CIRRELT

Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Agent-Based Simulation to Anticipate Impacts of Tactical Supply Chain Decision-Making in the Lumber Industry

Sébastien Lemieux
Sophie D'Amours
Jonathan Gaudreault
Jean-Marc Frayret

November 2009

CIRRELT-2009-51

Agent-Based Simulation to Anticipate Impacts of Tactical Supply Chain Decision-Making in the Lumber Industry

Sébastien Lemieux[1,2,*], Sophie D'Amours[1,2], Jonathan Gaudreault[1,2], Jean-Marc Frayret[1,3]

[1] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

[2] Département de génie mécanique, Université Laval, Pavillon Adrien-Pouliot, 1065, avenue de la médecine, Québec, Canada G1V 0A6

[3] Mathematics and Industrial Engineering Department, École Polytechnique de Montréal, C.P. 6079, succursale Centre-ville, Montréal, Canada H3C 3A7

Abstract. This paper proposes the integration of an Advanced Planning and Scheduling (i.e., APS) system in a multi-agent based simulation for the lumber industry. The goals are to simulate different market situations and evaluate the impacts of different supply chain planning tactics. In order to do this, we developed a simulator consisting of (1) a software agent emulating various customer types, connected to (2) a distributed APS made up of interdependent software agents that carry out operations planning of the different business units of the lumber supply chain and (3) simulator components such as time management and performance indicators used by the decision maker. A case study (Virtual Lumber Case) was developed in order to illustrate the potential of the simulator which represents an average Canadian enterprise in the softwood lumber industry. Two experiments are proposed to demonstrate the possibilities of this simulator. These experiments show that when planning using final customer demands, it is possible to increase the volume of non-discounted sales by up to 40% while improving the service level offered to the final customer.

Keywords. Distributed simulation, value creation network, advanced planning system, multi-agent, lumber.

Acknowledgements. This work was funded by the FORAC Research Consortium and the Natural Sciences and Engineering Research Council of Canada (NSERC). The authors would also like to acknowledge the impressive work of many FORAC research assistants and students who have worked on developing the FEPP over the years.

---

* Corresponding author: Sebastien.Lemieux@cirrelt.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec,
Bibliothèque et Archives Canada, 2009

1 Introduction

The Canadian lumber industry has been facing difficult challenges for the last few years as demand and prices for lumber products are decreasing in the North American market, mainly due to the financial crisis and the credit bubble which had caused the number of housing starts to decline. Faced with these challenges, lumber producers need to increase the operational efficiency of their supply chain in order to improve their service levels and increase value creation. Improving the design of these value creation networks and their planning and coordination strategies are possible means to attain these goals.

Consequently, there is a need for tools aiming to anticipate the economical impacts of these new designs. Adapted Advanced Planning and Scheduling (APS) systems can be part of these tools if coupled with simulators. APS systems provide companies with algorithms and models for planning their different activities from procurement to distribution (Stadtler 2005). These tools are well developed for discrete manufacturing and assembly types of production but less so for processes -based manufacturing and divergent production. This is especially true in the forest industry as reported by Rönnqvist (2003) and D'Amours et al. (2008), and the need for specific solutions is well acknowledged (Gronalt & Rauch 2008).

In previous work, an adapted distributed APS (dAPS) system used to plan and synchronize operations within a lumber supply chain was presented (Frayret et al. 2007). This FORAC Experimental Planning Platform (FEPP) allows the use of planning algorithms specific to each type of activity in the network as well as the use of coordination mechanisms to synchronize the different plans (Forget et al. 2006). In order to improve the design of the value creation networks, Cid Yáñez et al. (2009) exploited the FEPP in order to assess the impact of introducing demand information at various points in the planning process. These first contributions clearly demonstrated the value of demand information in terms of inventory decrease and increase of customer satisfaction.

Using simulation for such a decision-making context requires the ability to develop an operating simulation framework around the FEPP and the capability to verify the validity of the experiments carried out. Given the size of the problem to be addressed, these steps were particularly important in this study. It is indeed necessary to assess the validity of simulations (Balci 1986; Sargent 2005) and this validation should be done throughout the development of the simulations as proposed by Robinson (2002).

In this paper, a simulator which aims at evaluating the performance of a supply chain under different planning strategies is introduced. Building on the framework for distributed decision making proposed by Schneeweiss (2003), this work is a step towards the development of a flexible simulator consisting of APS tools and simulated customer agents to model a supply chain and its customers as introduced by Petrovic et al. (1998) and Chang & Makatsoris (2001). Figure 1 demonstrates how the simulator is composed of these modules (the FEPP, simulation components and a simulated customer agent) to simulate the scenarios and how it interacts with the user.

The FEPP serves to simulate the decision planning process of the business units involved in the supply chain. The *Simulated Customer Agent* is used to simulate the ordering behaviour of the customers. The simulation components provide the simulating environment and permit an accelerated simulation of the supply chain. Finally, the user controls the simulations and uses different KPI to evaluate the performance of the different planning strategies.

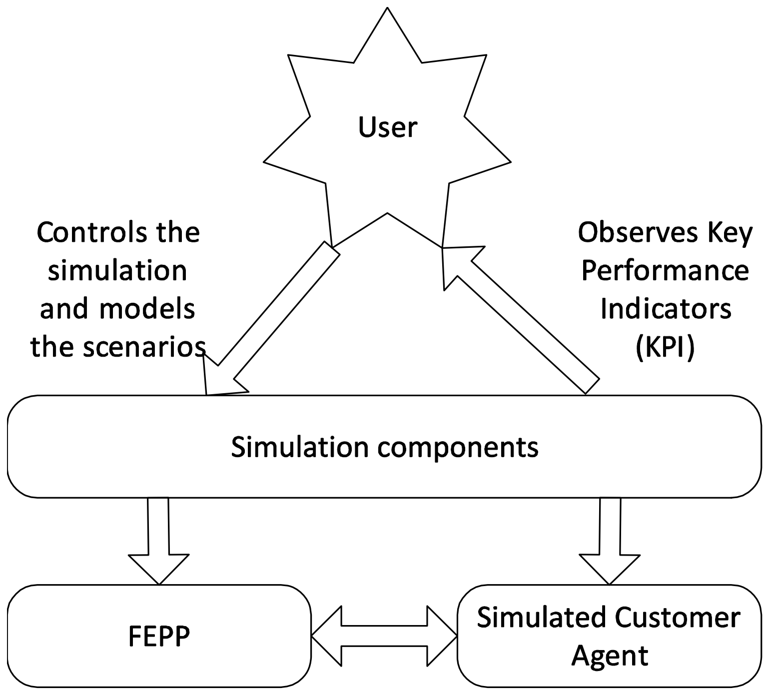

Figure 1 Conceptual simulator composition

In order to illustrate the contribution of this paper, an enterprise test bench representing a typical Canadian lumber production network called Virtual Lumber Case (VLC) is used. It is explained in Section 2 and is derived from actual data from the Canadian industry. The main advantage of using a virtual case is that it is possible to discuss all information regarding it without having to worry about disclosing trade secrets. The case information is available on request to the authors.

The proposed simulator is then discussed in Section 3. It consists of three main modules: the FEPP mimicking the behaviour of the planning process, different components such as time management and performance indicators used by the decision maker to compare different scenarios, and finally, a simulated customer agent simulating the behaviour of the final customers.

In Section 4, this simulator is used to assess the planning strategy of a lumber supply chain and more specifically, the impact of non-discounted sales volume using demand-driven approach compared to the more traditional supply-driven approach. The results are presented in Section 5. The last part (Section 6) will conclude with the different contributions of this paper and the future paths to explore.

2 Virtual Lumber Case

The Virtual Lumber Case is an enterprise test bench, which is basically a detailed supply chain model of a typical Canadian enterprise in the softwood lumber industry. It is a work in progress that includes data describing manufacturing processes, products and transportation resources that are based on different existing eastern Canadian sawmills. These data were used to configure a particular instantiation of the FEPP (D'Amours et al. 2006; Frayret et al. 2007). This platform is a dAPS system, in which each agent contains a detailed model of the organizational unit of the supply chain it represents. Furthermore, each agent is able to use specialized algorithms to plan operations and uses coordination mechanisms to synchronize its plans with other agents.

2.1 FORAC Experimental Planning Platform

The FEPP aims to create feasible and coordinated plans which are optimized to fulfil the demand of the final customers of the value creation network. These plans consist of detailed production schedules and take into account multiple constraints that are specific to each organizational unit. The FEPP is based on a multi-agent approach to coordinate the collective definition of each agent's production schedule, which uses coordination protocols (called conversations) to specify how agents should interact in various situations.

