Refinery)	controls	Raw materials storage tanks
Storage dept. (Refinery)	controls	End product storage tanks
Sales dept. (Refinery)	–	not applicable
Procurement dept. (Refinery)	–	not applicable
Logistics dept. (Refinery)	–	not applicable
3rd party logistics provider	–	not applicable
Shipper	–	not applicable
Supplier	owns	Oil wells
Jetty owner	owns	Jetty
Consumer	owns	Consumer installation

	the demand to be used for forecasting and later finalises the contracts for the actual sales. Properties include the prices of end products.
Procurement department	buys crude oil from the suppliers based on the forecast demand and the planned mode of operation.
Logistics department	is in charge of arranging shipping for the goods purchased by the procurement department as well as the shipment of end products to consumers.
Third party logistics provider	forms the link between the shippers and the logistics department.
Shipper	moves goods from one place to another.
Supplier	sells crude oils.
Jetty owner	is in charge of the jetty where crude oil tankers dock to unload the oil into the storage tanks. Furthermore, this actor manages the scheduling of vessels and operates the pump.
Consumers	have a demand for the end products of the refinery.

Other elements in the system are considered to be part of the “environment” of the model. This includes aspects which are assumed not to be influenced by components within the system. This means that, for example, the prices of crude oils and the demand for the various end products are assumed to be givens and are not generated within the model. Emissions during the supply chain activities (including production and transport) are not taken into account in this system description, as they are not relevant for the questions asked of this model.

5.3 Multi-plant Enterprise Supply Chain

In this section the first two modelling steps ('Problem formulation and actor identification' and 'System identification and decomposition') are followed for a case study of a multi-plant enterprise supply chain.

5.3.1 Step 1: Problem Formulation and Actor Identification

The *multi-plant enterprise* can be viewed as a *multi-level system*, whether hierarchically interconnected or decentralised, with a number of operational regimes at the various system levels. Usually, at each level of the decomposed system (i.e. functional level, plant level, enterprise level) local performance objectives are defined which should, preferably, not be restricted to the optimisation of local goals but rather aim at optimally contributing to the overall goal. However, the relation between local and overall system performance objectives may be rather fuzzy, especially since the overall objective is often not defined in detail and involves a longer time horizon. The local objectives are generally more detailed, concerned with a shorter time horizon and often with the specific interests of an individual actor (e.g. a business unit).

To facilitate the optimisation of the performance of the system as a whole, a coordinator may be required to supervise local decision-making in its relation to the overall goal. Therefore, a complex network of production plants needs a special form of operations management to coordinate plant-dependent activities and to ensure that the enterprise as a whole can realise its optimum performance (Bhatnagar et al. 1993). Operations management involves, among other things, decision-making regarding global enterprise resources planning systems: how to best operate a multi-plant enterprise and at the same time how to operate the separate plants in an optimal way.

These decisions can be categorised into long-, medium- and short-term decisions. The long-term decisions basically deal with operations strategy, product planning and the design of facilities and processes. The medium- and short-term decisions are concerned with the planning and control of production activities; more precisely, procurement and inventory management, sequencing and scheduling, quality management and maintenance management (Lee-Post and Chung 2008). Although some of those decisions can be made by an individual company, in a multi-plant enterprise each plant is a part of a corporate network and its relations with other plants are as important as its internal operations. With appropriate operations management, the enterprise can respond more quickly to dynamic market events, reduce its operating costs and increase customer satisfaction. Moreover, during an abnormal situation (e.g. unexpected shutdown of one production plant) it can be more resilient and recover from disruption more rapidly.

A multi-plant enterprise model can be helpful in finding a deeper view of network components' behaviours and their effects on the enterprise's overall performance.

Such a model can help decision-makers to study alternative scenarios and also to improve the operation with respect to some important performance indicators.

Even though the problem domain is different from the one discussed in Sect. 5.2 (the single production site vs. multiple sites within one enterprise being the main distinction) the same actors can be identified at the highest level (see Sect. 5.2.1). As opposed to the refinery manager in the previous case, the problem owner in this case study, the supply chain manager, is responsible not just for overseeing one specific link in the supply chain, but for the operation of a number of sites within one company.

5.3.2 Step 2: System Identification and Decomposition

The *multi-plant enterprise* that will be modelled in this section is a global lube oil supply chain with a global sales department which directly interacts with customers and three production plants at different geographical locations (Zhang et al. 2008). Each production plant itself has several functional departments, each having a specific role and performing certain tasks; the performance of each production plant is a result of its departments' behaviours and their interactions. All production plants operate on a "make-to-order" basis and can produce three types of products after receiving an order from the customer. Each product is produced from eight different raw materials. The goal is to fulfil a set of customer orders in the fastest possible time through assigning them to different production plants and coordinating the behaviour of different departments in each plant.

So, in this enterprise, the involved actors can be viewed on three levels:

- | | |
|------------------|--|
| Global level | There are three actors at the global level (the customer, the enterprise and the supplier). |
| Enterprise level | The manufacturing enterprise consists of the global sales department and a number of plants. |
| Plant level | Each plant has six different functional departments (scheduling, operations, storage, packaging, procurement and logistics). |

The behaviour of each of these actors is described in more detail in Fig. 5.2. Based on this system description, an agent-based model with ten agent types has been developed. Table 5.2 presents the social and physical sub-systems (i.e. agents and technologies) in the agent-based model. It is worth mentioning that the manufacturing enterprise is not considered to be a separate agent but a virtual one consisting of the plants' agents and the Global Sales Department (GSD).

5.4 Step 3: Concept Formalisation

The formalised concepts of the socio-technical systems domain, as described in the ontology in Sect. 4.2, were used for both supply chain case studies. For the development of the oil refinery supply chain model no major changes to the generic ontology

