

## PERFORMANCE ANALYSIS OF A CONFIGURED TO ORDER BUSINESS WITH A VARIABLE PRODUCT CONFIGURATION RECIPE

Soumyadip Ghosh  
Tom Ervolina  
Young M. Lee

IBM T.J. Watson Research Center  
Yorktown Heights, NY 10598 U.S.A.

Barun Gupta

IBM Integrated Supply Chain Group  
Hopewell Junction, NY 12533 U.S.A.

### ABSTRACT

We study a complex Configured to Order (CTO) business operating under variable product configuration recipes and high component commonality within product Bills of Materials. In particular, we conduct a series of simulation experiments to compare certain supply chain designs and investigate how their performance changes with changing variability of product configurations. The results indicate that under higher variability and component commonality conditions, a supply chain design that lays more emphasis on tracking the evolution of component demands and inventory positions performs better.

### 1 INTRODUCTION

Manufacturers of specialized products faced with tighter competition and falling margins have been placing an increasing emphasis in recent times on creating a wider portfolio of products as a means of providing better value to customers and boosting profit margins. A broader spectrum of specialized products often lead to higher market shares (Bagozzi 1986), but it may increase the operational cost and complexity of the manufacturer's business. A traditional response to these higher costs is to increase the number of common-use components in the product Bills of Materials (exploiting *component commonality*; Swaminathan and Tayur 1998, Jayaraman et al. 1998) and consolidating by having fewer product families, to realize economies of scale advantages. The product families typically allow customers more product configuration options in order to achieve a wider product portfolio.

The effect of such a consolidation with more common components on business performance may not be transparent in the presence of changing variability in the product configurations. Product demand uncertainty and configuration variability may interact in a manner as to render the demand for common components highly variable and hence suspect to inventory stock-outs and customer dissatisfaction.

These interactions typically tend to be complex and hard to model analytically.

We experiment with various designs of complex Configured to Order (CTO) businesses in such an environment by estimating their performance via stochastic simulations. Our work is motivated by a study of server products manufactured by IBM where server families tend to be highly configurable and are assembled from hundreds of components. A product Bill of Materials from this line is typically multi-tiered, but we shall consider a simplified product structure with only two tiers: the finished goods and the components that they are assembled from.

The entire supply chain for this business consists of various stages including demand forecasting, procurement, order processing and fulfilment processes; see Section 2. Business policies followed at each stage determine how the process executes and responds to changing environmental conditions. For instance, the policy followed at the demand forecasting stage determines how future projected demands for finished goods and their components are obtained from available order history information. Each process is also associated with an environment within which it functions. For instance, environmental assumptions for the supply planning phase specifies the stochastic behaviour of the suppliers, their constraints and capacities, etc.

We will be interested in evaluating the effect of control policies put in place at three stages of this supply chain (Section 2). We call a combination of policies chosen at each of these stages a *policy regime*. The management seeks primarily to determine a good supply chain design (a policy regime for the supply chain) that can respond favorably to changes in the variability of product recipes while maintaining a diverse product portfolio. A good control regime lets the marketing department offer customers a wider range of products via more latitude in configuring products with a good serviceability assurance.

Our goal in this work is to take a holistic approach of developing insights into how the complete supply chain performs under various control regimes and reacts to chang-

ing product profiles variabilities. Analytical methods are unfortunately limited to tackling only an aspect of the entire design problem. The approaches presented by Swaminathan and Tayur (1998) and Jayaraman et al. (1998) address how inventory costs and customer dissatisfaction can be minimized by an appropriate choice of common components. They determine the optimal choices of pre-configured sets of common components (they call these sets vanilla boxes) and their associated inventory parameters that minimize operational costs while providing better overall service to customers than a standard mad-to-order system. Song and Zipkin (2003) provide a good discussion on the progress being made in developing tools to analyze models that incorporate configuration variability as a consequence of correlated product demands. Lu and Song (2005) provide a recent perspective on this issue; they seek to find good approximations to optimal component inventory stocking policies for a system with common components and moderately variable product designs.

The key novelty of our simulation-based method is that it can quantify the dependence of the supply chain performance on the individual policies of each stage and the integrated process through which they work. This thus overcomes the difficulty of estimating the system performance by separately analyzing each sub-process in isolation.

