Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

A SIMULATION BASED SYSTEM FOR ANALYSIS AND DESIGN OF PRODUCTION CONTROL SYSTEMS

Corinne MacDonald
Eldon Gunn

Dept. of Industrial Engineering
PO Box 1000
Dalhousie University
Halifax, NS B3J 2X4, CANADA

ABSTRACT

We present aspects of a simulated based system for analyzing and designing production control systems. The core of the system is a simulation of a manufacturing system operating with the Production Authorization Card system. The simulation model is fast and flexible, making it attractive for generating large datasets for use in developing simulation metamodels of expected performance for a wide variety of production configurations. Details of the simulation system are provided, along with a discussion of the issues to be considered when using it to design production control systems.

1 INTRODUCTION

Production systems are complex, usually involving stochastic processing times, product travel times, cells with parallel machines, machine failures, assembly cells, etc. Control strategies, such as kanban, CONWIP, and base stock systems, are used to control the flow of information and material in the system in order to achieve some desired performance levels while providing customer service. For a given manufacturing configuration, the selection of the best control strategy depends on the desired performance of the system. Is the goal reducing inventory? Minimizing cost? Increasing throughput? Minimizing customer wait time? Or (most likely) some combination of two or more of such measures?

Each of these strategies has associated control parameters that must be set in order to run the system. For example, a kanban system requires that a number of cards be allocated to each production cell in the system; the base stock system requires a target stock level for each inventory buffer between stations. Thus, there are several parameters for each type of strategy, and the number of parameters increases with the number of processing cells in the system. Even for small systems, finding the optimal number of kanban cards is a difficult problem.

The goal of this work is to develop a framework to select the best type of strategy for a given manufacturing configuration, and the parameters for that strategy. This presents a number of problems. One is how to simultaneously compare multiple types of strategies in the same model (i.e. kanban and base stock). A second is to develop a model to estimate performance of a manufacturing system (throughput, cycle time, etc.), given that this is dependent on many system characteristics, which vary widely in the literature. The final challenge is how to use the performance model to find the best strategy and parameters for that strategy when the input space is extremely large.

To address the first issue, we have used a modeling framework capable of representing different types of manufacturing control strategies in a single model. The Production Authorization Card (PAC) system created by Buzacott and Shanthikumar (1992) offers this flexibility. They have shown that the PAC system, which uses a combination of process tags and target inventory levels to control production, is capable of representing kanban, base stock, CONWIP, and other systems, including hybrids.

The second challenge is to estimate system performance. There has been much work on techniques which provide approximate deterministic performance functions (e.g. DiMascolo, Dallery and Frein 1996, Gershwin 1987). Buzacott and Shanthikumar (1992) developed just such a method using the PAC system for two cell serial systems with exponential processing times and reliable machines. The obvious advantage is that the resulting performance functions can be used to analyze various configurations, and deterministic optimization approaches may be used to search for the best parameter combinations. However, in order to develop these approximations, the complexity of the original problem must be restricted through the use of simplifying assumptions.

There has been groundbreaking work where simulation is used to estimate performance but often again with quite restrictive assumptions. For example, Bonvik, Couch and Gershwin (1997) use simulation to compare control strategies, assuming a sequential production process with lognormally distributed processing times, exponential failure and repair times, and a constant demand rate. Demand not met immediately from stock is lost. Gstettner and Kuhn (1996) compare kanban to CONWIP on serial production lines with exponential processing times and saturated demand (demand for finished goods is assumed to be unlimited). Gaury, Pierreval and Kleijnen (2000) apply simulation and evolutionary algorithms to a production system with no failures, and lognormally distributed processing times and inter-arrival times for demand.

Because we want to integrate as many real-world complexities into the system as possible, we opted to use simulation to develop the performance model rather than approximation approaches. However, the method for determining the best control system using simulation raises further questions. Two common approaches for finding optimal inputs to a simulation model are simulation optimization and simulation metamodelling.

In simulation optimization, multiple measures of performance are usually combined into a single objective function (Fu 2001). This is itself problematic in an environment where relative costs are difficult to estimate and often scale dependent. This is why researchers such as Bonvik, Couch, and Gershwin (1997) and Liberopoulos and Koukoumialos (2005) have focused on tradeoffs. Also problematic is how to do simulation optimization with a parameter space of rather high dimension. Olafsson and Kim (2002) discuss some of the issues.

Simulation metamodels are deterministic functions designed to approximate the expected value functions over the domain of the input variables. Regression models and neural networks are two commonly used approaches. A dataset of design points and corresponding estimates of system performance are generated by the simulation, and then a function fit to this data. Because there can be very large numbers (millions) of valid design points for the production control strategy problem, a carefully selected subset of these design points must be used for the dataset, one which adequately covers the state space. Our approach has been to use a space-filling experimental design method (Kleijnen et al., 2005).