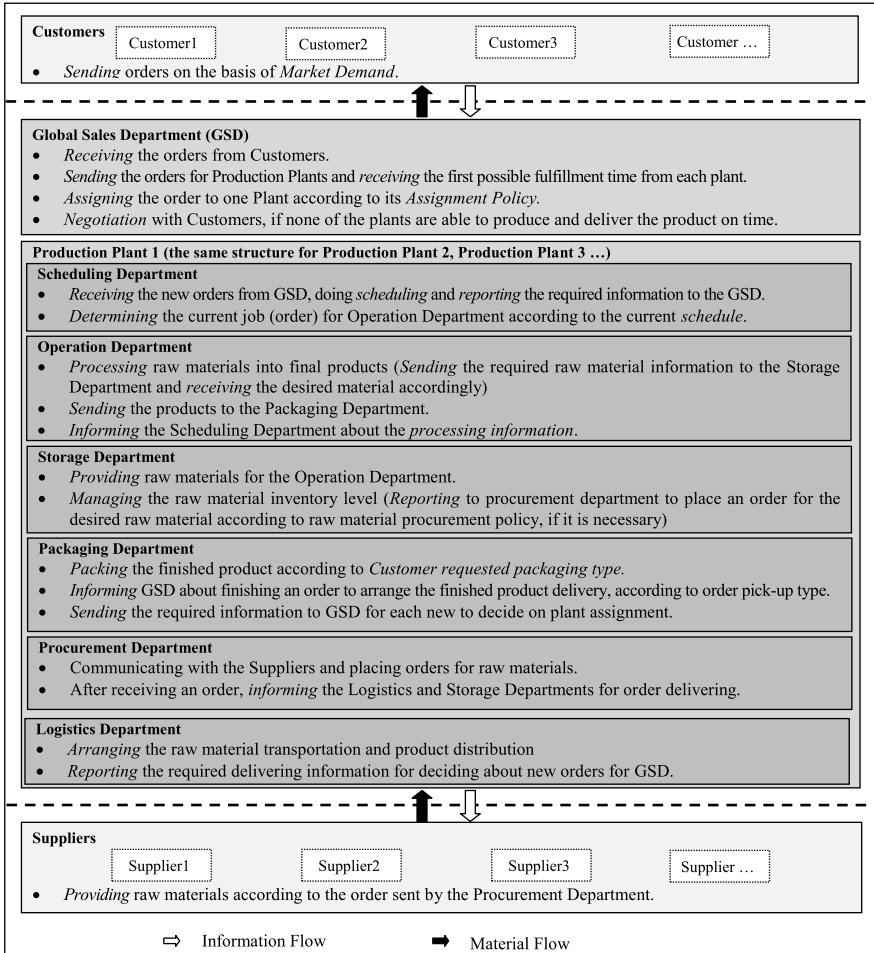


Fig. 5.2 Main actors and their behaviours/interactions in a lube oil multi-plant enterprise

were needed. All the key classes needed to define the system were already in place, based on other case studies (several of which are also included in this book). Mostly minor additions were needed, such as adding properties to the Transport Contract for more detailed registration of transport delays and payment. One of the more significant changes was the adoption of the Ownership class. Initially the Ownership edge represented a link between an Agent and a Physical Node (e.g. the ownership of a Technology by an Agent) only, but from the system decomposition as described in Sect. 5.2.2 it is clear that one Agent can also own another Agent. This is the case when a department of a company is a part of a larger enterprise. The ‘to’ property was thus relaxed to allow instances of not only Physical Nodes, but also Social Nodes. See Fig. 5.3 for a graphical representation of the new relationship.

Table 5.2 Agents and Physical Nodes and their relationships for the multi-plant enterprise case

Agent	Relationship	Physical Node
Customer	owns/controls	Customer facilities
Global sales department (GSD)	–	not applicable
Production plant	owns	Production facilities
Production plant	owns	Storage facilities
Production plant	owns	Packaging facilities
Scheduling dept. (Production plant)	–	not applicable
Operations dept. (Production plant)	controls	Production facilities
Storage dept. (Production plant)	controls	Storage facilities
Packaging dept. (Production plant)	–	not applicable
Procurement dept. (Production plant)	–	not applicable
Logistics dept. (Production plant)	–	not applicable
Supplier	owns/controls	Supplier facilities

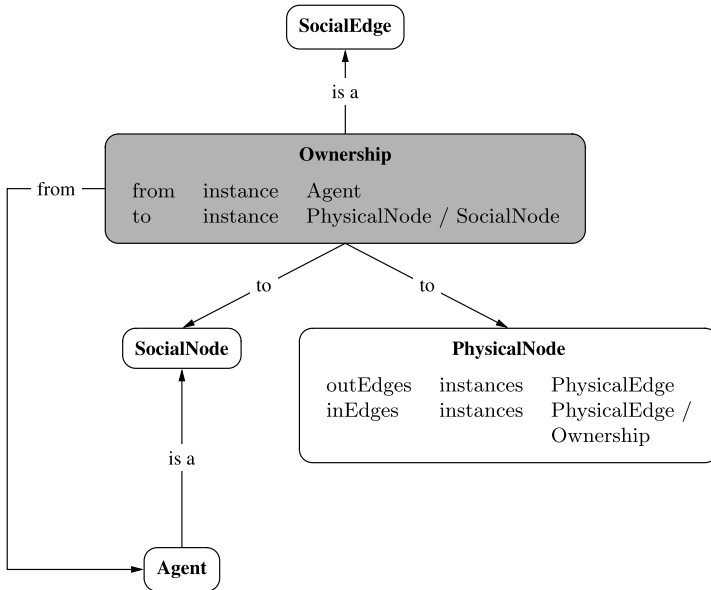
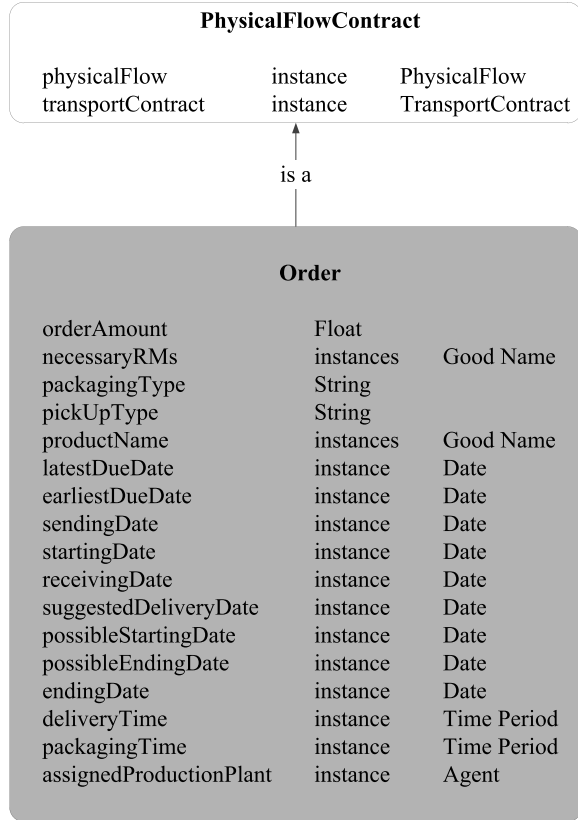


Fig. 5.3 An Ownership edge can connect a Social Node and a Physical Node, forming the link between social and physical networks. In this way we can define an Agent as being the owner of a Technology or another Physical Node. An Ownership edge can also be used between an Agent and another Social Node (e.g. to define that a company consists of several departments or that a department consists of specific people)

Fig. 5.4 Order ontology



The multi-site case study uses many of the same concepts as the refinery case. Furthermore, the formalisation of several concepts which had not previously been used in other models was required. As one example of a new concept that had to be added, let us take the concept of an *order*. The *order* concept is used to formalise the short-term purchasing transactions between different actors in the supply chain. So, for instance, a customer places an order for a specific final product with a production plant, which in turn further places orders for raw materials with its own suppliers. This concept is especially important when we aim to model a make-to-order production environment (such as a *multi-plant enterprise*). In this type of situation, every good or service can be made only after receiving an order from a customer. In our formalisation we consider *order* to be a sub-class of *PhysicalFlowContract*; therefore, it will inherit the attributes of its super-classes such as *from*, *to*, *physicalFlow*. In addition, it has its own specific attributes, including *deliveryTime*, *latestDueDate*, *necessaryRMs*¹. Other attributes of this concept are shown in Fig. 5.4.