2.1.1 Architecture Overview

The FEPP is composed of agents that interact with each other in order to solve the global lumber supply chain planning problem. Several conversation protocols and agent behaviours have been implemented in order to produce optimized solutions in various situations.

The configuration of the platform follows an organizational design approach that consists in the division of the supply chain into business units. In turn, this division into business units is based on the natural heterogeneity of the production process. In other words, the overall problem is split into several smaller sub-problems, each of which results in the managerial problem of a single organizational unit. Consequently, every agent is modelled after a specific organizational problem. This gives each agent the ability to solve a smaller scale problem using adapted tools. The collaboration, or the cooperation, of all agents present in the supply chain solves the value creation network global problem.

Many functionalities are shared by all agents. For example, they can interact with each other or perform tasks when necessary and by definition, they are autonomous. Moreover, by the way in which the FEPP is conceived, each agent can be represented as a modular block that can be assembled with others to model a supply chain.

Four basic types of agents were developed as seen in Figure 2: *Planning Unit*, *Source*, *Make* and *Deliver*. The *Planning Unit agent* leads agents in the global environment; it helps internal agents communicate with external agents, those that are in a different planning unit. A planning unit usually corresponds to a physical site in a company.

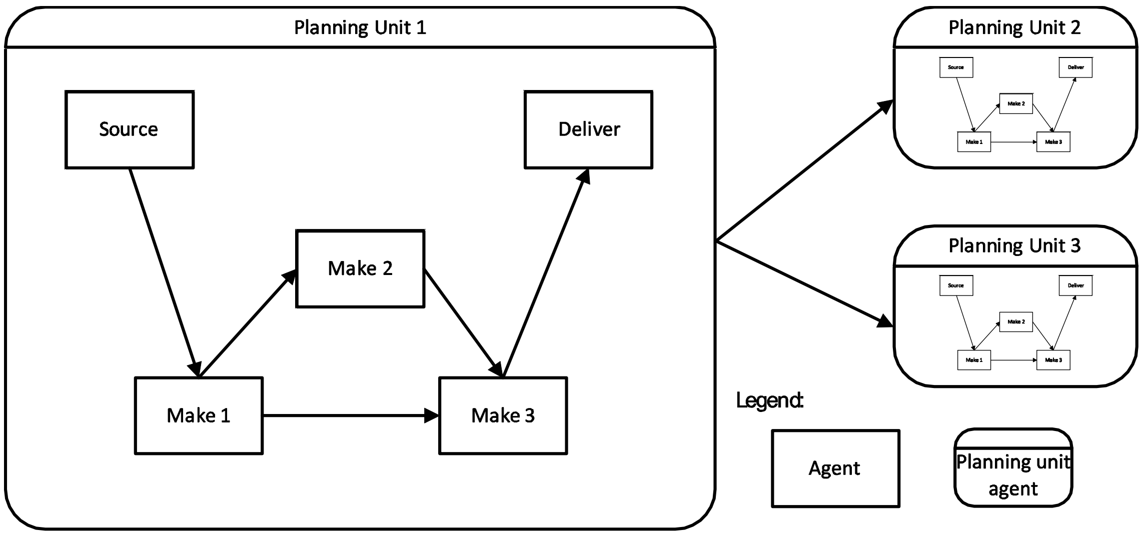

Figure 2 Generic overview of the FEPP

The *Source agent* plans the procurement of the products or raw materials needed by other agents in the planning unit and allocates them to the desired location. The *Deliver agent*, on the opposite side of the planning unit in Figure 2, is responsible for managing relationships with customers and selling manufactured goods. It optimizes the flows of material between an external customer and the supplier inside the planning unit. Finally, *make agents* are responsible for planning the operations of the production or distribution of products inside the planning unit. They are specifically designed to fit the production problem of each organizational unit. The most common types in lumber supply chains are the *Sawing*, *Drying*, *Finishing* and *Warehouse agent*. Each one uses operational research algorithms to plan their activities. The overall organization of the information exchanges between agents and the planning process is described in Section 2.2.1.

2.1.2 Workflow

In order to achieve their goals, these agents need to execute certain tasks, including planning local operations, computing local performance indicators or updating planned inventory data. In the FEPP, these tasks are enclosed within workflows. Basically, a workflow is a series of tasks and business rules.

From an implementation point of view, the agents tested in this study are reactive. Indeed, they listen to events that occur in a simulation. An event represents a change in the environment of an agent (e.g., an agent receives a new demand for a product, the current operational plan has changed, etc.). When an event is triggered, the agent launches all the workflows registered to this event. In turn, events can be triggered by workflows, database operations or conversations between agents. Events can also be parametrized in order to enable the tasks of the workflows to be properly carried out. The association of events to workflow is configurable by users in order to change the behaviour of agents.

An example of a workflow is shown in Figure 3. First, the agent executing this workflow loads a released plan containing the scheduled frozen jobs of the current active plan. If jobs need to be rescheduled because there is not enough inventory to start a job at the specified date, the agent executes the next task to reschedule the jobs when inventory becomes sufficient. Otherwise, the agent will continue and validate the current manufacturing plan template to ensure that each job can be executed with the specified resources. Next, the agent starts its planning engine to create a new operational plan. This plan is then inserted as a manufacturing plan template into the database. The last step is setting this plan as active to let the agent know that the current production schedule is now obsolete.

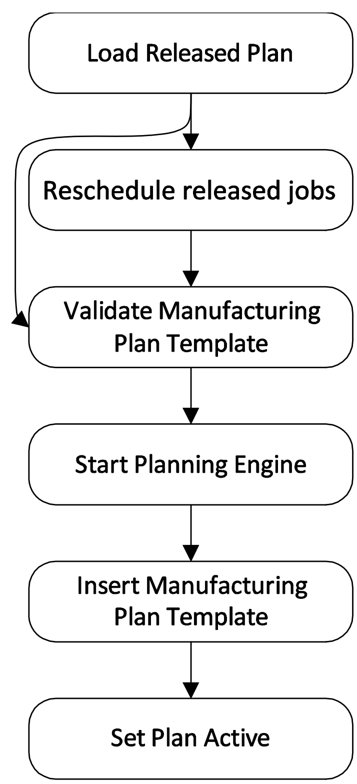

Figure 3 Example of a workflow

2.1.3 Conversation

In order to coordinate their plans, agents need to communicate with each other. Agents use conversation protocols that specify how to behave during an exchange of messages. These protocols model the role (i.e., sender or receiver) of each agent taking part on the corresponding exchange of messages. Each conversation protocol is composed of two symmetrical protocols (one for each agent). They can be described as a series of states between which particular messages are exchanged. During a given state, the active agent needs to do certain tasks which may in turn lead to a message to be sent. Once a message is sent, the sender agent is generally in a passive state, listening to new events or expecting an answer message from the other agent.

During a state, as proposed by FIPA[1], a message can be sent and must contain the following information:

- Sender agent name;
- Receiver agent name;
- Conversation protocol name;
- Conversation protocol version;
- Content of the message;
- Synchronous or Asynchronous sending.

The last information specifies to the sender if an answer from the receiver is expected. In such a case, the message will be sent synchronously. During the waiting period, the sender agent can execute a different task in parallel. Otherwise, the sender sends the message and then waits for the next task.

Conversations are triggered inside specific tasks. They can trigger an event in the receiver agent that can in turn start a workflow. An example of a simple conversation can be found in Figure 4. This protocol is used by an agent to “ask” another agent if it has a specific product in inventory. In the first state, the sender sends a message to another agent that has the desired products. Depending on the availability of the requested product, the two agents will fall into one of the two following states: “InventoryAnalysisDataRetrieved” if the receiver has these products in inventory or “InventoryAnalysisDataFailure” if not.