Software packages such as ARENA are powerful modeling tools, but the time required to run each scenario made these impractical choices for generating large sample datasets. Using a high-level programming language to develop a simulation model requires a larger investment of time in model development, but provides the ability to generate large datasets much more quickly. However, modifying the model directly to simulate specific system configurations is time consuming, and not an attractive option. A model framework with the flexibility to change the entire problem with little overhead is needed.

To address these issues, we have developed a framework to study the production control strategy problem, the heart of which is a simulation system (PACSIM) which simulates the performance of a manufacturing system operating under the Production Authorization Card (PAC) System. The discrete-event simulation is based on a model originally presented in Bielunska-Perlikowski and Gunn (2002). The system was originally written in FORTRAN, a high-level language, based on the SIMLIB routines found in Law and Kelton (2000); the current model retains this choice of language. The model is designed so that the details of the manufacturing configuration are provided as input, and can be easily modified. Specifically, it allows the number of workstations, process time per station (by product, as multiple products may be produced by a single center), routings, travel times, failure and repair rates by station, and other characteristics to be easily reconfigured. Thus, a four station transfer line with exponential distributions can be simulated, and then, by changing the contents of a text file in a matter of minutes, the problem can be changed to an assembly system with uniformly distributed processing times and random machine failures. Given the configuration of the manufacturing system, and the control parameters of the PAC system under which it will operate, the simulation model provides observations of system performance such as system throughput, cycle time, average finished goods inventory, average WIP, customer fill rate and average customer delay time.

In terms of the work reported here, the most important feature of the simulation model is that it is fast. Depending on the size of the manufacturing system being modeled, the simulation model is capable of several replications per second. This makes the construction of large datasets for metamodel construction feasible.

In this paper, we discuss the framework for analyzing manufacturing control systems. Details on the Production Authorization Card system, the architecture of the simulation system, its use in the construction of metamodels and issues that must be addressed are presented in Section 2. The metamodelling approach, which is based on neural networks, is also described. This approach has provided good results, but we do not preclude other metamodelling mechanisms. The results of some experiments conducted using this framework are presented in Section 3; one of these is an illustration of the use of the constructed metamodels in a tradeoff curve analysis to identify the best control system for a specific manufacturing configuration. We close by discussing future upgrades.

2 SIMULATION SYSTEM DETAILS

2.1 Production Authorization Card (PAC) System

Production systems produce goods in several stages. A cell refers to either a single machine or a group of identical machines capable of producing parts or finished products.

Stores are the physical locations where these parts are kept until required elsewhere in the system. The PAC system (Buzacott and Shanthikumar 1992) is a token-based system where production at a cell can only be undertaken if the required components or materials are available, and the cell has been given authorization to do so. This authorization comes in the form of a process authorization (PA) card. Figure 1 is an illustration of two sequential cells in a system controlled by the PAC system.

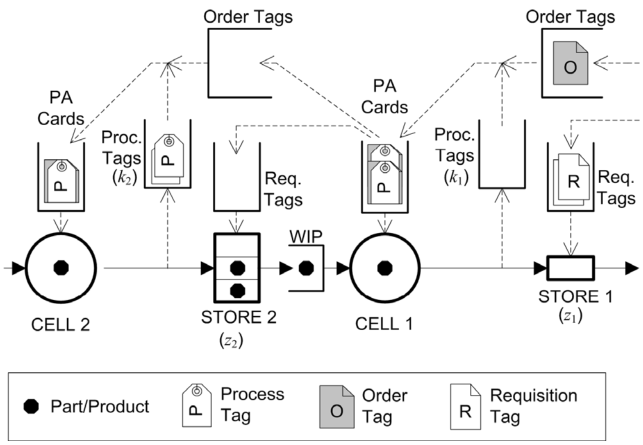

Figure 1: Schematic of the PAC System

When a cell receives a PA card, it immediately issues *order tags* to upstream stores for required parts produced at those cells. Customers may also generate orders to a finished goods inventory store. These order tags represent a notification that a part is going to be requisitioned from the store, and therefore production of a unit of this part should be started. Also kept at the store are *process tags*, which will limit the amount of jobs in production or in the production queue at the cell. If an order tag arrives and a process tag is available at the store, they are paired together to form a PA card; this authorizes the production cell to get the required components and start production on the part. If a process tag is not available, the order tag is placed in a queue to wait for the first available process tag.