¹RM stands for Raw Materials.

These new contributions could be re-used in future model developments, if applicable. This could be in additional supply chain models, but these concepts might also be useful for modelling other types of socio-technical systems.

5.5 Step 4: Model Formalisation

In this section the behaviour of the actors in the supply chain is formalised to create the model of the supply chain systems. Following the agent paradigm, the tasks are distributed between the agents. Some tasks therefore have to be split into several subtasks (requiring communication between the agents). A schedule is made so that some processes (e.g. procurement) only occur at certain intervals while others (e.g. production) happen at each time step of the simulation. Events such as the arrival of a VLCC at the jetty are monitored at each time step. After the brief introduction of the roles of the different actors (see for example Sect. 5.2.2), two typical tasks will now be considered and addressed in more detail, namely procurement from the refinery model and order assignment from the multi-site case.

5.5.1 Procurement

The procurement agent is responsible for buying a mixture of crudes well ahead of when these crudes would be required in production. The following narrative describes the steps this actor takes to complete the procurement task.

Before the procurement agent can place an order from the suppliers, it needs to determine which crudes to buy and how much of each product. First, the procurement department asks the sales department what the predicted demand is for each of the products produced by the refinery (i.e. gasoline, kerosene, diesel and fuel oil). After it has received the volumes that are expected to be produced, the procurement department determines what product has the highest crude demand, as this will be crucial in choosing the crude mixture and amounts. The average yields of the possible crude mixtures are used to calculate which product will require the highest amount of crude oil. The operations department is asked which recipe it will use to produce the product with the highest crude demand, considering all other products bought by the consumers as by-products, because this approach ensures that at least enough is produced. Next, from the recipe selected by the operations department, it can be deduced which crudes need to be bought and in which quantities.

The procurement department has multiple strategies to decide how much crude to buy, two of which are included in this model. In the first policy the amount of each crude to buy is solely determined by the recipe selected and the total amount needed to produce the expected order. In the second, more advanced, policy, the procurement department takes into account the current levels in the storage tanks (obtained from the storage department) as well as the amount in the tank which

is already reserved for the current production cycle (information obtained from the operations department). A safety stock level is kept as a minimum volume to be kept in the storage tank. After it has been determined which crudes need to be bought and how much of each one of these crudes is required, the procurement department approaches the market to find the best seller. It does this by asking all potential sellers for a contract and selecting the cheapest one, which is signed to form a trade agreement.

There are a number of prerequisites, such as that the sales department has completed its forecast of the demand. For the process to be completed several other tasks need to be performed, including asking the logistics department to arrange transport for the crudes that have been bought. These tasks are scheduled to be performed by other agents.

This narrative can be translated to pseudo-code as shown in Algorithm 5.1, considering only the first procurement strategy and omitting the selection of the best contract for simplification.

All other behaviours in the model are also split up in a similar fashion between the different agents. Another example of this is the selection of the production mode by the operations department, which is based on the forecasts made by the sales department and the crudes selected by the procurement department. Where possible, modelling the behaviour of the agents was based on existing behavioural rules for trading. However, additional rules had to be implemented: for example, for various procurement policies (e.g. forecasting of demand deciding on procurement), for scheduling which Operational Configuration to use, and for the activities of the jetty which had not been used in earlier models (van Dam et al. 2009a).

5.5.2 Order Assignment

Similar to those in the *oil refinery case*, the actors in a multi-plant enterprise do different sorts of activities and participate in different processes. Basically, the main determinant of their behaviour is the policies they have. As an example, for a procurement agent the “reorder point” is a policy for managing the raw material level and, according to it, this agent must contact another agent (supplier) to place raw material orders at appropriate times.

To illustrate how different agents may participate in a specific process and coordinate their activities in the multi-plant enterprise, the “order assignment process” is described here:

After receiving a new order from customers, the GSD will assign it to one of the available production plants. For this purpose, the GSD sends the order information to the scheduling department of production plants. Each scheduling department then replies with the earliest date by which the plant can make the product and deliver it to the customer. Based on the replies, the global sales department will assign the customer order to one production plant according to its assignment policy. This assignment policy can be a time-related or cost-related policy. For example, if the

Algorithm 5.1 Procurement process

```

1  get forecastDemands from SalesDept
2
3  set productWithHighestCrudeDemand to empty
4  set highestCrudeDemand to 0
5
6  for goodName from forecastDemands
7  do
8    set yield to averageYields for this goodName
9    set demand to forecastDemands for this goodName
10
11   if (demand/yield) > highestCrudeDemand
12     set highestCrudeDemand to demand/yield
13     set productWithHighestCrudeDemand to goodName
14   end if
15 end for
16
17 get operationalConfiguration from
    operationsDepartment
18
19 if procurementStrategy equals 1
20   for crude from operationalConfiguration
21   do
22     set amountToBuy to operationalConfiguration .
        relativeAmount * highestCrudeDemand
23
24     buy amountToBuy of crude
25
26   end for
27 end if

```

assignment policy is “First Completion Date Policy”, then the order is assigned to the production plants with first possible fulfilment date. If none of the plants are able to produce and deliver the product on time, the GSD will start the negotiation with customer for extending the order due date.

This process can be described with Algorithm 5.2. This algorithm describes the sequence of activities in the “order assignment process” and the actors which perform these activities. The behaviour of different actors and other processes in the system (such as “Order Replenishment Process” or “Abnormal Situation Process”) can be expressed in similar algorithms and pseudo-code.

Algorithm 5.2 Order assignment process

```

1 announceOrder by globalSalesDepartment
2
3 for ProductionPlant from AllAvailableProductionPlants
4 do
5     obtain Scheduling Department of Production
        Plant
6     obtain Schedule from Scheduling Department for non-
        Assigned Orders and new Orders
7     obtain Processing Timing from Scheduling Department
        for new Orders
8 end for
9
10 determineFirstPossibleDate for newOrder by
        globalSalesDepartment
11
12 if LatestDueDate of newOrder >= firstPossibleDate
13     assignTheOrderToTheFPTPlant
14 else
15     startNegotiationOnExtendingDueDate
16 end if

```

5.6 Step 5: Software Implementation

Both models have been implemented in a similar fashion using Java. The software implementation is based on the software framework including shared basic classes for the ontology concepts and several aids for visualisation and data analysis. This shared framework was updated following the specifications for these particular case studies. For example, following the extension of the Ownership class in Sect. 5.4, the Java class representing an agent was adjusted to implement the consequences of an agent being owner by another agent. For example, if an agent which is a subsidiary of another agent has to make a payment, it will use the “wallet” of the agent which owns it (this may be recursive).

