Proceedings of the 2004 Winter Simulation Conference
R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

SIMULATION, A FRAMEWORK FOR ANALYSING SME SUPPLY CHAINS

PJ Byrne
Cathal Heavey

Enterprise Engineering Research Group
Department of Manufacturing & Operations Engineering
University of Limerick
Limerick, IRELAND

ABSTRACT

The following paper briefly presents the formulation and development of a case study supply chain simulation model as developed for an industrial company. The case study company in question is considered to be a vertically integrated organisation, offering a complete range of its related industries products to a global marketplace.

The paper reviews the scale of supply chain system being analysed, the type of data required to populate such a model and the performance outputs from the model. These outputs include the percentage of demand that is both On Time and In Full (OTIF%), the days of inventory held in finished stock and also the finished stock quantities.

The paper also reviews the scope of such a model by reviewing some of the experimental work as carried out on this model and highlights the usefulness of such a model as an aid to supply chain decision making in a SME.

1 INTRODUCTION

The primary objective of this paper is to examine the usefulness of simulation for analysing SME (Small-Medium-Enterprise) industrial supply chains, using a real industrial case as shown in Figure 1, as a basis for this analysis. The organisation involved in the study is classed as a vertically integrated raw material supplier, offering a complete range of its related industries products to a global marketplace. Schary et al. (1995) define vertical integration as “the ownership by one organization of other firms in its supply or distribution network. The totally integrated firm is completely self-sufficient. The non-integrated firm is completely dependant on market forces and other organizations for its operations.” Christopher (1998) states, “Vertical integration normally implies ownership of upstream suppliers and downstream customers.”

SME - EU definition: an SME has fewer than 500 employees, has a turnover of less than 38 million ECUs per annum and is no more than one-third owned by an organisation larger than an SME (based on turnover and number of employees), unless it is a financial investor such as a bank or venture capitalist, obtained from Forfas <`http://www.forfas.ie/publications/archive/rd/gloss.htm`>.

In this paper, a brief introduction is given to the supply chain being analysed, a description of the data populating the model and the process detail incorporated into the model and the results being extracted from the model. Finally, a number of experiments are presented and their results reviewed.

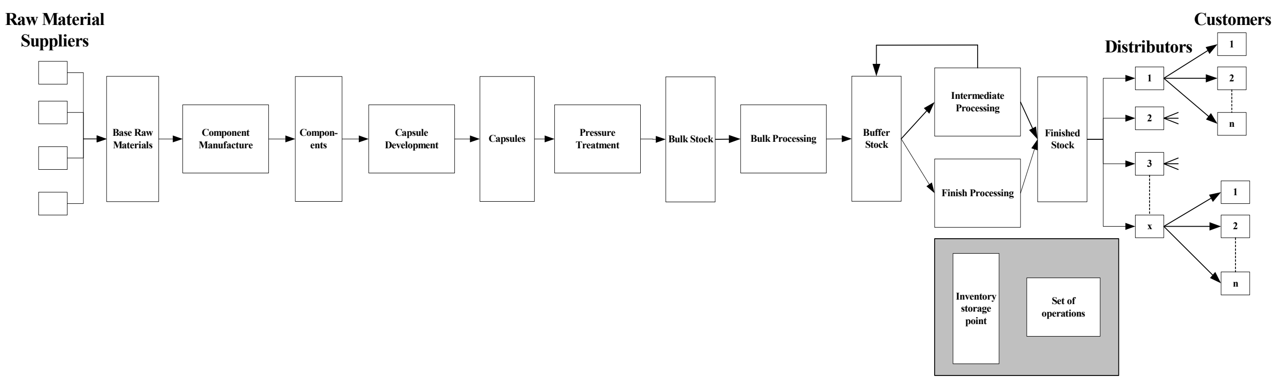

Figure 1: Supply Chain Structure

2 BACKGROUND

Although, as can be seen in the literature there is an extensive quantity of material reviewing supply chains, there are a number of areas, which require further study. It is one of these areas that this body of work is aimed at. The area presented includes the analysis of vertically integrated supply chains as opposed to logistical based supply chains, which have not been found in the literature to date. With vertically integrated supply chains, the production element becomes more important and has been addressed in this work. The second area addressed in this work is the application and implementation of supply chain modeling in Small-to-Medium sized Enterprises, as opposed to the much larger studies as carried out for example by IBM (Bagchi et al. (1998) and Lin et al. (2000)), Hewlett Packard (Lee et al. (1993a), Lee et al. (1993b) and Lee et al. (1995)) and Compaq (Ingalls et al. (1999)).

The companies supply chain structure as shown in Figure 1, can generally be regarded as an extension of the well studied serial multi-echelon supply chain, being supplied by a number of raw material suppliers and supplying a number of distributors, who in turn supply individual customers. As already stated the company is vertically integrated and owns and controls almost the entire supply chain across a global spectrum. The exceptions to this are a number of the multiple initial raw material suppliers and a number of the distributors. For the purpose of this review, the supply chain model has been analysed from the distributors to the bulk stock object to review the effect of the manufacturing function on the supply chain performance measures.