Either at the same time it issued order tags, or after some delay, the cell also issues *requisition tags* for all required parts to the corresponding supplying store(s) (or to raw materials inventory). A requisition tag authorizes each store to send the part to the requesting production cell, at which time the requisition tag is destroyed. If there is no inventory in the store at the time the tag is received, the tag is placed in a queue which represents the backlog of demands for components at that store. Outstanding requisitions are filled immediately upon the arrival of inventory at the store.

If a PA card is received by a cell, and all required components have been delivered to the cell, then the job is ready for processing on the first available machine. Once processing is complete, the PA card is destroyed, and the completed part and the associated process tag are returned to the store. If more orders are waiting, the process tag is used immediately to create another PA card.

Processing cells may produce more than one product, and will receive separate PA cards for those products. It is assumed that a separate store exists for each product produced by the cell. Some may produce products that require more than one unit of a part, and/or more than one type of part; upon receipt of a PA card for such a product, the cell will immediately issue multiple orders and requisitions for all required parts, and processing cannot begin until all the required components are received.

The operation of this system is dependent on four PAC parameters at each cell/store combination:

- z_i, the initial inventory at store i
- k_i, the number of process tags at each store i
- τ_i, the delay time between the issuance of an order tag and the corresponding requisition tag
- r_i, the packet size for transmittal of PA Cards (represents batch production)

The flow of parts and information is therefore controlled by setting these parameters (the combination of which will be referred to from this point forward as the control policy, or design point). The initial inventory parameter, z_i, represents the WIP cap at the product store. Another way to view this is as the target inventory level; if the amount of inventory in this cell is less than the target inventory value, the supplying production cell will be working (unless it is waiting for required parts or has failed). The number of process tags at a store limit the number of PA cards that may exist at that store, thus limiting the number of jobs in process. If the batch parameter, r, is greater than one, then order tags are accumulated at a cell until a batch of r tags have been received. At that stage, the batch of r PA cards is immediately created and sent to the production cells. The final parameter, τ, allows order tags to be sent in advance of the requisitions, thus requesting the upstream cell to start work on replacing a part that is about to be removed. This parameter is used when working with advance demand information.

What makes this system so attractive is that it can be used to model several well known control systems such as Base Stock, CONWIP (Constant Work-in-Process), kanban, and hybrids of these systems. For example, CONWIP systems can be studied by setting the initial inventory level at the final store equal to the desired constant WIP level. The inventory levels at the other stores are set to zero, and the number of process tags at each station is set equal to this constant WIP level, as it is possible for all jobs to be in process at a one station. In a kanban system, the number of units of stock in a store is equal to the number of available kanban cards; in the PAC system, a kanban system is modeled by setting the number of process tags at each store equal to the initial inventory parameter at the store. Therefore if there are any units in inventory at the store there must, by definition, be process tags available to form PA cards. Base stock systems have no limit on the number of jobs in production at a cell, but do limit the size of the inventory buffer between stations; in the PAC system, this means the initial inventory parameter is set equal to the base stock level at the store, and a very large number of process tags are then allocated at the store (such that there are always tags available when orders arrive). For definitions of more systems, such as OPT and MRP, see Buzacott and Shanthikumar (1992).

2.2 PACSIM Simulation System

The PACSIM simulation system (Figure 2) enables the simulation of a wide variety of manufacturing system configurations operating under the PAC control system. It is programmed in Fortran and was compiled with Absoft 10.0 ProFortran. The system provides the ability to set up a model of a particular manufacturing configuration in a short amount of time, and the ability of the model to simulate many different policies (design points) very quickly.

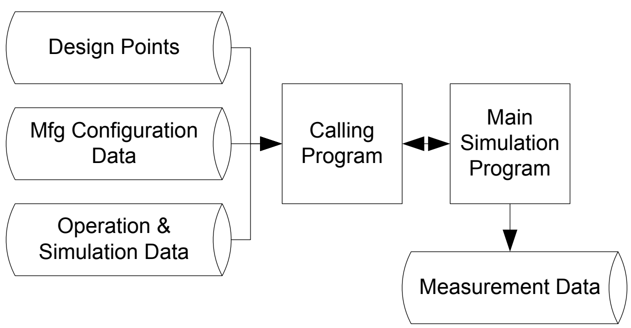

Figure 2: PACSIM Simulation System

The calling program reads all required information from the input files and calls the main simulation program to execute one replication of the simulation. The calling program determines the PAC parameters (design point) that will be used for each replication. These are provided to the program via a text file, which contains a list of design points to be simulated. In the case of a two cell sequential production line with no delays between orders and requisitions (so all τ values are zero) the text file would contain a set of points (z_1, k_1, r_1, z_2, k_2, r_2). We discuss the selection of these points later in this section.