There are no specific hardware requirements, as the models are designed to run on a standard desktop computer.

5.7 Step 6: Model Verification

The oil refinery model has been tested extensively. A benchmarking study was executed in which the agent-based model was compared with an equation-based model

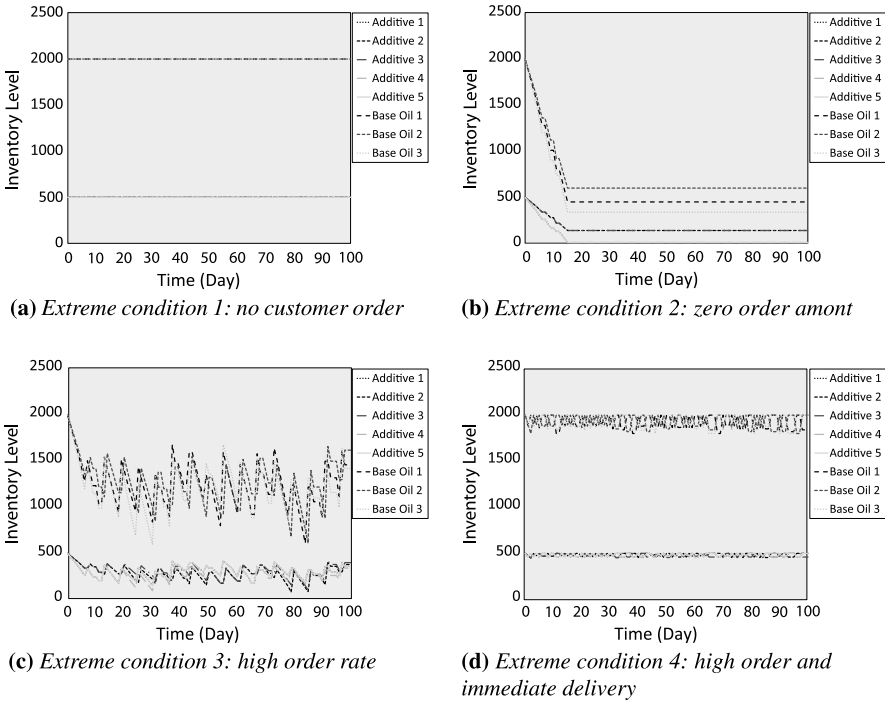


Fig. 5.5 Model behaviour (inventory level for production plant 1) under different extreme condition tests

of the same supply chain (van Dam et al. 2009a). The aim of the benchmarking study was to compare the modelling paradigms (i.e. the agent-based and equation-based paradigms) and to learn about the advantages and disadvantages of the different approaches. In order to benchmark the modelling paradigms, it was necessary to demonstrate that the models under study are comparable. During this process a number of differences between the results of earlier versions of the model were identified, which could then be checked in more detail. Several mistakes made during the formalisation and implementation phases were then able to be corrected in this way. Finally, a numerical analysis was performed, proving that the two models show the same behaviour (van Dam et al. 2009a), raising confidence that the agent-based model was implemented correctly.

For the multi-plant enterprise case, tests of extreme conditions and behaviour sensitivity were conducted. In these tests—in which selected parameters had extreme input values, such as zero—the model should behave according to our expectations. Figure 5.5 shows the inventory levels of different raw materials for production Plant 1 under four extreme conditions, as discussed below:

- The test related to Fig. 5.5a was conducted with the assumption of no consumer orders being placed for the duration of the simulation. So, as we expected, the

inventory level for raw materials is unaltered. The other two production plants produced similar results.

- In Fig. 5.5b, we set the reorder amount to zero. The results show that the production plant will use its initial raw material in the storage facilities, but as there is no replenishment, the raw material level will be zero for the rest of simulation time.
- Figure 5.5c shows a case in which the reorder point is set to 0.9; i.e. when the raw material inventory level falls below 90 percent of storage capacity, a new raw material order is placed. The expected result is frequent raw material orders, and this is confirmed in the figure.
- Figure 5.5d shows the results of a test similar to the previous case but with immediate raw material delivery (the raw material delivery time is set to zero). Again, the model behaves as expected and the stock is always near maximum capacity.

Similar tests can be done by using extreme values for other parameters (e.g. order processing time, storage capacity or plant availability) to study whether those factors have the expected effect on the model outcomes.

5.8 Step 7: Experimentation

Now that the models have been designed and implemented, they can be used to run experiments that can help answer the main questions posed in Sects. 5.2.1 and 5.3.1. Below, we first address the design of the experiments for the oil refinery case study, followed by a discussion for the multi-site case.

5.8.1 Experimental Setup for the Oil Refinery Supply Chain

In a similar way to that of the problem formulated in Sect. 5.2.1, an experiment will now be set up to evaluate the best response to a *disturbance*, namely a delay in a shipment of raw materials.² There are different options for the criteria with which to choose the best alternative. As examples, one can look at the overall profit of the refinery (for a certain time frame), profit during the production cycle affected by the disturbance, other financial measures, but also non-economical criteria such as customer satisfaction. Profit P of the refinery was chosen, at 14 days after a disruption took place. This means that the effect of a disturbance will be simulated over the next two cycles of operation, during which new raw materials are ordered and products are dispatched. It is assumed that the impact of the disturbance will have worn off by then.

The objective function for a period of 50 days is defined as follows:

²For simplicity, it is assumed that the magnitude of the disturbance is known as soon as the disturbance occurs. In reality, this may involve uncertainty. Furthermore, currently delays are in the order of magnitude of days, but the granularity could be adjusted so that a delay could be expressed in parts of a day (e.g. hours) instead of full days of 24 hours.

$$\begin{aligned}
\max_{\bar{x}} P(\bar{x}, \bar{d}) = & \sum_{t=1}^{50} (Income_{\text{sales}}^t(\bar{x}, \bar{d}) - Cost_{\text{procurement}}^t(\bar{x}, \bar{d}) \\
& - Cost_{\text{transport}}^t(\bar{x}, \bar{d}) - Cost_{\text{maintenance}}^t(\bar{x}, \bar{d})) \\
& + Value_{\text{product stock}}^{t=50}(\bar{x}, \bar{d}) + Value_{\text{raw materials}}^{t=50}(\bar{x}, \bar{d}) \quad (5.1)
\end{aligned}$$

The function for the transportation costs, and therefore profit, is discontinuous because these are a function of the amount of crude procured, the capacity of the ships and the travel time. The transport cost is calculated per vessel, which could either be a very large crude carrier for long-distance shipping or a general purpose tanker with much smaller capacity for short hauls in the case of emergency procurement, and the unit landed-cost of crude (i.e. procurement plus transport costs) therefore follows a saw-tooth pattern. This discontinuity makes it more difficult to determine the right amount to buy, especially in combination with other measures such as switching to another recipe in the refinery.