Section 3 describes the simulation experiments performed to evaluate these policies. Performance of the supply chain designs were evaluated using measures deemed most important by the management. Highly specialized and differentiated markets like the server business typically see fewer customers than a more commoditized one and hence maintaining good customer relationships is paramount. Customers request when they would like to receive the product that they order and are promised an order delivery date based on the requested date and manufacturer's projected availability of finished goods and components. Central to good performance is the ability to keep these promises and deliver the orders on time. Section 3 discusses this and other measures in detail.

Results of this study are presented and discussed in Section 4. We compare four policy regimes and find that the more the policy pays attention to tracking the demand and inventory availability information of components as it evolved over time, the better it performs when product recipes are highly varying. Section 5 summarizes the insights we obtain from this study and suggests certain important future enhancements to this model that should provide more value to this study.

## 2 THE SUPPLY CHAIN MODEL

The IBM server Configured to Order (CTO) supply chain can be abstracted to a supply chain model represented in Figure 1. The model consists of four main stages: demand

planning, supply planning, configuration planning and order scheduling. Each of these stages are governed by various policies that can be reconfigured as per the current supply chain design being studied, and work within a known or preset business environment. Our aim is to determine the effectiveness of the combination of policies being employed.

**Demand Planning:** This stage produces projected future sales quantities for the products. The demand forecast is based on the trend observed in the past business transaction data and is typically modeled in weekly buckets over a planning horizon. Our study assumes throughout that this horizon is a quarter year (three months or 13 weeks) in length. This forecast demand information can be at the finished goods level or the components level. The policy within this module sets a flag that indicates forecast requirement at the finished products level or components level.

Negotiations with suppliers in supply planning has to be carried out at the components level, which would thus require planning component forecasts. The first sub-process of the supply planning stage, the MRP process in Figure 1, creates projected demands for components from either of two methods. The first derives these from the forecasts of the products by via an explosion on their Bill of Materials (BoM). This step would indeed preserve the accuracy of the product forecasts in the components forecasts in a situation where the Bill of Materials is deterministic (each product needs exactly a known number of components). However, in a CTO world where end customers are allowed to vary the composition of the finished good, product-component relationships are stochastic in nature.

Historic customer ordering data can be used in the demand planning phase to gather information on anticipated usage of specific components in terms of product-component composition distributions. This information can, for instance, provide (fractional) usage rates called *attach rates* of components to products. These are essentially the mean number of components needed per product order, but can be conditioned by other business considerations like risk, etc. Component forecasts can be obtained in a straightforward manner by combining attach rate information with the generated product forecasts.

This procedure may not necessarily construct component forecasts that are as accurate as the product forecasts they are derived from since the attach rates used are themselves forecast parameters. The accuracy of any component that is common to many different products might be particularly suspect to this effect. For such components, an alternate procedure generates their projected demand from the historic data available on their usage in filling product orders. Ideally one would like to use this procedure for every component in the system, but this may not be viable since CTO systems tend to have a large total number of components. A practical alternative is to generate indepen-

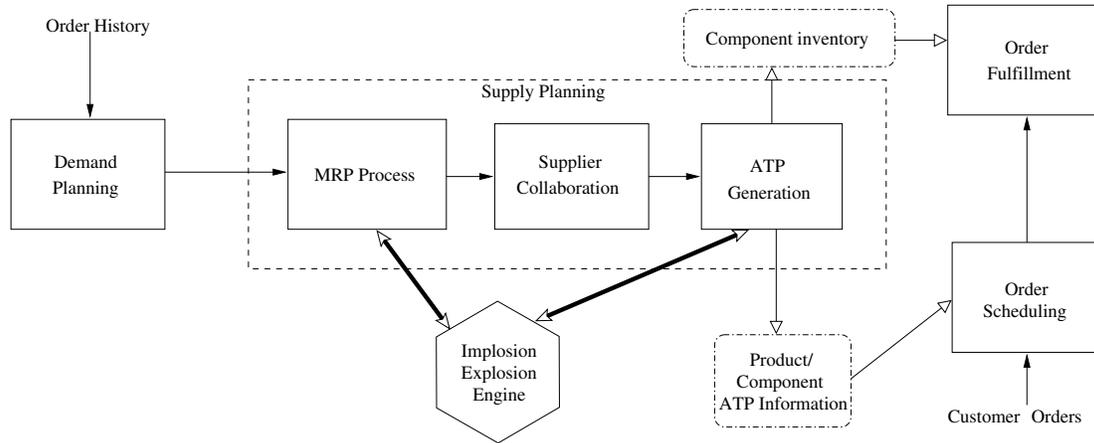


Figure 1: The Supply Chain Model

dent forecasts for an appropriately chosen set of components and derive the forecasts of the other components from the product forecasts in the MRP process.