3 SUPPLY CHAIN MODEL FORMULATION

In the previous section the problem of interest was formulated, through meetings and conversations with company personnel, who identified areas to be analysed and the level of detail required from the study, such as the effect of the manufacturing function on the overall supply chain performance. It was also at this stage that key performance parameters that were to be used in the model were identified, which included:

- On Time In Full (OTIF) – which is the percentage of orders received into finished stock that can be satisfied completely from finished stock, assuming a make-for-stock policy is being operated.
- Velocity levels – this is the processing time of an order through a processing area.
- Days of Inventory in Finished Stock – This is the number of days that current stock levels can satisfy, based on a rolling historical demand.

It was also at this point that it was decided that the model developed would be built using an object orientated simulation package, eM-Plant. This package was chosen for a number of reasons. The first being the fact that it is object orientated, which allows for easy replication of objects (such as work centres, etc.). The second reason is its ability to use an ODBC link (Open Database Connectivity Link) to connect the simulation model directly to an information database, and the third main reason was the flexibility offered by eM-Plant in relation to its programming language (Sim Talk) to customise objects to realistically resemble the existing physical supply chain. This was especially useful when modeling the manufacturing function and for examining the effect of these functions on the overall supply chain performance.

4 SUPPLY CHAIN MODEL INTRODUCTION

The following section will briefly outline the key aspects of the supply chain in question, but for a more complete description of the supply chain in question and subsequent model development, see Byrne et al. (2004).

It was noted from a general overview of the company, the issues the company would like to have analysed and the level of detail incorporated in the supply chain that it would not be possible to model this system accurately and efficiently using analytical techniques. Therefore the use of simulation was decided upon as a viable process for analysing the supply chain. However with this approach it is necessary to introduce some simplifying assumptions to the model, which ease the development of the model but do not take from the overall structure of the model or the validity of the results. A number of key simplifying assumptions, which were regarded as necessary and justified prior to the development of the simulation model were outlined as follows:

- The use of a single processor in each work centre, with a suitably adjusted processing time as opposed to splitting batches across a number of machines, to eliminate the need for batch splitting in work centres, which in turn reduces the number of entities in the model at any one time, thus speeding up the processing time of the model.
- Substitution is used across all products in the company when necessary, but there is no formal procedure for it. Therefore a number of substitution tables were developed for products with high demand to incorporate the concept into the model.
- Daily demand is sent into the model in relation to the earliest required date from the customer.
- Daily demand for same products is sent into the model as a consolidated work order, i.e. if two distributors order the same product on the same date the two orders are consolidated and the work order is given the earliest distributor required date as the due date for the entire work order.

- Demand is only satisfied when the entire order can be fulfilled.
- It is assumed that there will be an infinite stock of material in bulk stock, which is the most upstream stage in the model currently being analysed.

The model itself is significantly detailed in particular in relation to product movement through the manufacturing processes, taking into account stocking locations. The model itself consists of:

- Approximately 17,000 Order lines processed in system in one year, from 8 different distributors.
- 42 work centres, which consist of 276 machines operating on 6 different shift patterns.
- 10,000 finished stock parts of which approximately 3,000 are active at any point in time.
- 1,000 buffer stock parts from which the finished stock parts are processed from.
- Approximately 60,000 routing lines associated with the 42 work centres and all processed parts.

4.1 Data Transfer – ODBC Link

The simulation model itself is connected to a number of database tables (Figure 2), which are stored in MS Access, by means of an ODBC link between them and eM-Plant. This link enables easy transfer of data from the real system into the simulation model and can be set up to continually update and maintain the information in the model.

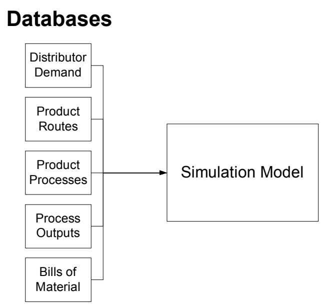

Figure 2: Data Transfer to Simulation Model

The databases, which are used in the model, contain information on the individual distributors demand, the product routes for parts as they flow through the system with associated standard times, the processes that produce different products and associated yield factors, the output from these processes and the overall bills of material.

5 MODEL EXPERIMENTATION

This section outlines the details and results of a number of experiments, which were carried out using the above-described model. The experiments have been constructed in such a way, so as to highlight the usefulness of such a model to a SME in a number of different areas. The experiments themselves were developed through consultation with the company. The experiments attempt to address issues, which are currently of interest to the company. Following this process, it was decided that three groups of experiments would be carried out on different aspects of the model. The first group of experiments examine the stocking policies and the planning process. The second group analyses the impact of production on system wide performance.

Although there are twelve distinct product families in relation to finished stock demand, only the top six product families will be analysed in the first experiment in this study and only the affected products in the remaining experiments.