[1] http://www.fipa.org/

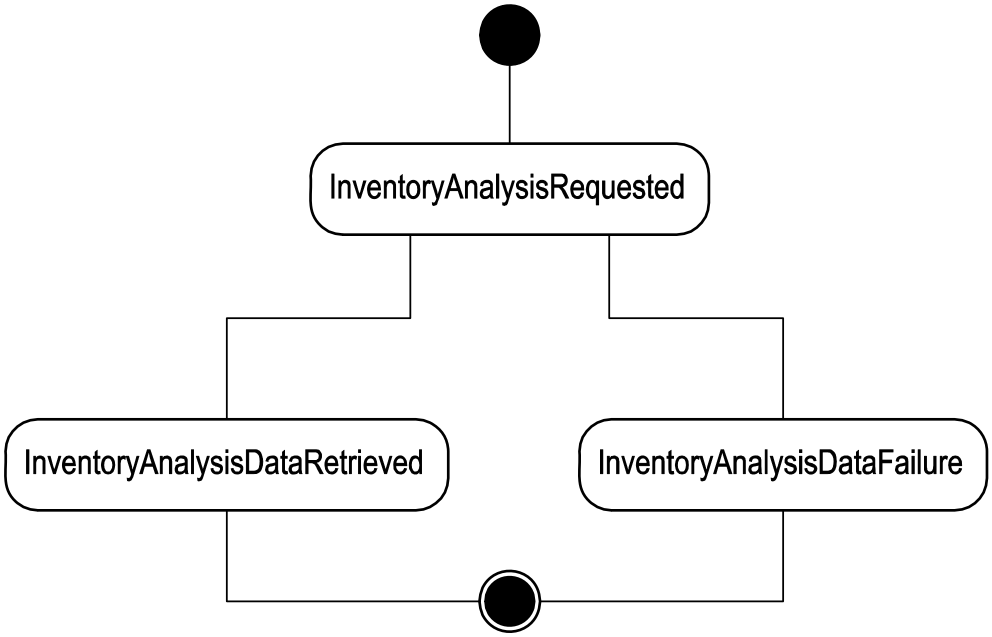

Figure 4 Example of states during a conversation between two agents

2.1.4 Implementation

The FEPP was developed over a seven-year period using mostly Microsoft .NET C# programming language with a service-oriented architecture (SOA) approach. In order to maximize efficiency and interoperability, some components have been also programmed using C++. Those components are related to the operations planning for the *Sawing*, *Drying*, *Finishing* and *Deliver agent* and use operational research tools such as ILOG CPLEX version 9.0.

The FEPP is composed of 99 conversation protocols to coordinate agents when communicating with each other. Depending on the agent type, agents can use between 5 to 21 workflows to define every possible action they can perform.

For every agent used in a simulation, a unique database is used to store all information created during a simulation. No information is deleted during a simulation. Instead each piece of information is tagged with a property that states whether the current information is active or not. So it is possible to “read” or compute the state (i.e. products inventories, planned operations schedule, performance indicator, etc.) of every agent at any moment during a simulation.

2.2 The Value Creation Network

Lumber supply chains differ from traditional supply chains due to the fact that raw material is heterogeneous. Furthermore, the procurement of eastern Canadian mills being mainly from natural forest through timber licences with the government results in mills having little control over the timber quality they are supplied with. Consequently, the output of a single log transformation process is rather difficult to predict. The mills can use process simulation tools such as Optitek developed by FPInnovations (Goulet 2006) but this cannot be applied in realtime operations. Eastern Canadian companies deal with this problem by forecasting the production of a large quantities of logs in order to know what quantities of each product can be sold, which is statistically more stable than the production of a single log.

2.2.1 Production and Distribution Planning Unit

The production and distribution network that was modelled in the FEPP is shown in Figure 5. Operations are planned from log reception to final delivery, using all agents from *Source agent* to *Deliver agent* and eventually to *Simulated Customers agent*.

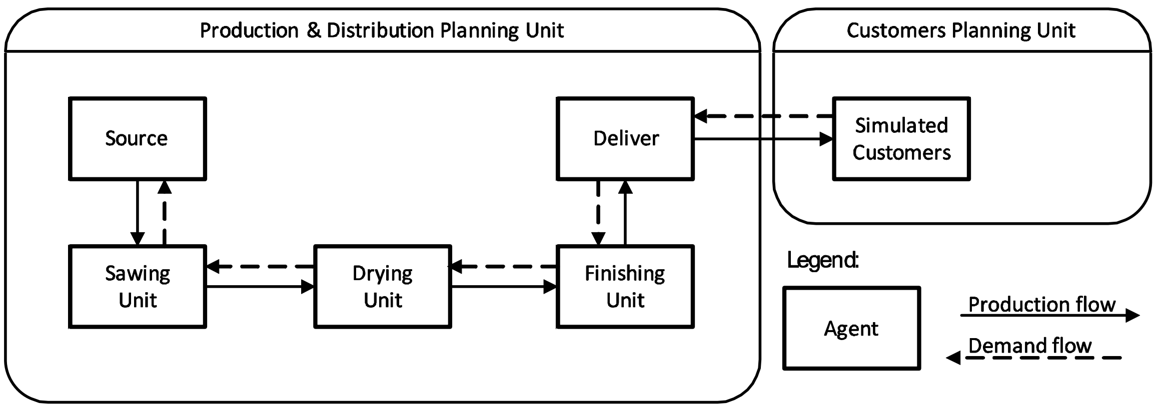

Figure 5 A configuration of the FEPP

As explained earlier, the role of the *Source agent* is to procure logs for the supply chain. These logs are then virtually sent to the *Sawing agent*. By deciding on the volumes and classes of logs to transform during each production shift and the sawing processes to implement with them, logs are planned to be transformed into multiple products: sawdust, chips and lumbers. Bundles of lumbers are then placed into a kiln according to a specific loading pattern. The scheduling of the finishing processes to apply to the dried products is then set by the *Finishing agent*, who determines the final characteristics of the products. The final products are then ready to be sold or delivered to different customers by the *Deliver agent*.

Table 1 summarizes the characteristic of the Virtual Lumber Case from the perspective of each organizational unit: products, processors and processes. The sawing unit can process eight different types of logs. It is composed of two sawing lines. For each type of log entering a sawing line, several cutting patterns are possible. The choice of the cutting pattern determines the mix of the output products.

For the drying unit, fifteen products can be dried in one of the seven kiln dryers depending on which loading pattern is selected from the 180 available. This unit is the bottleneck of the network.

Finally, the finishing unit can transform fifteen input products into 45 final products. For each input product a set of alternative processes is available. The algorithm selects the best sequences of processes respecting demands, inventories and resources constraints. The resulting products are transferred to the deliver unit, which is responsible for planning the deliveries to the customer.

	Sawing	Drying	Finishing	Deliver
Number of consumed products	8	15	15	45
Number of products made	15	15	45	45
Processors	1 Sawing Line 8' 1 Sawing Line 8'-16'	5 small kiln dryers 2 big kiln dryers Air dry area	1 Finishing Line	Not applicable
Processes	37	180	15	Not applicable

Table 1 Definition of sawing, drying, finishing and deliver units

The agents manage information about the products, inventory levels, processors and processes in each unit of the network. In addition, the *Sawing*, *Drying* and *Finishing agents* are able to plan their operations by using respectively, mixed integer programming, constraints programming and heuristics. These models are described in Gaudreault et al. (2009) and were validated weekly using a real industrial case with a sawmill during a one-year period. It is important to note that the drying and finishing models cannot create an optimal solution in a reasonable time. To compensate for this obvious lack, the agent generates a large quantity of feasible solutions, limited by computation time and/or solutions quantity, and picks the best one. Table 2 summarizes the planning problem of each agent.

	Sawing agent	*Drying agent*	*Finishing agent*
Objectives	Min. tardiness Max. production value Min. costs	Min. tardiness Max. production value	Min. tardiness Max. production value Max. resource utilization rate
Processes characteristics	Divergent product flow Co-production Alternative processes Only compatible processes can be executed within the same production shift	Divergent product flow Co-production Alternative processes	Divergent product flow Co-production Alternative processes Only compatible processes can be executed within the same production shift
Parameters	Machines capacity calendar Frozen jobs Maximum sales per product Inventory cost Raw product cost	Machines capacity calendar Frozen jobs Operations costs	Machines capacity calendar Frozen jobs Solution tree exploration mode Min. production duration by family
Optimization method	MIP	Constraint programming	Heuristic

Table 2 Definition of sawing, drying and finishing algorithms

2.2.2 Processes

In general, a transformation process defines how an operation is performed and what input is required for this particular operation. For a car assembly supply chain, the inputs of the process could be different parts, or platforms of pre-assembled parts and the output would be the ready-to-ship car. In contrast, the lumber industry is characterized by divergent processes where the number of output products is greater than the number of input products. When a log is cut, the output of this operation is sawdust, chips and amounts of lumber of different sizes, which can be desirable in the marketplace or not.