Data on the configuration and operation of the manufacturing system is provided in a second text file. This file is generated by an Excel interface to allow easier data entry. First, information on the configuration is required:

- Number of processing cells
- Number of products produced (finished goods and component parts)
- Length of simulation run and warm-up period
- The number of machines in each cell
- Customer order arrival information (mean time between arrivals, type of distribution, additional distribution parameter if required)
- Failure and repair information (mean time between failures and repairs for each cell, type of distribution and additional parameter, if required).

The distributions for random variates currently available in PACSIM are normal, lognormal, exponential, Weibull, uniform, and geometric; they can also be designated as constant values. Processing time for each product at each cell is a random variable; the mean processing times for each product can be different, but also the distribution from which the random variate is drawn may be different for each product. If machine failures are permitted, then failure events occur randomly with a known failure time distribution. All machines in a single cell are assumed to have the same failure and repair distributions.

Products of the system are numbered starting with finished products, then sub-assemblies, and finally raw materials. Product 1 is always a finished product. The information required on these parts is entered as follows:

- PAC parameters (initial inventory, processing tags, batch size and delay time)
- Processing time data (Mean service time, distribution type and additional parameters, if required)
- Routing information (cell where product is made, the products and/or materials required and their quantities, and time to changeover a machine to produce this product if at a multi-product cell)

Where a single cell produces multiple products, it is assumed in PACSIM that if a changeover is necessary to produce a product, the time for this changeover is the same regardless of which product was last produced. If the cell has more than one machine, then the model first tries to find an available machine that was previously producing that product (thus to avoid a changeover). If one isn't available, it selects an available machine and the setup time is added to the processing time.

The third primary data file required contains additional information on the operation of the system and some simulation parameters. The most significant information contained in this file includes:

- Are reports, measurement files, or both required?
- Number of runs per design point
- Does the system respond to customer orders, or is customer demand assumed to be unlimited?
- Are unsatisfied demands backordered or lost?
- Are machines subject to failures?
- Type of priority rules at multi-item queues

If demand is assumed to be unlimited, then a new customer order is created every time a job leaves the system. This ensures the final workstation is kept busy at all times.

The main simulation program is built using the event-driven SIMLIB routines of Law and Kelton (2000). The event graph and other implementation details are contained in MacDonald (2006). In addition to the end of simulation event, there are eight events:

1. Arrival of an order at a store
2. Arrival of a requisition tag at a store
3. Arrival of a PA card at a cell
4. Arrival of WIP at a cell
5. Departure of a completed part from a cell
6. Arrival of product and process tag at a store
7. Machine failure
8. Completion of a repair of a failed machine

Entities in this system wait in queues, and the program tracks statistical information on the contents of queues. For example, the 'part' entity created by a production cell waits in a store (queue) until requested by a downstream cell or the customer. The program tracks the time-averaged contents of this queue, and the minimum and maximum contents. This information is tracked for all entity queues. In addition, several statistical variables are tracked and reported at the end of each simulation run. They include average delay time for customer requisitions, product cycle time, and average machine utilization, among others. Another routine tracks the average time between arrivals of finished goods at the final inventory stores. This information is used to determine whether or not the system has reached steady state.

The measurement files are text files for each performance measure reported by the system. The observations are written to the text files after every replication. In the current version of PACSIM, these include, among others:

- Percentage of customer requisitions met immediately from stock
- Average inventory levels in finished goods
- Average WIP levels for each part
- Average delay in meeting customer requisitions

These estimates of performance produced by the simulation model, combined with the dataset of design points used to generate them, are necessary to construct a simulation metamodel. However, there are several additional steps with must be dealt with in order to generate the dataset for the metamodel.

2.3 Generating Data for Metamodel Construction

Prior to the generation of the dataset necessary to construct a system metamodel, typical decisions such as the appropriate simulation run length and warm-up period in any simulation analysis must be made. However, there are other issues that require much more consideration.

2.3.1 Determination of Valid Design Points

The PAC system controls the flow of information and material in the production system. It is possible to select parameters such that the system cannot achieve a throughput equal to the demand rate. This occurs when the number of process tags is too small and little or no inventory is kept. Obviously these policies cannot implemented, so the minimum values for the PAC parameters must be determined through experimentation. First, a dataset with combinations of small parameters is generated, and then simulated. A program then reads these files and generates a report, which includes information on each design point, such as the resulting system throughput (to be compared with average order arrival rate), and whether or not the average customer delay time appears to be increasing in time. This report must be reviewed manually to determine the minimum allowable PAC parameters.