5.1 Experimental Group 1 – Stock Policies (Part I)

The first set of simulation experiments were developed to investigate the effect of stocking policies on the supply chain performance parameters. The experiments were carried out in two parts. The purpose of the first part was to evaluate the extremities, in relation to the stocking policies under the following parameters.

- All products are MFS (All products are stored in Finished Stock), which was current practice.
- Limit MTO (selected products are stored in Finish Stock as determined by stocking policy rules, set out by the company, based on a four constraints, which include the following product analysis, (1) Pareto, (2) Volatility, (3) Peaky Demand and (4) a Declining Demand), which is the policy being analysed.
- All products are MTO (only excess produce is stored in Finish Stock), which can be used as a benchmark for the other two policies.

The second part reviews the recommended MTO limit rule and its individual constraints. Products are identified as MTO or MFS, depending on the outcome of four constraints, which were developed as part of a supply chain re-engineering exercise. The four constraints are outlined above.

5.1.1 On Time In Full (OTIF%)

It can be seen from Table 1 and Figure 3, that the change in stocking policy from All MFS, through a limited MTO to All MTO, caused a reduction in the OTIF%. It can be seen that the overall OTIF% drops from an average of 93.02% to 88.08%, if the stocking policy is to be changed from All MFS to a limited number of products manufactured as MTO, with the B Seg being affected most dropping from 70.4% to 43.8%.

Table 1: OTIF%

	All MFS	Limit MTO	All MTO
A Rnd	98.0	93.8	88.1
A Seg	93.1	90.8	90.9
B Rnd	89.9	79.1	78.0
B Seg	70.4	48.8	37.6
C Rnd	99.3	98.1	98.1
C Seg	84.9	85.3	81.4
Totals	93.02	88.08	84.93

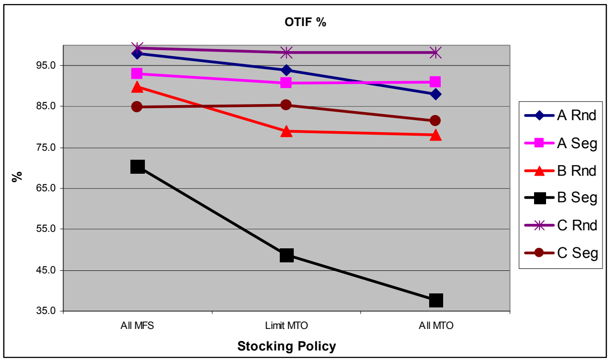

Figure 3: OTIF% - MFS – MTO

5.1.2 Average Days of Inventory

It can be seen from Table 2 and Figure 4, that there is significant reduction, in the range of 15 to 25 days in the average days of inventory in finished stock across each product family as the stocking policy moves from all items MFS to limited MTO.

Table 2: Average Days of Inventory

	A			B			C		
WEEK	All MFS	Limit MTO	All MTO	All MFS	Limit MTO	All MTO	All MFS	Limit MTO	All MTO
1	30	4	1	30	13	0	30	3	1
13	35	12	9	29	14	1	38	12	10
26	39	14	11	36	19	6	42	15	12
39	37	14	11	21	8	4	35	13	13
52	37	15	11	37	20	7	44	17	15

5.1.3 Finished Stock

It can be seen from Table 3 and Figure 5, that there is a significant drop in the quantity of finished stock held as the model stocking policy changes from all items MFS to limited MTO, with an average reduction of 24,000 products.

5.2 Experimental Group 1 – Stock Policies (Part II)

Whereas the previous section reviewed the use of stocking policies themselves using the simulation model, this section assumes that a MTO/MFS stocking policy based on a number of tests is to be used. With this procedure items, which pass the test will be manufactured as MFS items and the remaining items will be manufactured as MTO. The purpose of this section is to assess a number of different tests for categorising products as MFS/MTO and to evaluate these different tests in relation to the systems supply chain performance parameters.

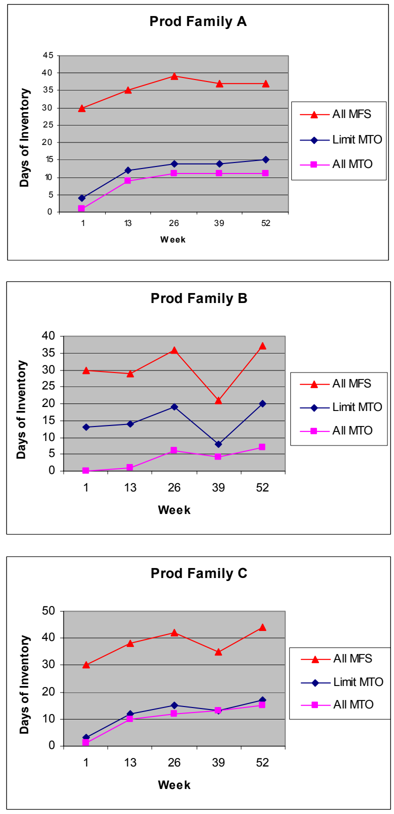

Figure 4: Average Days of Inventory – MFS – MTO

Table 3: Finished Stock

WEEK	All MFS	Limit MTO	All MTO
1	35472	14879	5661
13	31869	13191	4819
26	32405	13099	4570
39	31028	12378	4296
52	31769	12788	4264