Another particularity is that each log or tree can be cut following a set of different cutting patterns. Each association of logs and cutting patterns becomes a distinct process. In our model, the sawing processes are defined by average of inputs and expected quantities of outputs. Also, the production of these materials uses a machine during a finite time; therefore, the process is also characterized by an average time-based usage of a processor type. Table 3 presents a typical sawing process. The consumption is the input product, in this case 8' spruce log. For each 3.476 m^3 (i.e., a group of logs), a set of output products measured in fbm is produced[2] when this specific process (log class + cutting pattern) is performed.

[2] 1 fbm (foot, board measure) is equal to a board 12" wide x 12" long x 1" thick

Consumption:

Product	Quantity	Unit
8' spruce log	3.476	m^3

Production:

Product	Quantity	Unit
Chip	1.03	t
Sawdust	0.06	m^3
Plank	59.59	fbm
2x3 8'	381.5	fbm
2x4 8'	35.99	fbm

Resource utilization:

Product	Duration	Unit
8' Sawing Line	104.88	s

Table 3 Example of a sawing process

In fact, there are many possible cutting patterns to choose from to process a log. This leads to the need to select which cutting pattern to use. Figure 6 illustrates how each log can be cut using a different cutting pattern. But not every cutting pattern may be used with every log; some may not be feasible or may not generate enough value to compensate for the cost of the log. Cutting pattern selection changes over time as it is governed by inventory level, demand profile and resource availability.

At the drying unit, each drying process corresponds to a loading pattern. A loading pattern defines how green lumber is loaded into the dry kiln in order to maximize the use of available space. A loading pattern is created by filling out the volume of a dry kiln with bundles of lumber (see Figure 7). However, there are constraints that apply to these processes. First, the duration of the drying operation is defined by the thickness of the boards to dry and their moisture content. Moreover, different species of wood require different drying times. Therefore, it is optimal to have a homogeneous “mix” of lumber in the kiln. The drying process defines the control parameters over time (e.g. temperature, pressure). Table 4 describes an example of the drying process.

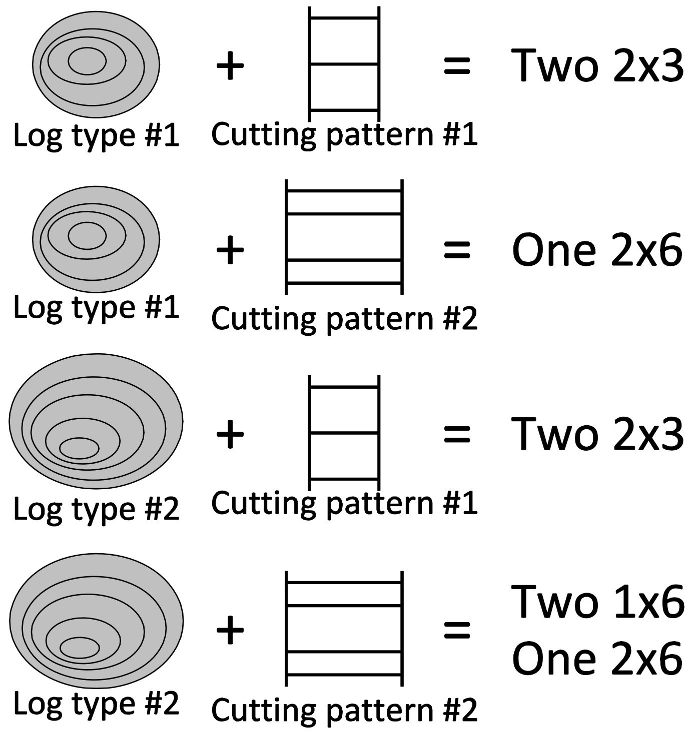

Figure 6 Example of a sawing decision

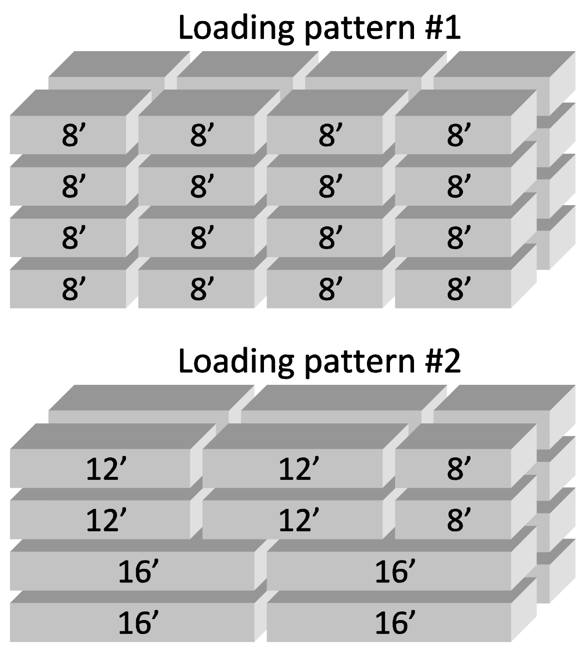

Figure 7 Example of possible loading patterns

Consumption:

Product	Quantity	Unit
2x3 8' Green	31,680	fbm
2x3 10' Green	79,200	fbm
2x3 12' Green	118,800	fbm

Production:

Product	Quantity	Unit
2x3 8' Dry	31,680	fbm
2x3 10' Dry	79,200	fbm
2x3 12' Dry	118,800	fbm

Resource utilization:

Product	Duration	Unit
Kiln Dryer #2	216,000	s
	60	h

Table 4 Example of a drying process

In this unit, there is another kind of drying process which is not carried out in a dry kiln. This process is called air drying and it is mainly used in summer. It occurs in the lumber yard, the duration is counted in weeks, it generates a better quality product and uses less energy to produce. However, air dried products must still go into a kiln dryer for the final step of the process for a smaller duration.

The last step is the finishing unit. At this stage, dried lumber is ready to be planed, sorted and packaged. This operation sets the final attributes to the lumbers (i.e., length, width, thickness and quality). As illustrated in Figure 8, this process is divergent (i.e., several final products are produced). Each rough piece of lumber is optimized based on its defects in order to decide whether that piece of lumber needs to be cut in two and where or if trimming can remove a defect on one side or the other in order to increase its market value.

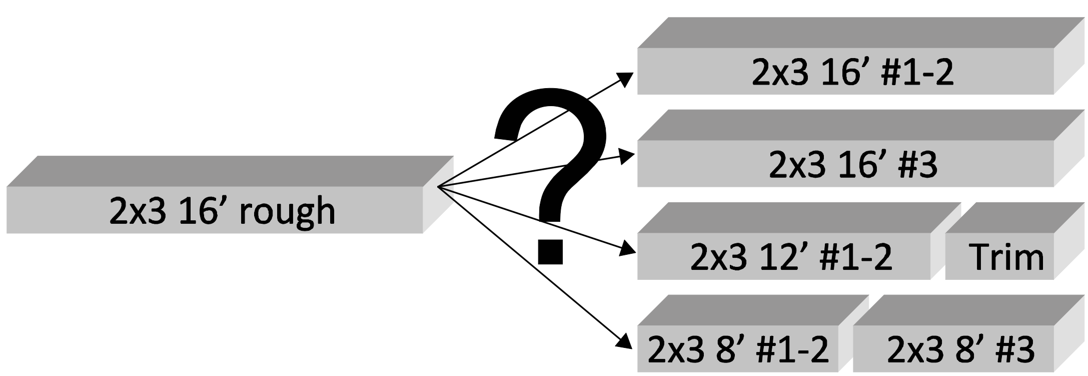

Figure 8 Example of possible outputs for a rough lumber

Because a large quantity of lumber is produced every day, it is possible using production reports or simulation data from Optitek to model the output of a process as a percentage that expresses the expected output distribution in terms of lumber final characteristics.

An example is shown in Table 5. For a 2x3 12', more than half of the pieces of lumber will conserve the same length (12') and will be of the indicated quality (#1-2, #3 or #4) after the planing process. It is also important to note that the finishing process creates trim.

Consumption:

Product	Quantity	Unit
2x3 12'	1	fbm

Production:

Product	Quantity	Unit
2x3 8' #1-2	0.007125	fbm
2x3 10' #1-2	0.12065	fbm
2x3 10' #3	0.0133	fbm
2x3 10' #4	0.004465	fbm
2x3 12' #1-2	0.665	fbm
2x3 12' #3	0.07505	fbm
2x3 12' #4	0.06365	fbm
Trim	0.05076	fbm

Resource utilization:

Product	Duration	Unit
Finishing Line	0.22	s

Table 5 Example of a finishing process

2.2.3 Customer Planning Unit

Two types of customer have been considered for the *Simulated Customers agent*: the spot market customer who expresses a one-time need for a specific volume of a particular product, and a contract-based customer who is a regular customer and whose needs are filled within a previously agreed upon contractual setting.