In Eq. (5.1) the monetary value of the product inventories and raw material stocks at the end of the simulation run are included. The consequences of the disruption to future cycles within the time horizon are thus included in the cost function. One such example might be that if the response is to switch to another mode of operation without any emergency procurement, it is possible that during a later cycle the planned operation cannot be met. However, no new decisions following such new disturbances are assumed; a single response is formulated.

Faced with a disturbance, the problem owner must make a number of choices. Firstly, it has to determine if the disturbance has a significant effect on the operation of the supply chain. If the effect is deemed minor, no action may be necessary, but if it is not able to execute the previously planned schedules due to insufficient crude, corrective action may be required. A disturbance in the supply of crudes can be addressed by changing the operating mode or the throughput, or by emergency crude procurement. Often a combination of these actions is needed.

For the emergency procurement *EmPr*, the procurement department can contact a local supplier to buy crude at a much higher price but with a shorter lead time. The procurement department has to ask the logistics department for the expected delay to be able to make this decision. Furthermore, the operations department can choose to change the operational configuration (*COC*), meaning that a different recipe is selected using crudes that are still in stock but which result in yields that are not ideal compared to the scheduled operation. Finally, the operations department can change the operational scale (*COS*) to run the refinery at a lower throughput, producing fewer end products but avoiding having to shut down the plant when crude runs out (or postponing plant shutdown, for example to allow emergency procurement crudes to arrive).

The degrees of freedom are thus defined as follows:

$$\bar{x} = (EmPr_1, EmPr_2, EmPr_3, EmPr_4, EmPr_5, COC, COS) \quad (5.2)$$

For emergency procurement the degree of freedom for each of the five crudes is between 0 kbbl and the amount that could reasonably be available on short notice, which is assumed to be 600 kbbl.³ Furthermore, the number of different recipes in the refinery is assumed to be four, one of which is always selected as the current operational configuration. The CDU in the refinery has a minimum capacity as one of its design parameters; below 40 % of the maximum throughput the process no longer works and the refinery has to be shut down. These constraints are defined as follows:

$$0 \leq EmPr_i \leq 600 \quad \text{in kbbl, for each of the } i = 5 \text{ crudes} \quad (5.3)$$

$$COC \in \{R1, R2, R3, R4\} \quad \text{discrete choice between operating modes} \quad (5.4)$$

$$40 \leq COS \leq 100 \quad \text{percentage of CDU throughput capacity} \quad (5.5)$$

The scope is limited to disturbances \vec{d} dealing with the supply of crude oil to the crude distillation units. These are defined by:

$$\vec{d} = (ShipDelay, StorageProblem) \quad (5.6)$$

with:

$$ShipDelay \in \{0, 1, 2, \dots, n\} \quad \text{in days, for 1 ship for 1 cycle} \quad (5.7)$$

$$StorageProblem \in \{0, 1\}^m \quad \text{for each of the } m = 5 \text{ crude storage tanks} \quad (5.8)$$

For this case study it is assumed that a disturbance to the system occurs on day $t = 22$. A ship at sea is delayed for 30 days, but there are no problems with any of the storage units. This means that the disturbance is defined as $\vec{d}(22) = (30, 0, 0, 0, 0, 0)$.

The Nelder-Mead optimisation method (Nelder and Mead 1965) is used here as the search strategy, to determine *which* experiments need to be performed with the model. The method is based on identification of the best, the worst, and the second worst outcomes in each iteration for the pre-defined simplex (a set of experiments). See the box below for more details. This method has not only been chosen because of its ability to deal with discontinuous objective functions, but also to illustrate how an optimisation method that is commonly used in process systems engineering using mathematical models (Biegler and Grossmann 2004) or samples from experiments in a real system can also be a powerful approach when combined with an agent-based model.

³kbbl stands for 1000 standard oil barrels.

Nelder-Mead optimisation method

In the Nelder-Mead optimisation method (Nelder and Mead 1965), an initial simplex S is defined as a convex hull with $n + 1$ vertices $\{\bar{x}_j\}_{j=1}^{n+1}$ in an n -dimensional space \mathbb{R}^n (with n equal to the number of degrees of freedom in \bar{x}). These vertices satisfy the non-degeneracy condition, meaning that the volume of the simplex hull is non-zero. For every next iteration $j + 1$, the values for $\{\bar{x}_j\}_{j=1}^{n+1}$ are determined by comparing the objective-function values followed by replacement of the worst vertex by another point. The simplex adapts itself to the local landscape and finally contracts to the (local) optimum. In other words, for an initial collection of chosen values for each of the variables one may wish to adjust (i.e. the “degrees of freedom”) it is determined, in this case through simulation, what the performance of the system is. The search algorithm then determines, based on these findings, which variables to adjust and in which direction, through extrapolation of the previous generation. By repeating this process and running new simulations using the values suggested by the algorithm, the variable landscape is explored until a suitable solution is found.

An initial simplex S for a 6-dimensional space (the number of degrees of freedom in \bar{x}) for the Nelder-Mead optimisation method needs to contain 7 vertices. Table 5.3 (rows 1 to 7) on page 175 shows the initial values that were chosen, based on a first analysis of the problem by an expert. These values for the degrees of freedom provide enough variation for the search algorithm to proceed. The stop condition for the search is when all \bar{x} in the population have prevented a shutdown of the refinery and no better values for the criteria are found through a new iteration.

5.8.2 Experimental Setup for the Multi-plant Enterprise

The agent-based model developed for the multi-plant enterprise can be used to set up many experiments and to study important factors that influence the behaviour of each plant separately as well as the performance of the enterprise as a whole. For this purpose we need to define some performance indicators at the enterprise and plant levels. In general, the performance of a supply chain can be analysed in terms of customer service (e.g. tardiness, number of late orders), financial aspects (e.g. profit, overall operational cost), or a combination of both. The performance indicators considered in the experimental setup in this chapter are *number of late orders*, *total tardiness (in days)* and *average inventory level*. It should be stressed that the flexibility of an agent-based model guarantees an easy extension of the original performance indicator if it appears that additional analysis is necessary. The experiments with the developed model can be defined by implementing differ-

ent behaviours for the different departments (i.e. changing their properties or their working policies).

Meanwhile, the objective of the experimental setup can be related to the normal operation of this multi-plant supply chain (i.e. to evaluate the dynamic behaviour of supply networks under different conditions or to study the effect of different policies on system performance in order to find the most appropriate policy). For instance, the effect of changing the reorder point (procurement policy) on system performance is studied. For this purpose, this parameter was changed from 10 to 25 percent, so as to find an appropriate value that minimises the *number of late orders* and *total tardiness*. Similarly, changing the policies for different departments may have many different effects on different system performance indicators, and accordingly, the normal operation of the supply chain under different scenarios can be analysed.