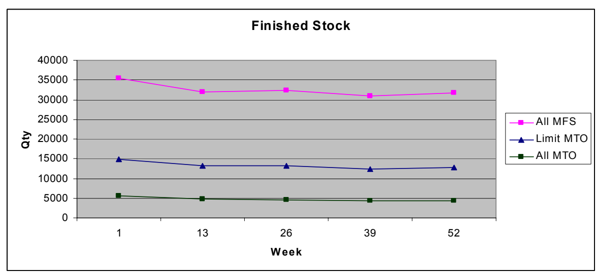

Figure 5: Finished Stock – MFS – MTO

The five test policies to be evaluated are, YEEE, YEEN, YNEN, YENN and YNNN, where Y = Yes, N = No, and E = Either Yes or No. So for example rule 2 (Y E E N), will stock items (MFS), which are pareto items (signified by the Y in position 1) and not on a downward slope (signified by the N in position 4), but doesn't take into account the volatility or the excessiveness of the product (signified by the Es in position 2 and 3).

5.2.1 On Time In Full (OTIF%)

It can be seen from Table 4 and Figure 6 that the general trend is a consolidation or slight reduction in the OTIF percentages as more products are added to the MTO category (i.e. as the stocking policy moves from YEEE to YNNN). There is a slight exception to this on the YEEN run where a number of the product families OTIF percentages increased as the number of MFS items decreased. However this can be attributed to a different product mix in the system caused by a different production policy. However it should be noted that the OTIF percentages ranged from a high of 86.91% to a low of 84.59%, which is a drop of 2.32% or 283 out of 12195 sales order lines.

Table 4: OTIF Percentages – MTO Limits

	YEEE	YEEN	YENN	YNEN	YNNN
A Rnd	93.98	93.98	93.83	93.98	93.83
A Seg	91.30	91.52	90.78	91.16	90.78
B Rnd	84.36	86.08	79.06	82.14	79.06
B Seg	55.63	59.66	48.83	53.93	48.83
C Rnd	98.13	98.13	98.13	98.13	98.13
C Seg	85.29	85.29	85.29	85.29	85.29
Totals	86.29	86.91	85.61	84.77	84.59

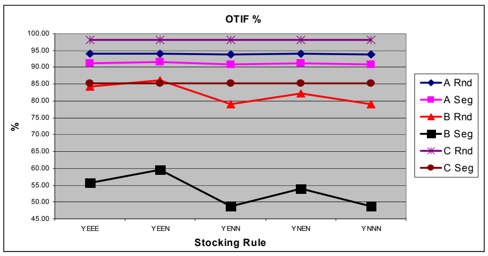

Figure 6: OTIF Percentages – MTO Limits

As can be noted from the data for the major product families, there is very little change in the OTIF percentage for product family A and C, but with a significant difference in product family B. One of the major reasons for this is the typical demand pattern for the different product families. Product families A and C tend to have more frequent smaller orders, with family B having less frequent but larger orders, therefore the reduction in MFS items impacts family B to a greater extent.

5.2.2 Average Days of Inventory and Finished Stock

For this part of the experiment the average days of inventory in finished stock and the average quantity of FREs in finished stock shall be analysed together. Table 5 and Figure 7 present the average days of inventory in finished stock for each product family and Table 6 and Figure 8 present information on the quantity of finished stock for each run of the model.

It can be seen from the observed data in Table 5 and Figure 7 that as the stocking policy moves from YEEE to YNNN (i.e. a reduction of products which are MTO) the trend is that of a reduction in the average days of inventory in finished stock. It can be seen that there is on average a reduction of 6 days inventory for product family A and B, and 2 days for product family C as the stocking policy moves from YEEE to YNNN. Similarly the total quantity of finished stock drops on average 11000 as the stocking policy moves from YEEE to YNNN.