There are other constraints which can be applied to the selection of design points to reduce the size of the input space. Buzacott and Shanthikumar (1992) recognized that the number of process tags at an upstream cell does not need to be any larger than the maximum possible number of outstanding orders at the downstream cell. Any tags in excess of this amount would never be used. Any such design points should be eliminated from the input space.

When batching is permitted in a production environment, there are several constraints which must be met to ensure the system will operate properly. Clearly, the number of process tags must exceed the batch size, or a sufficient number of matched orders would never be reached. There is also the possibility of deadlocking (e.g. Wysk, Yang, and Joshi 1991), where all system resources have been seized by two or more processes, but none of the processes have seized enough of the resources needed to complete the current task, so that processing comes to a halt. In the PAC system, this can occur when batching is employed at two concurrent cells, and one cell is waiting for more orders to form a batch, while the other waits for more completed products to arrive and free up process tags so that more orders can be sent.

These are just some of the issues that must be considered to ensure all design points are feasible and not unnecessarily increasing the size of the input space. A much more complete discussion of allowable parameter combinations can be found in MacDonald and Gunn (2006).

2.3.2 Experimental Design

Our simulation model enables us to generate a large dataset quickly; therefore, when determining an appropriate experimental design for selecting design points, efficiency is no longer a primary concern (Kleijnen et al. 2005). However, since there may be several PAC parameters, a good space-filling design, where the design points are scattered throughout the input space with minimal unsampled regions (Cioppa, 2002), is employed. In our experiments, the limits of the input space must first be established, as well as the valid parameter combinations, before generating the design points. Our selection approach (MacDonald and Gunn 2007) is a factorial design where the factors are actually ranges of input data (low and high, or low, medium and high) such that the ranges cover the entire range of allowable values for each parameter; the factor combinations are then generated, and then the parameter values are randomly drawn from the appropriate range. This is a modified version of the approach used by Hurrion (1997).

The generation of the observations for the dataset raises an interesting question of how to expend simulation effort. One alternative is generate a large number of design points and simulate each just once. Another is to use a smaller number of design points and replicate each point several times to develop more accurate point estimates. Our experience to date is that if we use a space-filling design, we can build accurate metamodels by generating a large set of design points and then simulating each long enough to ensure an unbiased sample of steady state. This question requires further research.

3 APPLICATIONS OF THE MODEL

3.1 Comparing Pull-type Control Systems

A problem well discussed in the literature is the direct comparison of different types of traditional strategies to determine which is best for a particular type of system. For example, Duri, DiMascolo, and Frein (2000) compare kanban, base stock and the generalized kanban policy, which was originally introduced by Buzacott (1989) and is actually the first version of PAC (without batching or time delays). Therefore, all three of these strategies can be modeled with the PAC system. The authors use an approximation method to predict performance of a manufacturing system with processing stations in series, a Poisson arrival process for demand, exponentially distributed processing times and reliable machines. They then use this approximation to find the parameters for each of these three control systems to minimize a cost function of finished goods and work in process inventory, subject to a minimum constraint on customer service probabilities. Each strategy must be evaluated with different models. They then find the optimal by enumerating almost all valid policies to find the one with the minimum cost.

In a similar experiment, a three station configuration (Figure 2) was analysed to find the best kanban, CONWIP, and general PAC policies. The details of this work may be found in MacDonald (2006). The first station of this line produces two parts, each required by a different production cell. Orders arrive according to a Poisson process, with different arrival rates for the two products. Processing at the three cells has a Weibull distribution. The first cell produces two parts, each one required by a different downstream cell; PA cards are processed on a first come, first served basis. There is time associated with moving units between stations. These few complications preclude the use of approximation methods for performance evaluation.

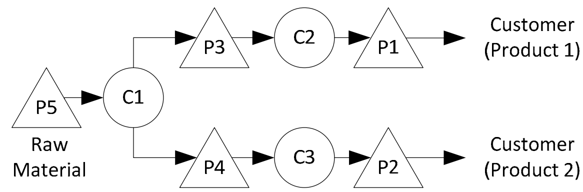

Figure 3: A three station configuration

We assumed that batching of orders going to the first cell was permitted; therefore, a total of ten PAC parameters are required. Through the experimentation process discussed earlier, the constraints on the input space were determined, after which it was determined that the input space consisted of over 800,000,000 valid design points. A separate program generated potential design points using a 3-factorial type design (as discussed in Section 2.3), and then eliminated any combination that did not meet the constraints. As a result, 6,086 design points were generated and then simulated with PACSIM (at a rate of 7 replications/second on a 2.0 GHz PC). We then used neural networks to create metamodels for the design points and each measure of performance. Rather than optimizing a single performance function, we developed optimal policy curves (Starr and Miller 1962) for each type of strategy. We used a simulated annealing algorithm applied to the metamodel to find the policies with the best combinations of total average inventory and customer service percentage. The resulting curves are shown in Figure 4. The general PAC policies were, of course, the best, since both CONWIP and kanban are restricted versions of PAC control; however, the CONWIP policies were very close to optimal, likely easier to implement (Hopp and Spearman 2001), and superior to kanban policies for this example.