In addition to those for the normal behaviour of a multi-plant enterprise, many experiments can be formulated to study the enterprise performance during different, abnormal, situations and to find effective strategies to cope with them. As an example, we assumed an unexpected shutdown in Production Plant 1 at the 70th day of the time horizon. Consequently, starting with the 71st day, all orders must be fulfilled by Production Plants 2 and 3. As expected, the overall performance was affected by this disruption, resulting in a number of non-finished and late customer orders, and the nominal optimum reorder point was not optimal for the disturbed situation. To handle this disruption, the enterprise could adapt its different policies (e.g. its procurement policy) by changing its nominal reorder point. The effect of this adaptation is also studied with the model.

5.9 Step 8: Data Analysis

Next, the results of the experiments described in Step 7 are presented and analysed for both case studies.

5.9.1 Delay in Shipment in the Oil Refinery Supply Chain

Figure 5.6a illustrates the crude stocks of the refinery under normal operation, and Fig. 5.6b shows the operational scale (both planned and actual throughput) over time after disturbance $\vec{d} = (30, 0, 0, 0, 0, 0)$ which occurs on day $t = 22$. This disturbance resulted in a loss of \$24 million because of lost production hours, as illustrated by the gap between planned and actual throughput. Furthermore, the agent-based simulation model supports the decision about which response is the most appropriate given the many degrees of freedom. The change in operational configuration was not included in this experiment; rather, the planned production mode (in this case $R4$) was used. For each possible $COC \in \{R1, R2, R3, R4\}$ an optimisation for $EmPr_i$ and COS had to be performed because choosing a different recipe would necessarily influence the criterion surface. Preliminary results for the decision on emergency

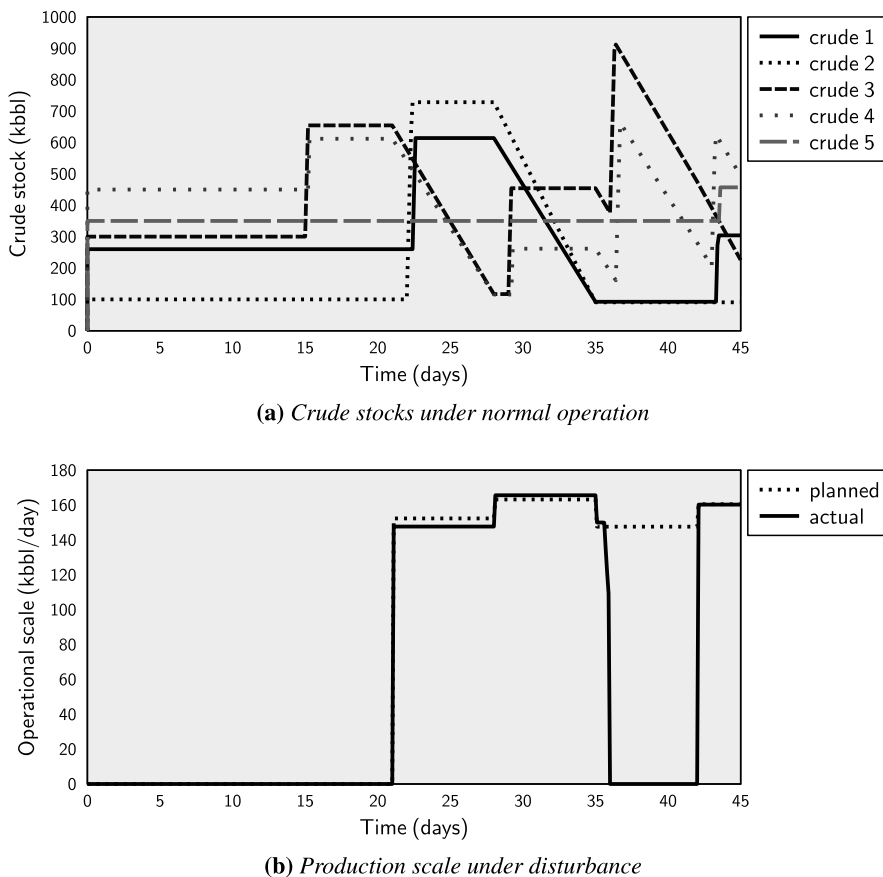


Fig. 5.6 Results from the oil refinery supply chain case study

procurement and the change of operational scale are shown in Table 5.3. After 20 iterations no further improvement was made, so the search was terminated. The proposed solution prevents a shutdown of the refinery by purchasing emergency crudes to make up for the delayed shipment and by slightly reducing the throughput. The loss caused by the disturbance was reduced by \$14.7 million (excluding penalties to be paid by the shipper for delays) which is more than 60% of the \$24 million in damage caused by the disturbance (van Dam et al. 2009b).

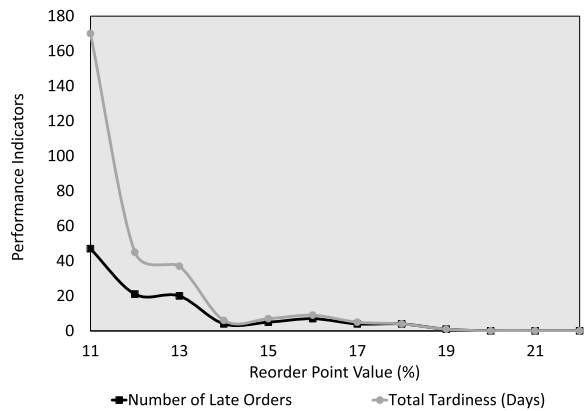
5.9.2 Normal and Abnormal Behaviour Analysis for the Multi-plant Enterprise

Figure 5.7 shows the effect of changing the reorder point (procurement policy) on plant-level and enterprise-level performance. Accordingly, as the reorder value in-

Table 5.3 Outcome of the Nelder-Mead simplex optimisation algorithm with the agent-based model

	$EmPr_1$	$EmPr_2$	$EmPr_3$	$EmPr_4$	$EmPr_5$	COS	$P(\bar{x}, \bar{d})$
1	0	0	0	0	0	60	-5.69 E8
2	0	0	100	100	0	60	-3.65 E8
3	0	0	200	300	0	60	-1.61 E8
4	0	0	300	250	0	55	-1.64 E8
5	0	0	300	250	0	50	-6.19 E7
6	100	100	100	100	100	50	-3.70 E8
7	10	10	500	600	10	50	1.39 E8
final	4	4	349	382	4	54	1.45 E8

Fig. 5.7 The effect of the procurement policy (reorder point) on overall system performance (total tardiness and number of late orders)



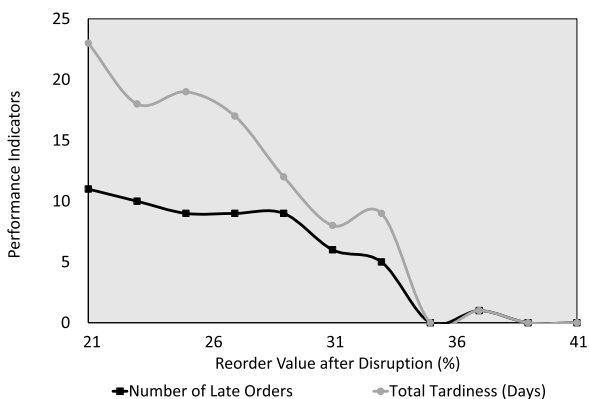
creases, the overall performance of the enterprise improves (*total tardiness* and *number of late orders* decrease); since without raw material, the operation department must pause its execution of an order and wait to receive the raw materials from the supplier. This causes delay in fulfilling customer orders that are assigned to that production plant. Generally speaking, with a higher reorder point the raw material availability can be a less important bottleneck for the production plant. Figure 5.7 also suggests considering 20 percent as an optimum reorder point for this set of customer orders. This is because 20 percent is the first reorder value at which both *total tardiness* and *number of late orders* are zero; all customer orders can thus be fulfilled without any delay.