Table 5: Average Days of Inventory

	A					B				
WEEK	YEEE	YEEN	YENN	YNEN	YNNN	YEEE	YEEN	YENN	YNEN	YNNN
1	11	7	4	6	4	15	15	13	13	13
13	18	15	12	13	12	16	16	14	14	14
26	21	17	14	16	14	21	21	19	19	19
39	21	17	14	16	14	10	10	8	8	8
52	21	17	15	16	15	22	22	20	20	20

	C				
WEEK	YEEE	YEEN	YENN	YNEN	YNNN
1	9	7	3	4	3
13	18	15	12	13	12
26	21	19	15	16	15
39	17	14	13	14	13
52	23	21	17	18	17

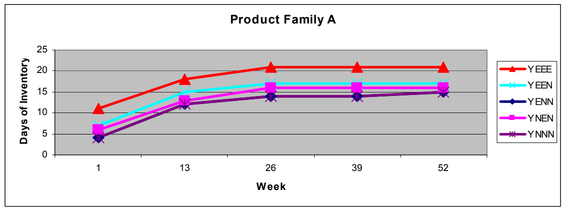

Figure 7: Average Days of Inventory – Example

Table 6: Finished Stock

WEEK	YEEE	YEEN	YENN	YNEN	YNNN
1	26864	24382	15309	20393	14879
13	23868	21567	13786	17954	13191
26	23636	21230	13562	17301	13099
39	23143	20689	12930	16811	12378
52	23542	21069	13201	17144	12788

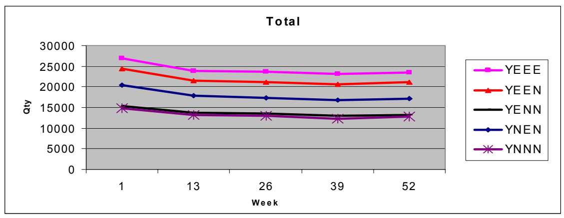

Figure 8: Finished Stock – MTO Limits

5.3 Experiment 1 – Conclusions

It is clear to see from the simulation experiment that there was a significant difference in the results obtained under these three stocking policies (Part I) in relation to the OTIF percentages, the average days of inventory in finished stock for each product family and the average finish stock levels, as would be expected and as can be seen in Table 7.

Table 7: Experiment 1, Part I Sample Results

	OTIF% for A Rnds	Avg Days in Finished Stock for A	Total FREs
All MTO	88.14	11	4722
limit MTO	93.83	14	13267
All MFS	97.97	38	32509

Although it is possible that all products could be manufactured as MTO or MFS it is very unlikely in this industry due to customer requirements and commitments. Therefore the second part of the simulation experiment was set up to review the “limit MTO” stocking policy itself and to vary its constituent parts.

For an example of these results view Table 8 for product family A and B. Whereas each product family followed the same overall patterns as in part I of this experiment, not all product families follow the same pattern in this the second part of the experiment. For example the OTIF percentage for the A Rnds remained constant across each run of the experiment, while the OTIF percentage for the B Rnds fell as more products were categorised as MTO (moving from YEEE to YNNN). The general trend was that of a reduction in the number of days on inventory in finished stock for A and B Rnds, and the total FREs stored in finished stock generally decreased. It can be seen however from the more detailed data that there is a general decrease across the range in the overall OTIF percentage (from a high of 86.91% to a low of 84.59%, which is a drop of 2.32% or 283 out of 12195 sales order lines). There is also on average a reduction of 6 days inventory for product family A and B, and 2 days for product family C as the stocking policy moves from YEEE to YNNN. Similarly the total quantity of finished stock in FREs drops on average 11000 as the stocking policy moves from YEEE to YNNN.

Table 8: Experiment 1, Part II Sample Results

	OTIF% - A Rnds	Avg Days in FS - A Rnds	OTIF% - B Rnds	Avg Days in FS - B Rnds	Total Finished Stock
YEEE	93.98	21	84.36	9	24211
YEEN	93.98	17	86.08	9	21787
YNEN	93.98	14	82.14	5	13758
YENN	93.83	16	79.06	6	17921
YNNN	93.83	14	79.06	4	13267

5.4 Experimental Group 2 – Production Controls (Part I)

The purpose of this set of simulation experiments was to determine the effect of certain production controls on the overall supply chain performance parameters. The previous experiments were run with no production rules in place. In other words, orders simply moved as soon as they could from work centre to work centre, as was the procedure being used in the SME, *Company X*. In reality this causes large quantities of WIP in the bottleneck centres. In order to limit this, *Company X* have recently introduced the idea of buffer restrictions ahead of work centres to limit this build up.

As with the last set of experiments, this experiment has been conducted in two parts. The first part introduces the concept of buffer restrictions to the model. This model was run on five different scenarios, setting the maximum allowable work orders in any finishing work centre to 5, 10, 20, 30 and 40 work orders with the exception of the manual inspection work centres.

5.4.1 On Time In Full (OTIF%)

The OTIF percentages using the buffer restrictions are shown in Table 9. It can be seen from this data that when the system is restricted to 5 or 10 work-orders the OTIF percentage is relatively low. Once the work-order limit is increased to 20 and above the OTIF begins to level out. It was found that the ECSyn work centre, reaches approx 100 work-orders using no buffer restrictions (see Figure 9), which when limited causes blockages throughout the rest of the model for numerous products and product families, thus affecting the overall OTIF results. However the experiment itself is useful in that it highlights the effect of increasing buffer sizes in such a situation. At some point the buffer size is not the problem but the capacity of the bottleneck work centre itself. This can be seen in Figure 9, where ECSyn is running with the buffer spaces all full from week 25 to the finish of the model thus causing blockages throughout the entire system.