The sales of lumber products can be achieved through the spot market, which is common in commodity industries. In this context, the customer expresses its needs to several forest companies sporadically and at random intervals. The customer and chosen supplier then negotiate the final price, the volume and a delivery date. The customer and the producer both exhibit opportunistic behaviours. They have no obligation to one another, other than to fulfil the terms of this unique sale.

A contract involves a relationship and certain expectations on both sides. The contract may specify several requirements such as guarantees of purchase or deliveries, conditions of pricing and payments. It combines repeated deliveries over a set time horizon. In addition, demands from the contract customer occur at relatively regular intervals compared to the spot customer where it is more sporadic by nature.

There is little data available to model these two types of customer behaviours. Generally, companies only keep the information on shipped quantities, not demand information. In particular, if a substitute product is shipped, the original demand information disappears from the database. Therefore, it is very difficult to replicate exactly how customers behave in real life without making assumptions.

To get around this problem we made the assumption that deliveries were the same as demands to statistically model the demand processes just described. For each product, the total demand volume over the planning horizon bounds the summation of the different demands occurring within the planning period. This parameter is based on the historical deliveries of an enterprise in the industry.

The total amount requested was fixed as a percentage of total capacity of the mill in supply-driven production. To compute this capacity, the FEPP was configured in order to maximize the expected value, subject to production and supply constraints, and no demand constraints. From the resulting production figures, the maximum production capacity was derived from the total products volume. The total spot customer demand corresponds to about twice the capacity of the plant.

3 Simulation Components

The simulator consists of three main modules: the FEPP intended to simulate the detailed planning process of the production and distribution network, a customer agent simulating the behaviour of the market, and the simulator components, used to create, analyze and compare different scenarios. The representation of this simulator can be found in Figure 1.

3.1 Time Management

Time management is an issue in a distributed simulation because each agent handles local events that follow a global precedence network that, in turn, must be coordinated. Furthermore, because agents use complex planning algorithms that usually take time to produce a complete operations plan, there are two time scales to coordinate: the simulated time that triggers external events such as the expression of a demand, and the real runtime of each agent computing new plans.

In the proposed simulator, the simulated time is coordinated through a centralized clock. All agents use the same clock in order to make sure they all have the same perception of the sequence of simulated events. Thus, there is no risk that an agent is left behind (or races ahead) of others. Another solution often used in the literature is to allow each agent to use its own clock. In this case, messages must contain the date and time of transmission to allow for coordination between agents. However, this latter solution was not practical as the architecture of the simulator did not permit quick changes with a sufficient degree of confidence. Agents could react negatively to a sudden change in time, and would not be able to adapt to a date received if it were "in the past".

The simulation uses discrete events to run (Ross 2002). The simulator module contains a list of every action to be performed by the agents present in the simulation that enables the simulator to control the simulation. Each action is parametrized by the following characteristics:

- Agent name;
- Action name;
- Specific parameters;
- Trigger time.

The simulator has a list of all agents involved in a simulation and knows their state (working or standby) at all times. Time advances in real time when at least one agent is working. When all agents are in a standby mode, the simulator looks for the next action to achieve, sets the clock ahead accordingly and then asks the agent to follow-through with the targeted action.

When an action is triggered, the simulator delivers instructions to the target agent concerning what to do with specific parameters. Those parameters can contain information how to perform a specific task or an object. When an agent receives such information, it replies to the simulator by indicating that it is in a working mode and then performs the requested action. Once the task is completed, the agent reports to the simulator and returns to a standby mode.

The action carried out by an agent can trigger an action by another agent using a conversation. The first agent sends a message to the second. The latter receives the message and adds a request for the new action in the simulator actions list. The action may be triggered immediately or later depending on the needs of the requesting agent.

3.2 Key Performance Indicator

Key performance indicators are a necessity to determine which scenario is better. The first one that comes to mind would be to use financial indicator. The challenge with this approach is that there are a lot of imponderables (e.g., how to quantify a continuous refusal of sales or systematic late deliveries). These impacts are very difficult to calculate for a supply chain, even in real-life, while it is very easy to understand that a enterprise with a poor service level will incite its customers to look elsewhere.

In order to cope with these flaws, an indicator based on the volume of non-discounted sales will be used as shown in Equation (1). A non-discounted sale is a single demand that is planned to be fulfilled on time without using a substitute product as is done in this industry. The higher the total volume of non-discounted sales is, the better the value creation network is able to fill the needs of its customers.

$$\sum_{\text{Demand} \in \text{DemandsFulfilledOnTime}} \text{Volume}_{\text{Demand}} \quad (1)$$

To counteract the obvious case of accepting all demands and planning big delays in the fulfilment of these demands, two more indicators will be used to measure the service level offered to the customer.

The first one is the fillrate. It evaluates the number of demands that will be fulfilled on time (see Equation (2)) expressed in percentage. Even though the customer will only accept a demand if it is on time, the value creation network has a flexibility to change a scheduled delivery in case a more important (higher value) demand arrives.

$$\frac{\sum_{\text{Demand} \in \text{DemandsFulfilledOnTime}} \text{Demand}}{\sum_{\text{Demand} \in \text{ConfirmedDemands}} \text{Demand}} \quad (2)$$

The last one is the backorder. This indicator evaluates how long a demand will remain in backorder when not fulfilled on-time (see Equation (3)).

$$\frac{\sum_{\text{Demand} \in \text{LateConfirmedDemands}} \text{BOLateness}_{\text{Demand}}}{\sum_{\text{Demand} \in \text{LateConfirmedDemands}} \text{Demand}} \quad (3)$$

3.3 Random Number Generator

A single *Simulated Customers agent* emulates the behaviour of multiple customers. It is responsible for generating the characteristics of all demands that will be used in the simulation. This agent is parametrized using random variables that represent different ordering aspects.

The rather large amount of random numbers necessary for this kind of simulation implies the need for a rapid, robust and malleable generator capable of serving the needs of different models. The uniform random number generator used was the Mersenne Twister (Matsumoto & Nishimura 1998). This generator can produce random numbers with a very large cycle (the quantity of generated number before it loops to the first one) without slowing down the performance.

However, numbers that follow a uniform distribution are not useful to establish customer demands. Therefore, it is necessary to transform the output of the generator into numbers that can follow different types of distribution, such as triangular, normal or exponential (Ross 1993).

The method used in this application is simple but effective. The first step is to transform the density function of the desired distribution into a discrete form. For example, to generate a number following a normal distribution, a graphic representation of the distribution is divided by the number of intervals. Each interval is associated with a probability which is calculated by the area under the curve divided by the total area. Thus, by choosing a uniform number in the interval [0, 1], it is possible to find the corresponding value using the cumulative weight of the interval. An example is provided in Figure 9; if the chosen number is 0.47, then it is the value associated with the second interval that will be selected. Of course, it is always possible to approximate the areas under the curve if the number of intervals tends to be high. From this point, a random number will correspond to a random number following the desired distribution.

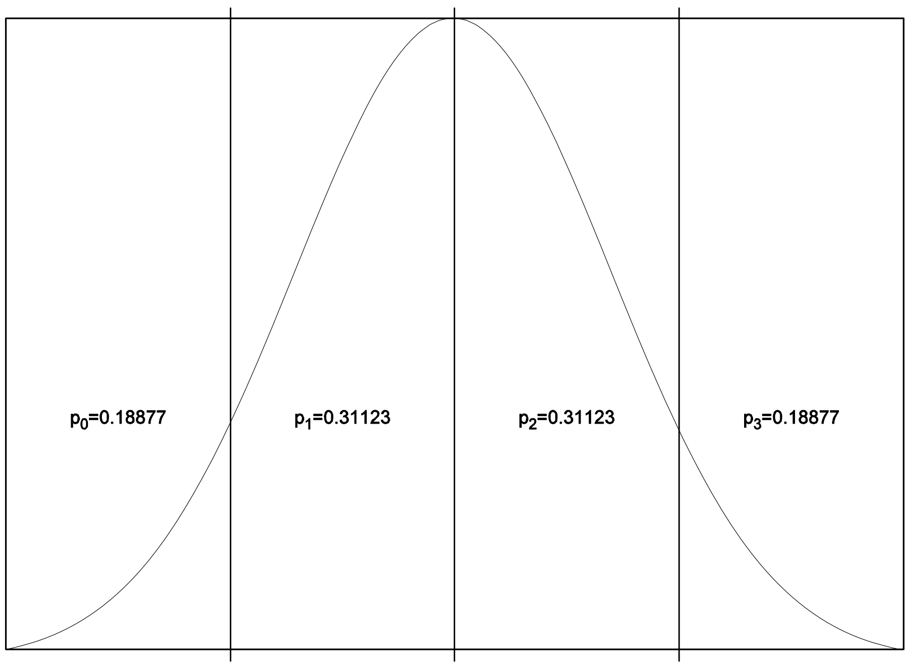

Figure 9 Example of a normal curve

3.4 Customer Generation Model

Instead of generating demands during the simulation, the *Simulated Customers agent* creates them before the start of a simulation. This approach facilitates the analysis, transformation and replication of demands without going through a whole simulation run. Furthermore, the validation of a modelled customer is easier to do since the simulation does not have an impact on the generation of demands. This generation is based on the demand modelling defined by Chang & Makatsoris (2001). Consequently, it is necessary to add to these pre-generated demands some information that will be used by the agent to “hide” futures demands.