Table 5.4 illustrates how an unexpected shutdown in Production Plant 1 at the 70th day of the time horizon will affect the performance of multi-plant enterprise. The assumption was that all production plants were working according to their optimum reorder point (20 percent). Because of this plant disruption, there are 14 late orders with 34 total tardy days. As mentioned in the previous section, one way of managing this abnormality is to adapt the policies for different departments. One possibility would be to change the procurement policy for Plants 2 and 3 immedi-

Table 5.4 Performance data for the enterprise in a normal and an abnormal situation (Plant 1 shutdown); the reorder point in both cases is 20 percent

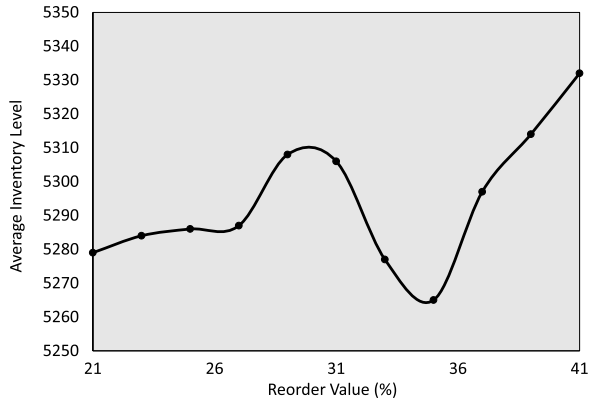
Performance indicator	Normal operation	Abnormal situation
Orders assigned to ProductionPlant1	33	22
Orders assigned to ProductionPlant2	33	39
Orders assigned to ProductionPlant3	34	39
Orders assigned to All Production Plants	100	100
Late Orders by ProductionPlant1	0	0
Late Orders by ProductionPlant2	0	7
Late Orders by ProductionPlant3	0	7
Late Orders by All Production Plants	0	14
Total tardiness for ProductionPlant1	0	0
Total tardiness for ProductionPlant2	0	16
Total tardiness for ProductionPlant3	0	18
Total tardiness for All Production Plants	0	34

Fig. 5.8 The effect of the procurement policy (reorder point) on overall system performance (total tardiness and number of late orders) during an abnormal situation (Plant 1 shutdown)



ately following the disruption in Plant 1. The relation between changing the reorder point for Production Plants 2 and 3 after disruption and the total tardiness as well as the number of late orders is shown in Fig. 5.8. In general, increasing the reorder point improves these two performance indicators; but on the other hand, a higher reorder point will increase *raw material holding cost*, too (Fig. 5.9). In total, if the reorder point is increased, the *average inventory level* must also be increased; but the inventory level is also dependent on the number of completed orders. Having more finished orders means more raw material consumption and a lower average inventory level. Based on this analysis, setting the reorder point for the two available production plants to a value of approximately 35 percent will result in desirable levels for both *customer satisfaction* and *raw material holding cost*. Of course, for a real case, according to the storage-related costs, a more precise value can be determined. Meanwhile, an enterprise can react to an abnormal situation with many

Fig. 5.9 The effect of the procurement policy (reorder point) on overall system performance (average inventory level) during an abnormal situation (Plant 1 shutdown)



other strategies, and changing the reorder point can be a part of a comprehensive crisis management package.

5.10 Step 9: Model Validation

Following the views on validation presented in Sect. 3.10 in Part I, the main concern is how to be certain that the model is applicable for its domain and can aid the understanding of the problem under study. This application relevance can be shown by different experiments done with the model and the insights gained regarding, for example, the dynamic behaviour of the system. The previous subsections in this chapter show the application of both developed models for the problems defined in Sects. 5.2.1 and 5.3.1 by defining different experimental setups. Furthermore, when it comes to expert validation, for both models presented in this chapter, the domain expert validation was done in all stages of model development. As mentioned in Sect. 5.7, the refinery supply chain model was benchmarked against another model. That equation-based model has been validated against the real system and was applied in order to offer decision support (Pitty et al. 2008). Since the agent-based model of the refinery supply chain was successfully compared with an already validated mathematical model, the agent-based model that was based on it can be considered valid also. Similarly, feedback from domain experts on the multi-plant case gives confidence that the model is valid and useful.

5.11 Step 10: Model Use

The models presented in this chapter have been used for *offline* decision support, meaning that they are used to test different scenarios in a simulation, after which the “best” solution can be implemented in the real system. In an *online* controller the results of the model could be directly used to affect the real system. However, it

should be stressed that the same model of a supply chain could be used in different ways and to solve different problems. Perhaps some changes would have to be made to be able to include other types of disturbances and add new degrees of freedom to test out new scenarios, but the basic foundation of the model will remain useful and applicable.

The *actors* represented in the models do not include the role of the supply chain manager itself, even though that role was chosen as the problem owner and the main actors to benefit from the decision support offered by these models. This means that the *agents* in the model, representing the actors, do not themselves take decisions about how to deal with a witnessed disturbance, but it is up to the *user* to decide which experiments to run and how to interpret the outcomes.

Various roles of simulation in supply chain management can be considered, including (Merkuryev et al. 2009):

- Back-up decision-making
- Validate algorithms
- Convince users of supply chain management approaches
- Educate

Models of supply chains have been successful in offering decision support. As an example, consider early supply chain models which could be used to visualise the so called *bull whip effect* in which the effects of decisions are propagated through the supply chain (Forrester 1958). Furthermore, different policies and algorithms have been tested with such models (Julka et al. 2002b). Simulations of supply chains have also been used for educational purposes to elucidate the dynamics within a chain of different entities (Holweg and Bicheno 2002).

One specific issue with the visualisation of data and simulation outcomes of supply chain models is making sure that justice is done to the wide range of actors or geographical locations. As stated above, supply chains are complex systems with multiple social entities intertwined with many physical systems. Solving a problem in one aspect of the supply chain (e.g. a red alert for a storage tank) might require actions elsewhere in the supply chain. As of yet this is an unsolved challenge in the two models presented in this chapter. This would, however, be an important step towards the development of decision support tools which can be used by different stakeholders, even if they are not modellers themselves.