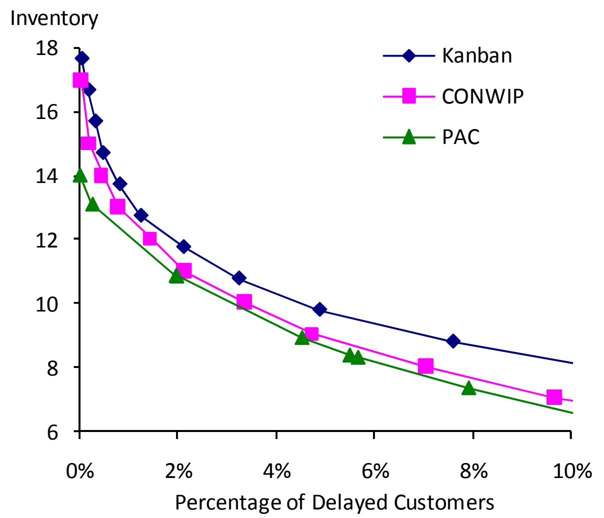

Figure 4: Optimal policy curves for three control strategies

3.2 Simulating Transfer Lines

In order to demonstrate the flexibility of the PAC methodology, it is worth considering transfer lines. A transfer line is a special type of production line where all parts flow through the same series of sequential processing steps. It is assumed that in this type of system there is a limit to the size of the buffer (store) between each successive processing station. If a cell completes processing of a part only to find the buffer in which it should be placed is full, the part must be kept at the cell until a space becomes available, and further processing is halted. Hence the cell becomes blocked. A cell will become starved if it runs out of jobs to process. Exact analytical models of system performance do not exist for all but the simplest of cases. There is extensive literature on developing approximation models for estimating system performance, but fewer papers on determining the best allocation of buffer spaces (e.g. Gershwin and Schor 2000).

This type of system can be studied using PACSIM. Buzacott and Shanthikumar (1992) refer to the use of buffers to control product flow as Local Control. To model this control system, the initial inventory parameter of the final store (finished goods) is set to zero, the initial inventory parameter of all other stores is set to one more than the buffer size (to account for the extra space at the machine), and all process tags are set equal to one (so that no material is moved out of an upstream store unless the machine is ready to process it). This type of system does not respond to customer demand, but instead produces as much as possible, and it is assumed that final products immediately leave the system. As described earlier, the program will create a customer order at the beginning of the simulation, and then every time this order is satisfied (a product is completed and leaves the system) the model will immediately generate a new customer order, so that the final cell will continue processing as long as there are components available to process. Since there is only ever one customer order outstanding, the average customer delay time will represent the average time between final product departures. Therefore the inverse of this measure will be the average system throughput. Dallery, David, and Xie (1989) discuss some transfer line examples, with deterministic processing times at each single-machine station. The machines are subject to random failures; the mean time between failures and the mean time between repairs are exponentially distributed. They use a simulation model of a transfer line to compare to their approximation approach for finding the average throughput and work in process inventory in each buffer. One example presented is a four station transfer line with buffer sizes 20, 0 and 20. The throughput (TH) and average buffer contents (B) for both the approximation method and a simulation model are shown in Table 1, along with the corresponding average results of 25 replications using PACSIM. The PACSIM results are in close correspondence with the Dallery, David and Xie simulation but are produced with modest computational effort.

Table 1: Performance results of a 4-station transfer line

	TH	B_1	B_2	B_3
DDX – Approx.	0.430	14.9	0	5.1
DDX – Simulation	0.431	15.0	0	5.0
PACSIM	0.432	14.9	0	5.1

4 FUTURE UPGRADES AND OPPORTUNITIES

PACSIM provides the ability to study manufacturing systems of moderate complexity. To further develop our approach to designing appropriate manufacturing control strategies for complex manufacturing systems, and to improve the understanding of the effects of such strategies on the various measures of system performance, we would like to integrate more complexities into PACSIM, and continue to examine various types of problems using this system.