3.4.1 Spot-based Demand

The general goal is to represent the behaviour of a typical sporadic customer. In order to do this, a demand process (i.e., the process that a demand would go through) was first developed. According to this demand process, the customer calls the sales department in order to know if a quantity of product is available for a given date. Then, if the supplier replies with a positive answer, the customer confirms its need, which becomes a formal order. Otherwise, the assumption is made that the customer withdraws its demand and chooses another suppliers to fulfil the demand.

In summary, for each customer-supplier-product relationship, the agent generates a list of demands which have several characteristics:

- Date of delivery (when the customer will want the product);
- Quantity;
- Effective Date (when the customer will ask for a demand);
- Status (planned, confirmed or removed and initially set to planned).

The demand generation algorithm is illustrated in Figure 10. The first step of the algorithm is to generate a random number that corresponds to the total amount that will be demanded for this customer-supplier-product relationship under this behaviour over the whole simulation horizon. As long as there is quantity left from this total amount, the algorithm creates a new demand. Each demand created will be characterized by a random quantity, a delivery date influenced by the seasonality of the desired product and a lead time that will be drawn to create the effective date by subtracting it from the delivery date. The effective date corresponds to the date when the customer sends the demand (e.g., the date of the call in the demand process). This information is crucial for the simulation process as the demand will be announced to the supplier agent only when the clock reaches the effective date of the demand.

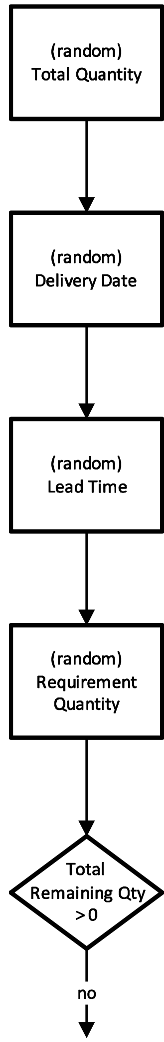

Figure 10 Spot demands generation

3.4.2 Contract-based Demand

The contract-based customer, as opposed to the spot market customer, cannot cancel a demand if it is not filled in time. Each demand must be filled in chronological order and the supplier is entitled to penalties in the case of late fulfilment. The assumption is made that an agreement binds the supplier to the customer and the supplier should be able to meet the demands of its customers or pay penalties.

The generation process of contract-based customer demands also differs from the process used for spot market customers. The first part of the algorithm, as outlined in Figure 11, is represented by the first three boxes which are used to draw up a list of dates that will represent the dates when demand must be fulfilled. From the start date of the simulation, the list is completed by drawing an interval with the next demand until the expected date of a demand is no longer within the time horizon of the simulation.

Once this list is complete, a global quantity will be drawn and distributed using seasonality of the product on the randomly established delivery dates. For each demand date in the list, the algorithm will assign a percentage of the total volume representing the “weight” of the date in the seasonal curve of total weightings.

The last step is to introduce in these demand figures a certain degree of “noise” that represents the numerous uncontrolled variations of demand in the market. To do this, we introduce a Gaussian noise, centred at an average of one and a standard deviation representing the degree of variation, which will multiply the quantity.

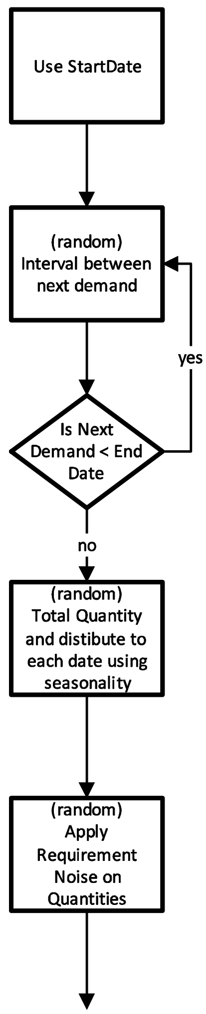

Figure 11 Contract demands generation

3.5 Customer Behaviour Model

During a simulation, the Simulated Customer agent adopts certain behaviours to better mimics a real customer. The model uses demand data, generated from the actions described in the previous section (3.4).

3.5.1 Spot Customer

At the beginning of a simulation, all demands that must be sent before the start date of the simulation, those that the effective date is before the start date, are assembled into a demand plan. A demand plan is a group of demands sent to a supplier agent. This demand plan is the initial plan that is sent to the *Deliver agent* of the supplier planning unit (Figure 12). According to its planning strategies, the Production and Distribution Network (see Figure 5) will conceive a supply plan to the best of its abilities and will try to satisfy all demands.

In turn, the customer reacts according to the “quality” of the proposed supply plan. For each demand, the customer calculates the date when the current demand is completely satisfied. If the date is later than the required date, the customer will remove this demand from the list. The supply allocated to the cancelled demand will be reallocated to forthcoming demands. On the contrary, if the date is acceptable, the demand will change status from “proposed” to “confirmed”. This indicates that both the supplier and the customer are aware of this demand and the supplier must do everything possible to maintain this planned delivery date.

At this point, if the simulation run is not completed, the Simulated Customer agent creates a new demand plan. This plan contains the confirmed demands from the previous plan generated during the previous planning cycle. Furthermore, any new demands for which the effective date falls before the current simulation date is added to this demand plan.