Table 9: OTIF Percentages – Buffer Restrictions

	5 Work Orders	10 Work Orders	20 Work Orders	30 Work Orders	40 Work Orders
	OTIF%	OTIF%	OTIF%	OTIF%	OTIF%
A Rnd	43.95	70.34	78.21	79.62	73.84
A Seg	25.94	36.31	56.03	59.27	60.34
B Rnd	39.90	58.23	54.18	60.76	61.06
B Seg	25.16	47.45	46.17	45.96	45.53
C Rnd	55.32	90.35	90.82	90.82	90.82
C Seg	58.42	79.90	80.39	82.35	82.84
Totals	**36.80**	**55.86**	**66.80**	**69.16**	**67.34**

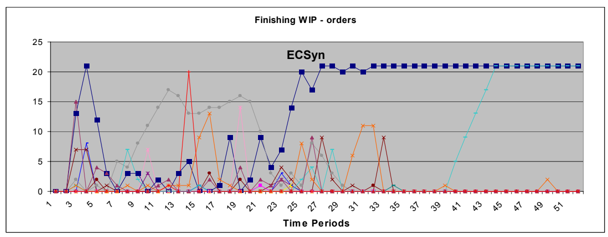

Figure 9: WIP Level - Buffer Restriction of 20 Work Orders

5.4.2 Average Days of Inventory and Finished Stock

Figure 10 presents a sample overview of the average days of inventory in finished stock for product families A, B and C, using the five different buffer restrictions. It can be seen that in general the average days of inventory across all products increase as the buffer restrictions increase from 5 to 10 work-orders in the finish processing area. However increases beyond this to the buffer limits produces little or no further effect, with the exception of Product Family A. It can be noted from the graphs that, in general there is a drop off in days of inventory after week 26, which corresponds with the information shown in Figure 9. Therefore, the finished stock inventory is being used to satisfy demand after this point without sufficient replenishment thus reducing the average days of inventory in finished stock from this point on.

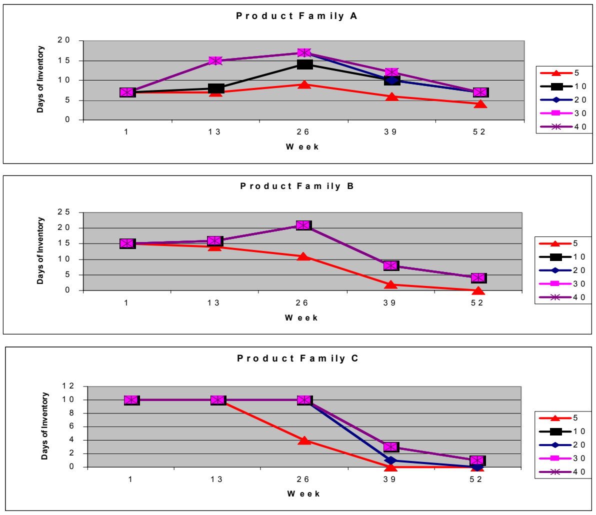

Figure 10: Average Days of Inventory in Finished Stock – Buffer Restrictions

Figure 11 presents information on the quantity of finished stock throughout the run of the simulation model. It can be seen from this graph that there is not a significant change in the quantity of stock held in finished stock as the buffer restrictions change, with the exception of the use of a buffer size of 5 work-orders. As was seen with the previous performance indices a significant drop in finished stock can be noted from week 26. This again can be attributed to an increase in demand at the ECSyn work center.

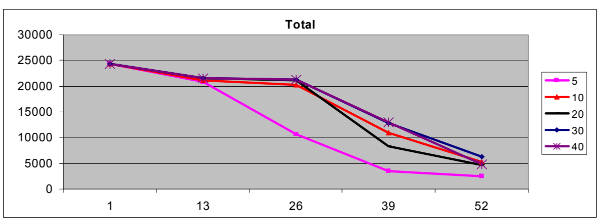

Figure 11: Total Finished Stock Qty – Buffer Restrictions

5.5 Experimental Group 2 – Production Controls (Part II)

After the first part of the simulation experiment was completed where the WIP was analysed for each work centre, it became clear that it was the ECSyn work centre that was experiencing the highest build up of WIP (see Figure 12). It can be seen that the rest of the work centres, which have not been identified individually on the graph have a much lower build up of WIP. The only other work centre with high levels of WIP is the FG work centre, which is the first centre in each parts product routing. Taking this information into account the second set of simulation experiments was run using the original configuration, with an additional 5 and then 10 machines in the ECSyn work centre.