4.1 Upgrades to the Simulation Model

Additional functionality identified for PACSIM, to enable the study of a wider variety of manufacturing systems, is outlined below:

- Instead of assuming that parts are processed one at a time, the model should be expanded to allow for machines that produce parts simultaneously in batches, where the batches need not always be the same size. For example, in a heat treatment operation, several parts can be treated simultaneously.
- Issues with respect to imperfect quality should be included. In the current model, it is assumed that

all produced parts are good. However, there are many ways that quality issues can impact performance; through the addition of inspection processes, the assumption that a random number of parts are defective, the addition of rework, etc.

- Material movement time is currently part of the model, but it is assumed that the time is known with certainty, and that parts are moved immediately. This is rarely true. The model should allow for transport vehicles that can move more than one part at a time (but may have limited capacity), that may arrive at random intervals.
- At cells where more than one product are produced, additional priority rules (besides FIFO and oldest job first) should be included.
- Adding the ability to treat certain system characteristics (e.g. priority rules, average processing times) as inputs to the model in the same way the PAC parameter are now. This would enable experiments whereby not only the control strategy can be changed, but also the system itself.

4.2 Opportunities for Future Research

With some recent upgrades to the model, and the additional functionality discussed above, we foresee several opportunities for future research.

One opportunity is to study tandem production lines with 10 or more productions stations, and different production time distributions and machine failures, to gain further insight into systems such as Kanban, CONWIP, etc. on system performance. There are several variants of this type of analysis to be explored – saturated (unlimited demand for finished goods) versus unsaturated (responding to random customer demand) systems, systems with batching versus without, and lost sales versus backordering.

We have not yet studied the effect of advance demand information (e.g. Liberopoulos and Koukoumialos 2005, Krishnamurthy and Claudio 2005) on PAC controlled systems. This presents at least two opportunities for future work. The first would involve determining a method to compare MRP type systems using the PAC scheme, including imperfect forecasts and the use of the delay parameter, with pull type systems where the delay parameter is not employed and the system typically responds to actual demand; an extension of this would be some understanding of how accurate a forecast needs to be so that systems responding to the forecast would perform better than a pull type system. The second would involve the analysis of an MRP system specifically, and whether or not such a system would see an improvement in performance if limits on the number of available process tags were placed on some or all of the production cells. One criticism of MRP is that jobs are released into the system without regard for the number of jobs already in the system (Hopp and Spearman 2001) therefore, a limit on information flow could reduce average work in process inventory and provide a lower, and perhaps less variable, average lead time for jobs. Both of these issues should be investigated.

We also see a need to further investigate the deadlocking issue mentioned earlier, and develop rules on the assignment of parameters to avoid this situation. Some of these situations been identified in Buzacott and Shanthikumar (1992), and we have identified others through analytical techniques and rules to prevent their occurrence (MacDonald and Gunn 2007); however, for more complex systems, especially in the cases of multi-product machines and assembly systems, automatic detection mechanisms should be developed.

5 CONCLUSION

Developing optimal control strategies for production systems is a complex problem. The PAC system enables the study of different configurations and different control systems in a single framework. PACSIM enables us to configure and simulate moderately complex manufacturing systems operating with PAC control. In order to develop the simulation metamodels required for optimization and tradeoff analysis, a very fast, highly flexible model that can rapidly evaluate a large number alternate control parameter settings is required. PACSIM has proven to meet many of these needs.

We are continuing to develop the system to allow for the analysis of more complex manufacturing systems. As discussed in Section 4, there are a number of important configurations found in the literature, and we continue to add these modeling opportunities to PACSIM. Current developments are aimed at enabling the study of systems with advance demand information, varying batch sizes, different priority rules, less than perfect yield, and random material movement times.

REFERENCES

Bielunska-Perlikowski, K., and E.A. Gunn. 2002. Structure of a Simulation Model for a Manufacturing System Coordinated by Production Authorization Cards. *Journal of Design and Manufacturing Automation* 1(4):231-245.

Bonvik, A.M., C.E. Couch, S.B. Gershwin. 1997. A Comparison of Production-line Control Mechanisms. *International Journal of Production Research* 35(3): 789-804.

Buzacott, J. A. 1989. Queueing Models of Kanban and MRP Controlled Production Systems. *Engineering Costs and Production Economics* 17:3-20.

Buzacott, J. A., and J. G. Shanthikumar. 1992. A General Approach for Coordinating Production in Multiple-

Cell Manufacturing Systems. *Production and Operations Management* 1(1):34-52.

Cioppa, T. 2002. *Efficient Nearly Orthogonal and Space-Filling Experimental for High-Dimensional Complex Models*. Ph.D. Thesis, Naval Postgraduate School, Monteray, CA.