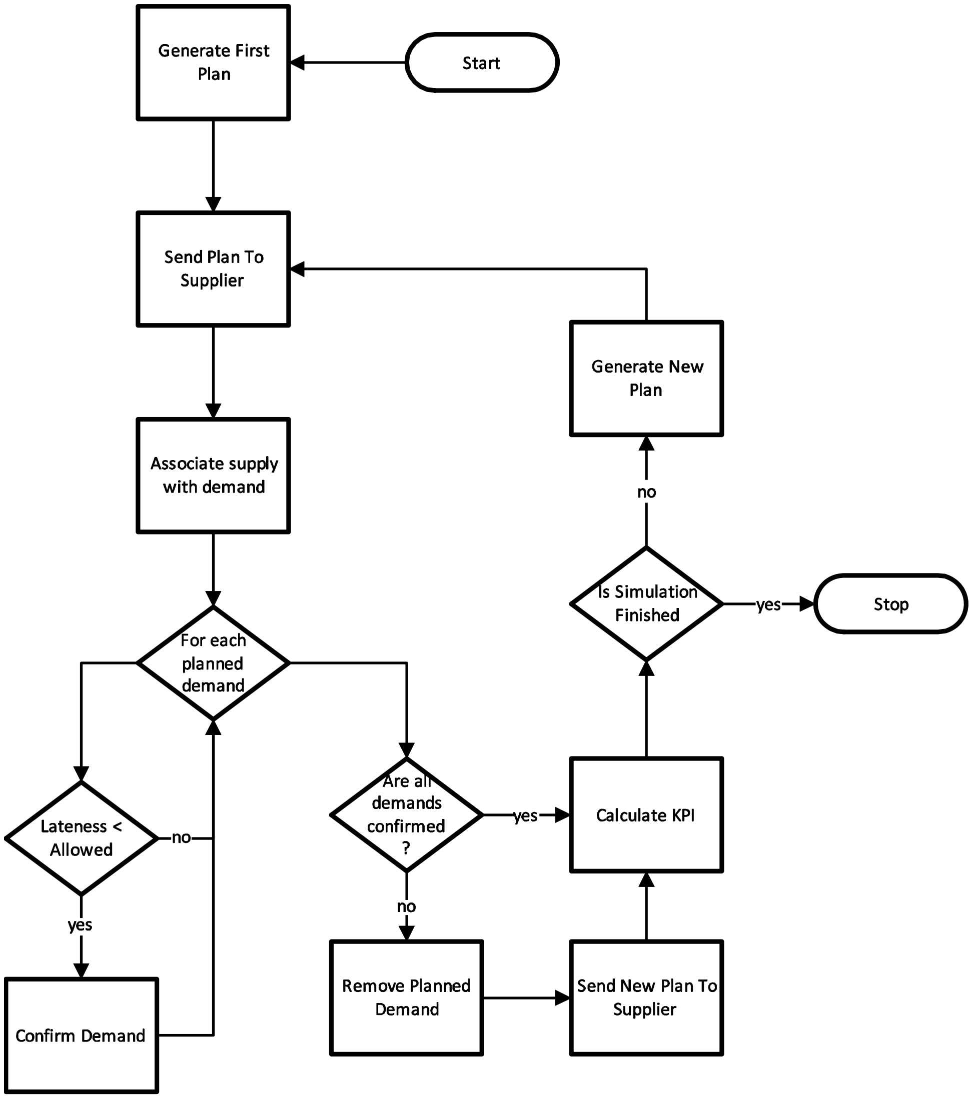

Figure 12 Behaviour of Spot Customer

3.5.2 Contract Customer

The planning horizon is divided into three subhorizons: frozen, evolving and provisional. These subhorizons are used to introduce perturbations in the demand plan using the principle that the farther the demand is in the time horizon, the greater its uncertainty is. In the frozen horizon, which starts with the current simulation date, demands are fixed and cannot be changed for any reason. At the other end of the planning horizon, the provisional horizon of the demand plan can change completely since there are no restrictions to changes. Between these two horizons, the quantity of each demand may change, but the increase or the decrease is bounded.

During the initialization procedure of the simulation, the Simulated Customer agent creates an empty initial demand plan for the contract customer. To this plan, the customer adds demands that fall into the planning horizon. This plan is then sent to the Production and Distribution Network which will try to meet the needs (Figure 13).

Then, if the simulation is not finished, a copy of the last plan is made. This operation makes it possible to keep information about previous plans. The horizons are adjusted to the current simulation date in the copied plan and new demands that fall into these horizons are added. Existing demands present in the evolving and provisional horizons are modified with respect to their bounds.

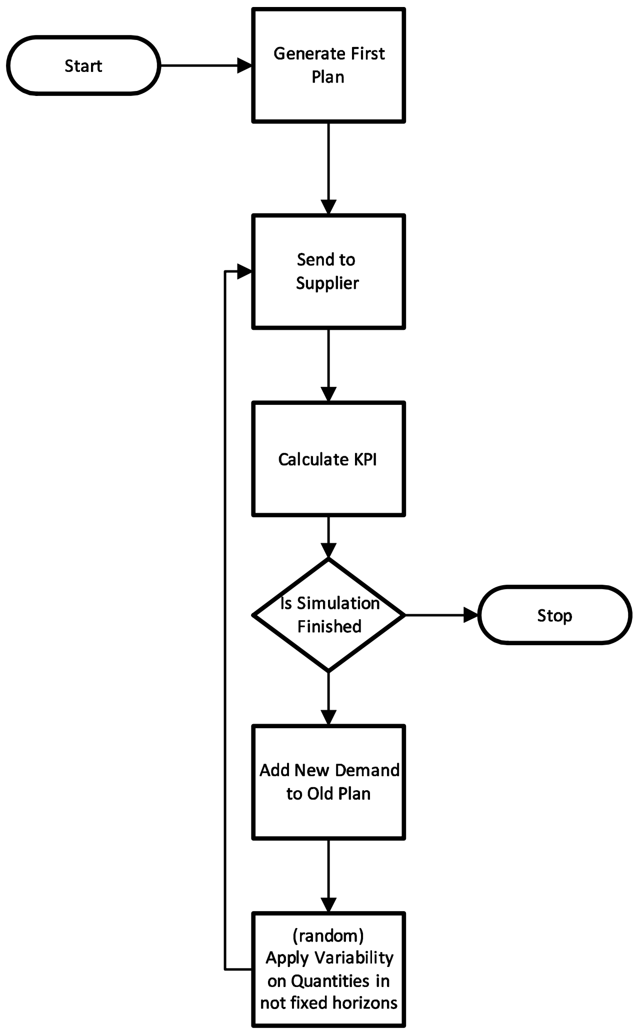

Figure 13 Behaviour of Contract Customer

4 Experiment

The objectives of this paper are to evaluate the impacts of different planning strategies in a supply chain and simulate them using a multi-agent simulator integrating an APS. In order to do this, two experiments were developed to evaluate and compare different scenarios using a total of 13 agents representing business units in a lumber value creation network and essential agents needed by the FEPP. The goal of these experiments is to observe the impact of a demand-driven approach in a lumber supply chain. In other words, these experiments propose to study the effects on a supply chain when a company plans its operations taking into account customer demands compared to a supply chain which doesn't, as is often the case in a commodity based industry. In such a case, the company aims to produce higher value products without considering actual demand. It is only afterward that they balance demand with supply using substitutions, discounting or inventory strategies.

In order to achieve a long horizon simulation, the simulator uses the rolling horizon concept. This is the main difference with Cid Yáñez et al. (2009). From a given date, the planning is done for four weeks. After every agent has completed all of their active taskflows, the simulator sets the clock ahead one or two weeks depending on the experiment (see Table 6). For each agent, the new on-hand inventory condition is computed using the inventory at the former start date, plus the planned production of the previous period (it is assumed that production happened exactly as planned), minus the quantity shipped to the next agent in the supply chain.

For these experiments, a six-month horizon was chosen. Over the course of a simulation, each agent will need to perform the local planning of its operations. The total number of times it happens depends on the strategy used in the experiment and can vary from 12 to 48 during the simulation horizon.

4.1 First experiment

In this experiment, the spot customer is used in order to anticipate the performance variations of a lumber supply chain when demands are propagated throughout the network (demand-driven) versus when the production is pushed (supply-driven).

The supply-driven approach is traditionally the one used in the North American lumber industry. In a supply-driven approach, the corporate production planner sets goals for each production facility in the form of production mix and yield targets over a given time period. In each facility, the local production planner uses these targets in order to set operational plans. Even if a demand variation occurs in the market, the production will stay the same and may only change the next time the corporate planner revises these targets.

In contrast, demand-driven occurs when the planner takes customer demands into account. This approach raises planning challenges especially in a context of divergent production processes. It is indeed extremely difficult to create manually a production plan in order to meet all customer demands. From the point of view of the

Production and Distribution network, the solution to this problem is to find a compromise between all the demands from the customer and the available log supplies in order to create an operational plan that respects the production constraints.

In this experiment, the Virtual Lumber case was used and the two different planning strategies were tested. In the first one, the production planning unit takes a supply-driven approach. Only the *Deliver agent* knows the exact customer demands. The sales department uses a strategy called Available-to-Promise (ATP) which is the most common approach in the lumber industry. Its motive is to sell the lumber produced by the mill. This strategy is best explained in Ervolina et al. (2006). It is used as the reference case in this experiment.

The second strategy tested is the demand-driven approach. Here, the different agents optimize their local production plans according to the customer demands and anticipate future demands by using the remaining capacity to produce the expected higher value products. In this approach, before giving a response to a customer demand, the agents verify that the current demand can be fulfilled by making changes in the production plan or changing the allocation of the production. All agents are called upon to validate or plan their operations again in order to find the best synchronized plan to meet final customer demands.

4.2 Second experiment

The idea in the first experiment was that a demand-driven approach was better to serve the customer but the results were not conclusive. So, a second experiment was necessary and it is derived from the first. It uses the same two planning strategies (supply-driven and demand-driven) and integrates a new one. This new strategy involves answering customer demand with on-hand inventory the same way as the supply-driven approach. But, when the customer confirms a demand (see Section 3.5.1 for more information), this demand is used to influence the production planning like the demand-driven approach.

This filtered demand-driven approach (only confirmed demands are propagated through the supply chain) should better serve the customer than the demand-driven approach by reducing backorder. Also, by using filtered demands propagated through the value creation network, it should help stabilize the production while influencing this planning process.