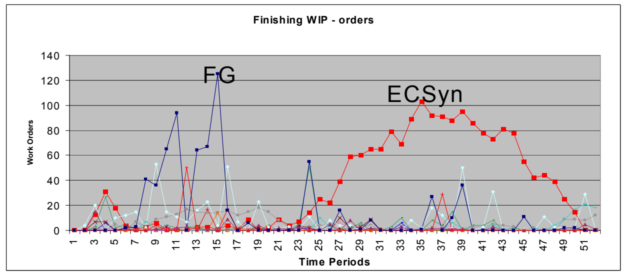

Figure 12: Build-Up of Work-Orders for YNNN Experiment

5.5.1 On Time In Full (OTIF%)

Table 10 presents a summary of the OTIF percentages using the same product families as before. It can be seen from this table that when an additional 5 machines (from 20 to 25 machines) are added to the ECSyn work-centre there is a significant increase in the OTIF percentage for all work orders that use that work centre, which includes the segment end of product families B, D, E and F (see Figure 13), but not on the other product families.

Table 10: OTIF Percentages – Increased Capacity

	Orig	Add 5 EC mcs	Add 10 EC mcs
B Rnd	84.0	84.0	84.5
B Seg	55.6	70.1	75.8
D Rnd	83.7	87.0	87.3
D Seg	63.2	84.3	90.4
E Rnd	92.3	93.7	94.5
E Seg	78.7	90.4	95.7
F Rnd	95.8	95.8	96.6
F Seg	47.7	68.2	93.2
Totals	**87.88%**	**89.40%**	**89.69%**

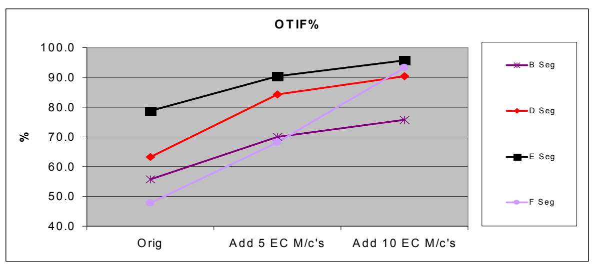

Figure 13: OTIF in Relation to Increased Capacity

5.5.2 Average Days of Inventory and Finished Stock

A sample of the average days of inventory in finished stock for each product family, under different capacity levels in the ECSyn work centre is presented in Figure 14, where product families B, C and D are examined graphically. It was noted that there was no change in the average days of inventory for the product families that don't use that work centre and only a small change in the product families that do use this work centre. It should also be noted however that as the capacity for the ECSyn work centre is increased the effect of increased demand around week 26 has a reduced effect on this KPI due to the extra resources available. For example product family C's average days of inventory fell from 21 to 10 days between weeks 26 and 39 due to this peak in demand using the original configuration. But with an increase of capacity by 5 and then 10 machines the fall was reduced from 21 to 15 and 20 days respectively.

It was seen from these simulation experiments that the number of machines in the ECSyn work centre has very little effect on the overall quantity of stock in finish stock, however there are small changes in the individual categories.

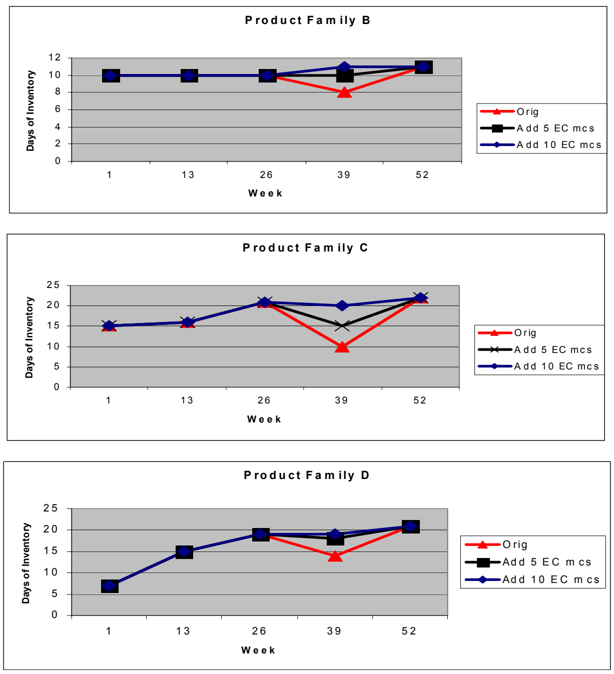

Figure 14: Average Days of Inventory in Finished Stock – Increased Capacity

5.6 Experiment 2 – Conclusions

This simulation experiment was carried out to review the effect of changing the models detail in relation to its production controls on the supply chain performance parameters. There are numerous experiments, which can be carried out in this area, for example production capacities, buffer restrictions, work order prioritisations, batch movement rules. So, for the purpose of illustration two such sets of experiments have been completed. The first set reviews the effect of introducing buffer restrictions in an attempt to minimise the level of WIP in the system. The second set was designed to review the effect of increased capacity in a bottleneck resource, as identified in the first part of the experiment.