Dallery, Y., R. David, and X. L. Xie. 1989. Approximate Analysis of Transfer Lines with Unreliable Machines and Finite Buffers. *IEEE Transactions on Automatic Control* 34(9): 943-953.

Di Mascolo, M., Y. Frein, Y. Dallery. 1996. An Analytical Method for Performance Evaluation of Kanban Controlled Production Systems. *Operations Research* 44 (1):50-64.

Fu, M. C. 2001. Simulation Optimization. In *Proceeding of the 2001 Winter Simulation Conference*, ed. B.A. Peters, J.S. Smith, D.J. Medeiros, and M.W. Rohrer, 53-61. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Gaury, E.G.A., J.P.C. Kleijnen, and H. Pierreval. 2001. A Methodology to Customize Pull Control Systems. *Journal of the Operational Research Society* 52:789-799.

Gershwin, S. 1987. An Efficient Decomposition Method for the Approximate Evaluation of Tandem Queues with Finite Storage Space and Blocking. *Operations Research* 35(2), 291-305.

Gershwin, S., and J. E. Schor. 2000. Efficient Algorithms for Buffer Space Allocation. *Annals of Operations Research* 93:117-144.

Gstettner, S. and H. Kuhn. 1996. Analysis of production control systems kanban and CONWIP. *International Journal of Production Research* 34: 3253-3273.

Hopp, W. and M. L. Spearman. 2001. *Factory Physics*, 2nd edition. Boston: McGraw-Hill.

Hurrion, R. D. 1997. An example of simulation optimisation using a neural network metamodel: Finding the optimum number of kanbans in a manufacturing system. *Journal of the Operational Research Society* 48: 1105-1112.

Kleijnen, J. P. C., S. M. Sanchez, T. W. Lucas, and T. M. Cioppa, 2005. A User's Guide to the Brave New World of Designing Simulation Experiments. *INFORMS Journal on Computing* 17 (3):263-289.

Krishnamurthy, A., and D. Claudio. 2005. Pull Systems with Advance Demand Information. In *Proceeding of the 2005 Winter Simulation Conference*, ed. M.E. Kuhl, N.M. Steiger, F.B. Armstrong, J.A. Joines, 1733-1742. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Law, A. M., and W. D. Kelton. 2000. *Simulation modeling & analysis.* 3rd ed. New York: McGraw-Hill, Inc.

Liberopoulos, G., and S. Koukoumialos. 2005. Tradeoffs between base stock levels, numbers of kanbans, and planned supply lead times in production/inventory systems with advance demand information. *International Journal of Production Economics* 96:213-223.

MacDonald, C. 2006. *A Framework for the Analysis of Manufacturing Systems Controlled by a Production Authorization Card Scheme*. Ph.D. Thesis, Dalhousie University, Halifax, Nova Scotia, Canada.

MacDonald, C., and E. Gunn. 2006. *Determination of Valid Policies for the Production Authorization Card (PAC) System*. Working Paper #06-06, Department of Industrial Engineering, Dalhousie University, Halifax, Nova Scotia, Canada.

MacDonald, C., and E. Gunn. 2007. Analyzing manufacturing systems using neural network metamodels. In *Proceedings of the Industrial Engineering Research Conference*, ed. G. Bayraksan, W. Lin, Y. Son, and R. Wysk, Nashville, TN.

Ólafsson, S., and J. Kim. 2002. Simulation Optimization. In *Proceeding of the 2002 Winter Simulation Conference*, ed. E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 79-84. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Starr, M.K., D.W. Miller. 1962. *Inventory Control: Theory and Practice*. Englewood Cliffs, New Jersey: Prentice-Hall.

Wysk, R.A., N. S. Yang, and S. Joshi. 1991. Detection of Deadlocks in Flexible Manufacturing Cells. *IEEE Transactions on Robotics and Automation* 7(6): 853-859.

AUTHOR BIOGRAPHIES

CORINNE MACDONALD is an Assistant Professor in the Department of Industrial Engineering, Dalhousie University. She received a BEng in Industrial Engineering from the Technical University of Nova Scotia and a PhD in Industrial Engineering from Dalhousie. Dr. MacDonald's research and teaching interests include modeling and design of production systems, work design, simulation and metamodeling. She is a member of CORS, INFORMS, a Senior Member of IIE, and a registered professional engineer.

ELDON GUNN is a Professor of Industrial Engineering at Dalhousie University. He received a BSc from Mount Allison University, an MASc from Dalhousie University, and a PhD in Industrial Engineering from the University of Toronto. Dr. Gunn's research interests involve the application of operations research models in a variety of settings, ranging from natural resources to manufacturing. He is a former President of the Canadian Operational Research Society and a Fellow of the IIE.