5.12 Conclusions

This chapter has described the application of agent-based modelling as a promising approach to capturing the dynamic and complex behaviour of an industrial supply chain. Meanwhile, the two cases studies presented in this chapter have shown that models developed using the modelling steps from Chap. 3 can be used to support decision-makers for managing their supply chains, both during normal operation and when facing abnormalities.

Each supply chain is a decentralised, complex and adaptive system with many heterogeneous actors and physical components interacting through different flows, including material, information, monetary and social flows. Furthermore, decision-making in a supply chain is distributed among different actors, and each of these actors has its own objectives and procedure for decision-making. The collective decisions made by these autonomous actors at various levels of the system result in the overall system behaviour. So the need for a modelling approach that can capture all these complexities and this decision structure is clear. As discussed in this chapter, agent-based modelling seems to be one of the best candidates for addressing these issues. In addition, agent-based modelling is flexible and can define a broad range of experiments with different scenarios to answer “what if” questions; this is critical for decision support under disruptions or in the design phase. The models used in this chapter were developed in a bottom-up fashion, making it relatively simple to change the configuration: it is easy to include new actors in the system (e.g. more suppliers with different prices and lead times in the supply chain) or to adjust the physical configuration (e.g. extra storage tanks for the refinery or production plants).

However, there are some barriers that may affect the willingness to use this modelling approach: one of the main challenges concerns validation, and this is one of the main sources for much of the criticism that agent-based modelling has received as a research method in the literature. Obviously, these criticisms are not solely about using agent-based modelling for supply chain management. In addition, considering the novelty of this approach (especially as a decision support methods for complex adaptive systems), there is a need for more research on finding some universally accepted validation methods.

Acknowledgements The authors would like to express their thanks to Rajagopalan Srinivasan and Arief Adhitya for their support in the development of the models used in both case studies.

References

- Behdani, B., Lukszo, Z., Adhitya, A., & Srinivasan, R. (2010a). Agent-based modelling to support operations management in a multi-plant enterprise. *International Journal of Innovative Computing, Information and Control*, 6(7), 1–12.
- Behdani, B., Lukszo, Z., Adhitya, A., & Srinivasan, R. (2010b). Performance analysis of a multi-plant speciality chemical manufacturing enterprise using an agent-based model. *Computers & Chemical Engineering*, 34(5), 793–801.
- Bhatnagar, R., Chandra, P., & Goyal, S. K. (1993). Models for multi-plant coordination. *European Journal of Operational Research*, 67(2), 141–160.
- Biegler, L., & Grossmann, I. (2004). Retrospective on optimization. *Journal of Computers and Chemical Engineering*, 28(8), 1169–1192.
- Forrester, J. (1958). Industrial dynamics: a major breakthrough for decision makers. *Harvard Business Review*, 36(4), 37–66.
- Gjerdrum, J., Shah, N., & Papageorgiou, L. G. (2000). A combined optimisation and agent-based approach for supply chain modelling and performance assessment. *Production Planning & Control*, 12, 81–88.

- Holweg, M., & Bicheno, J. (2002). Supply chain simulation—a tool for education, enhancement and endeavour. *International Journal of Production Economics*, 78(2), 163–175.
- Julka, N., Srinivasan, R., & Karimi, I. (2002a). Agent-based supply chain management—1: framework. *Computers & Chemical Engineering*, 26(12), 1755–1769.
- Julka, N., Karimi, I., & Srinivasan, R. (2002b). Agent-based supply chain management—2: a refinery application. *Computers & Chemical Engineering*, 26(12), 1771–1781.
- Kaihara, T. (2003). Multi-agent based supply chain modelling with dynamic environment. *International Journal of Production Economics*, 85(2), 263–269.
- Lee-Post, A., & Chung, C. H. (2008). *Systems for supporting operations management decisions*. Berlin: Springer.
- Mele, F. D., Guillén, G., Espuña, A., & Puigjaner, L. (2007). An agent-based approach for supply chain retrofitting under uncertainty. *Computers & Chemical Engineering*, 31(5–6), 722–735.
- Merkuryev, Y., Merkurjeva, G., Piera, M., & Guasch, A. (2009). *Simulation-based case studies in logistics*. Berlin: Springer. ISBN: 978-1-84882-186-6.
- Min, H., & Zhou, G. (2010). Supply chain modeling: past, present, and future. *Computers and Industrial Engineering*, 43(1–2), 231–249.
- Moyaux, T., Chaib-draa, B., & D'Amours, S. (2006). *Supply chain management and multiagent systems: an overview*. Berlin: Springer.
- Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, 7(4), 308–313.
- Pitty, S. S., Li, W., Adhitya, A., Srinivasan, R., & Karimi, I. (2008). Decision support for integrated refinery supply chains—1. Dynamic simulation. *Computers & Chemical Engineering*, 32(11), 2767–2786.
- Siirola, J. D., Hauan, S., & Westerberg, A. W. (2003). Toward agent-based process systems engineering: proposed framework and application to non-convex optimization. *Computers & Chemical Engineering*, 27(12), 1801–1811.
- Srinivasan, R., Bansal, M., & Karimi, I. (2006). *A multi-agent approach to supply chain management in the chemical industry*. Berlin: Springer.
- Swaminathan, J. M., Smith, S. F., & Sadeh, N. M. (1998). Modeling supply chain dynamics: a multi-agent approach. *Decision Sciences*, 29(3), 607–632.
- Tayur, S., Ganeshan, R., & Magazine, M. (1999). *Quantitative models for supply chain management*. Norwell: Kluwer Academic.
- van Dam, K. H. (2009). *Capturing socio-technical systems with agent-based modelling*. PhD thesis, Delft University of Technology, Delft, the Netherlands. ISBN 978-90-79787-12-8.
- van Dam, K. H., Adhitya, A., Srinivasan, R., & Lukszo, Z. (2008). Benchmarking numerical and agent-based models of an oil refinery supply chain. In B. Braunschweig & X. Joulia (Eds.), *Computer-aided chemical engineering: Vol. 25. Proceedings of the European symposium on computer aided process engineering ESCAPE 18* (pp. 623–628). Lyon: Elsevier.
- van Dam, K. H., Adhitya, A., Srinivasan, R., & Lukszo, Z. (2009a). Critical evaluation of paradigms for modelling integrated supply chains. *Journal of Computers and Chemical Engineering*, 33(10), 1711–1726. 2009.01.023. doi:10.1016/j.compchemeng.
- van Dam, K. H., Lukszo, Z., & Srinivasan, R. (2009b). Abnormal situation management in a refinery supply chain supported by an agent-based simulation model. In R. M. de Brito Alves, C. A. O. do Nascimento, & E. C. Biscaia Jr. (Eds.), *Proceedings of the 10th international symposium on process systems engineering—PSE2009*, Salvador, Bahia, Brazil.
- Zhang, H., Wong, W. K., Adhitya, A., & Srinivasan, R. (2008). Agent-based simulation of a speciality chemicals supply chain. In *Proceedings of the first international conference on infrastructure systems (INFRA 2008): building networks for a brighter future*, Rotterdam, The Netherlands.