Experiment	Strategies used	Planning interval
First	Supply-driven, Demand-driven	2 weeks
Second	Supply-driven, Demand-driven, Filtered Demand-driven	1 week

Table 6 Experiments summary

5 Results

The results for the first experiment can be seen in Figure 14a. The results are computed on the basis of non-discounted sales over 10 simulation replications. These sales correspond to exact needs expressed by the customer as described in Section 3.2. On Figure 14a, it is possible to see that using a demand-driven approach increases the volume of non-discounted sales.

But this increase of sales puts more pressure on the supply chain and this leads to a worse service level for the customer. In fact, the fillrate decreases, which leads to a greater number of demands fulfilled late and a longer time in backorder as can be seen with the backorder indicator.

In theory, a demand-driven approach should be better adapted to fulfil customer demands. The problem in this experiment might come from the delay between two planning cycles. For example, if a demand is planned to be fulfilled late, the next chance it would get to be on time would be in two weeks. To cope with this, more frequent planning would provide a better way to increase sales while maintaining a good service level.

That was the premise of the second experiment. And as can be seen on Figure 14b, a better control of the planning process increases the volume of sales while maintaining a relatively good service level over 12 simulation replications.

Moreover, by using a filtered demand-driven approach, it is possible to increase the total volume of sales while increasing the service level. This can be explained by the fact that the supply chain is not influenced by undesired demands compared to the demand-driven approach. In the latter strategy, all demands influence the planning process. Many of these demands will be removed by the customer and contribute to create some sort of noise in the supply chain planning process when using the demand-driven approach. When using the filtered demand-driven approach, the *Deliver agent* acts as a demands filter for the whole supply chain and helps improve all the indicators of the whole value creation network.

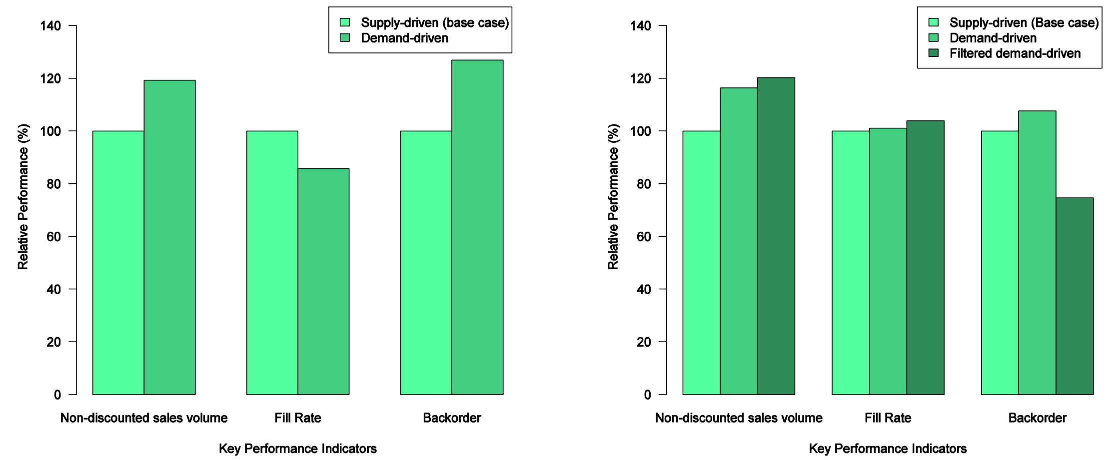

(a) First Experiment (b) Second Experiment

Figure 14 Experiments Results

However, these results should not be interpreted as a recommendation for the global lumber industry. Each sawmill is different and would probably achieve a different result for this experiment. Moreover, it is possible to use a multi-agent simulation to analyze such a tactical problem, and consequently, better analyze its impacts in a value creation network and be in better position to make a decision.

6 Concluding Remarks

As demonstrated in the previous section, it is possible to use the FEPP as a simulation framework to conduct robust experiments in a variety of fields. This is a solid and very versatile platform that can accommodate studies of revenue management, planning problems, strategic decision and forest planning.

To demonstrate the simulation capabilities of the FEPP, another experiment is available in Lemieux et al. (2008) where the contract customer is present. The objective of this experiment was to find the optimal proportion of contracts a value creation network can accept without decreasing the total sale value of all its customers.

Furthermore, the sawmills used in these studies can in fact be modelled after real sawmills. This creates an opportunity to study strategic questions without impacting the day-to-day operations. This would enable decision makers to simulate different scenarios before making a binding decision, resulting in a better informed decision.

Consequently, future work will include the validation of the results presented in this paper. This objective will need an increased number of replications of the simulations described to ensure a statistical quality of the result obtained. A system is being developed to manage the simulations. This system will be able to launch a predefined number of simulations in parallel and will incorporate a service to gather the resulting key performance indicators of a simulation. Once a simulation is complete, it will start a new one without the need of interaction with the user.

Acknowledgements

This work was funded by the FORAC Research Consortium and the Natural Sciences and Engineering Research Council of Canada (NSERC). The authors would also like to acknowledge the impressive work of many FORAC research assistants and students who have worked on developing the FEPP over the years.

References

Balci, O., 1986. Credibility assessment of simulation results. In *Proceedings of the 18th conference on Winter simulation*. Washington, D.C., United States: ACM, pp. 38-44. Available at: http://portal.acm.org/citation.cfm?id=318242.318258&type=series [Accessed November 3, 2009].

Chang, Y. & Makatsoris, H., 2001. Supply chain modeling using simulation. *International Journal of Simulation*, 2(1), 24–30.

Cid Yáñez, F.[. et al., 2009. Agent-based simulation and analysis of demand-driven production strategies in the timber industry. *International Journal of Production Research*, 47, 6295-6319.

D'Amours, S. et al., 2006. IFIP International Federation for Information Processing. In Boston: Springer, pp. 17–26.

D'Amours, S., Rönnqvist, M. & Weintraub, A., 2008. Using Operational Research for Supply Chain Planning in the Forest Products Industry. *INFOR: Information Systems and Operational Research*, 46(4), 265-281.

Ervolina, T.R. et al., 2006. Simulating order fulfillment with product substitutions in an assemble-to-order supply chain. In *Proceedings of the 38th conference on Winter simulation*. Monterey, California: Winter Simulation Conference, pp. 2012-2020. Available at: http://portal.acm.org/citation.cfm?id=1218479 [Accessed November 3, 2009].

Forget, P., D'Amours, S. & Frayret, J., 2006. IFIP International Federation for Information Processing. In Boston: Springer, pp. 89–98.

Frayret, J. et al., 2007. Agent-based supply-chain planning in the forest products industry. *International Journal of Flexible Manufacturing Systems*, 19(4), 358-391.

Gaudreault, J. et al., 2009. Distributed Operations Planning in the Lumber Supply Chain: Models and Coordination.

Goulet, P., 2006. Optitek : User's Manual.

Gronalt, M. & Rauch, P., 2008. Vendor managed inventory in wood processing industries-a case study. *SILVA FENNICA*, 42(1), 101.

Lemieux, S. et al., 2008. Intégration d'outils APS dans une simulation multi-agent: Une application à l'industrie du bois d'oeuvre. In *Proceeding of the 7e Conférence Internationale de Modélisation et Simulation*. pp. 708-715.

Matsumoto, M. & Nishimura, T., 1998. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.*, 8(1), 3-30.

Petrovic, D., Roy, R. & Petrovic, R., 1998. Modelling and simulation of a supply chain in an uncertain environment. *European Journal of Operational Research*, 109(2), 299-309.

Robinson, S., 2002. General concepts of quality for discrete-event simulation. *European Journal of Operational Research*, 138(1), 103-117.

Rönnqvist, M., 2003. Optimization in forestry. *Mathematical Programming*, 97(1), 267-284.

Ross, S.M., 1993. *Introduction to Probability Models* 5th ed., San Diego: Academic Press.

Ross, S.M., 2002. *Simulation* 3rd ed., Academic Press.

Sargent, R.G., 2005. Verification and validation of simulation models. In *Proceedings of the 37th conference on Winter simulation*. Orlando, Florida: Winter Simulation Conference, pp. 130-143. Available at: http://portal.acm.org/citation.cfm?id=1162708.1162736 [Accessed November 3, 2009].

Schneeweiss, C., 2003. *Distributed decision making*, Springer.

Stadtler, H., 2005. Supply chain management and advanced planning--basics, overview and challenges. *European Journal of Operational Research*, 163(3), 575-588.