After completing this experiment it was found that there was a large build up of WIP in a number of work centres in the model. In the finish processing area this was particularly evident in the ECSyn work centre. As outlined above the build of WIP in this centre was quite high. When compared with the rest of the model this work centre appears to be excessively high. On further analysis of this area it was found that there were certain discrepancies with the standard times (see validation in Byrne (2004)) that were being used and there was also a number of errors in relation to product routes going through the system, where work centres that no longer existed in the real system were being used in the input data to the model. Although there appears to be erroneous data in the model, this experiment is still useful in that it provides indications as to the trends created from production control alterations.

6 SUMMARY AND CONCLUSION

Simulation has only in recent years been applied to the analysis of real supply chains (Ingalls et al. (1999) and Lin et al. (2000)). For example, Lin et al. (2000) reported the application of the Asset Management Tool (AMT) to the analysis of a number of IBM's divisions. While AMT consists of a number of components, its core element is its simulation engine. There also exists dedicated supply chain simulation software, for example Simflex (Flextronics (2003)) and great success has been reported in the application of these tools.

The work carried out here adds to the body of knowledge in the supply chain simulation area in a number of different ways. The first is the analysis and development of a supply chain simulation model for a real vertically integrated system as opposed to the primarily logistical based models as found in the literature. With a vertically integrated structure the production elements of the system become more important and are analysed in this body of work. This is in contrast to the above studies where simulation was applied to systems with more simplified manufacturing steps and external suppliers and distributors.

The second topic is the implementation and application of supply chain simulation in SMEs. The division under analysis can be classed as a SME, which is in contrast to the real systems, which have currently been reviewed in the literature. IBM, HP, etc. have designed large-scale systems, which require a lot of resources and which is not a viable option for many SMEs. However, it is important that they have the ability to study their supply chains on some level in order to be able to compete in today's marketplace.

REFERENCES

Bagchi, S., S. J. Buckley, M. Ettl and G. Y. Lin 1998. Experience using the IBM supply chain simulator. *Pro-*

ceedings of the 1998 Winter Simulation Conference, ed. D.J. Medeiros, E.F. Watson, J.S. Carson and M.S. Manivannan. 1387-1394. Piscataway. New Jersey: Institute of Electrical and Electronics Engineers.

Byrne, P. J. 2004. Supply Chain Analysis of a Vertically Integrated Raw Material Supplier Using Simulation. *Manufacturing & Operations Engineering*, University of Limerick.

Byrne, P. J. and C. Heavey 2004. An Industrial Supply Chain Simulation Model. *Proceedings of International Simulation Conference 2004, Malaga*.

Christopher, M. 1998. *Logistics and Supply Chain Management, Strategies for reducing cost and improving service*. England, Financial Times Professional Limited.

Flextronics 2003. Available online via <`http://www.flextroics.com/simFlex/products.asp`> [accessed July 13, 2003].

Ingalls, R. and C. Kasales 1999. CSCAT: The COMPAQ supply chain analysis tool. *Proceedings of the 1999 Winter Simulation Conference*, ed. P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, 1201-1206. Piscataway. New Jersey: Institute of Electrical and Electronics Engineers.

Lee, H. L. and C. Billington 1993a. Material Management in Decentralized Supply Chains. *Operations Research* 41 (5): 835-847.

Lee, H. L. and C. Billington 1995. The Evolution of Supply-Chain-Management Models and Practice at Hewlett-Packard. *Interfaces* 25 (5): 42-63.

Lee, H. L., C. Billington, B. Carter and H. E. Edmondson 1993b. Hewlett-Packard Gains Control of Inventory and Service through Design for Localization. *Interfaces* 23 (4): 1-11.

Lin, G., M. Ettl, S. Buckley, S. Bagchi, D. D. Yao, B. L. Naccarto, R. Allan, K. Kim and L. Koenig 2000. Extended-enterprise supply-chain management at IBM personal systems group and other divisions. *Interfaces* 30 (1): 7-25.

Schary, P. B. and T. Skjøtt-Larsen 1995. *Managing the global supply chain*. Copenhagen, Handelshøjskolens forlag; Distribution, Munksgaard International.

AUTHOR BIOGRAPHIES

PJ BYRNE, Ph. D., is a Post Doctorate Fellow in the Department of Manufacturing and Operations Engineering in the University of Limerick, Ireland. He is a Production Engineering graduate from the University of Limerick and holds a Ph.D. from the same college. His main research interests are Simulation Modelling of Discrete Event Systems, Modelling and Analysis of Supply Chains and Manufacturing Systems His email and web addresses are <`pj.byrne@ul.ie`> and <`www.ul.ie/eerg`>.

CATHAL HEAVEY, Ph. D., is a Senior lecturer of Operations Management in the Department of Manufacturing and Operations Engineering at the University of Limerick. He is an Industrial Engineering graduate of the National University of Ireland (University College Galway) and holds a Ph.D. from the same college. His main research interests are decision support systems and simulation modelling. His email and web addresses are <`cathal.heavey@ul.ie`> and <`www.ul.ie/eerg`>.