The possible policy alternatives to be compared then are

- D1** forecasting product demands and attach rates for each relevant component-product pairing, and
- D2** forecasting product demands and demands for a set of components independent from the product forecasts, and attached rates for the rest.

**Supply Planning:** This stage encompasses three sub-processes that are common to most CTO manufacturing environments: Material Requirements Planning (MRP) process, Supplier Collaboration process, and Available to Promise (ATP) Generation process. The supply planning process is executed by invoking these sub-processes in sequence.

The *MRP process* translates the product forecast information to component-level projected demands of the entire horizon of thirteen weeks. The required quantities of components are either obtained from component forecasts when these are generated independently, or otherwise computed by performing an explosion on the Bill of Materials of the products. The MRP Process calls an Explosion Engine where the attach rates are set as in the demand planning stage. Various policies can be applied to the MRP sub-process depending on the business conditions.

The *Supplier Collaboration* process transmits the component demand quantities to the suppliers and receives a supply commitment from them to the demand request. It is typical that for at least one of the components, the supplier will not be able to commit the full amount requested. This condition leads to a component shortfall which subsequently means the supply chain is constrained and will not be able to meet the demand forecast as planned.

Historic data suggests that supplier constraints are typically tight in the early part of the horizon, but suppliers

can be flexible farther out in the horizon, in that they can plan for and meet larger component requirements if needed by the manufacturer to fulfill delayed orders. The supplier collaboration business environment describes the flexibility, responsiveness, and capacities of the suppliers, and the terms, conditions, and timing of the communication process between manufacturer and suppliers. An appropriate environmental assumption can be to use a randomized supplier commitment where every component supply commit has a uniform variability around the current request with no memory of previous commitments.

The supply commitment for the current week corresponds to the actual component replenishment supplied. The *ATP Generation* process seeks to determine, from the supply commitment for the future weeks in the horizon, the best possible availability of end-products towards satisfying the demand forecast. This information plays an important role in ensuring that customer orders are scheduled and handled in a timely fashion. ATP Generation is a complex optimization problem and typically requires an Implosion Engine. The ATP Generation sub-process can take into consideration many attributes that include whether exact optimal or approximate heuristic solutions should be generated, factors like costs, revenue, penalties for delayed availability and partial shipments, risk factors, demand priorities, supplier sourcing preferences and costs, penalties for using substitute parts, and assumptions about flexibility in the attach rates. Each possible (reasonable) configuration of attributes leads to a different implosion procedure. We use a procedure that generates approximate solutions from heuristics while ignoring all costs and revenues, setting equal priorities for all end-product demands, sourcing from single suppliers, applying no penalty for substitute parts and delayed availability, and allowing no flexibility in attach rates. This is implemented in an implosion/explosion tool named the Supply Capability Engine.

Many components that go into a product can be considered optional for the purpose of the implosion for a variety of reasons. For instance, the supply of these components are known to be non-constraining in the sense that they may either be externally sourced in a crunch, or other similar components may be substituted at a moderate cost. Thus, the finished goods ATP the Implosion Engine calculates can be optimistic. However, the availability of some components might need be tracked along with those of the finished goods. This might for instance be because of the critical nature of the component to the final product (a server blade to the final server), or because it is a commonly used by many products (hard drives). Our system's ATP generation process considers the following policies where ATP information is tracked for

- A1** only finished goods, conditioned on the supply of all components,
- A2** only components, and
- A3** a chosen set of components and finished goods, conditioned only on the supply of components that are not chosen.

Components ATPs are the same as the supply commitments elicited for the components. We shall choose the set of independently tracked components to be the same as the set of components forecast independently in the Demand Planning stage.

**Order Scheduling:** Customers request an order delivery date that is typically days to weeks in the future of the order placement date. This might be because these expensive purchases require advance planning and budgeting on their behalf, and some products might have significant production lead times. Each customer order is processed and an order ship date is scheduled based on requested delivery date and the expected availability of products or components. Orders can be checked against various combinations of components and finished goods availability. We shall compare the following scheduling policies where order requirements are matched against:

- O1** finished goods ATP only,
- O2** components ATP only, and
- O3** finished goods ATP and the ATP quantities of a chosen set of components.

This stage's policy determines factors that influence the promised order ship date, for instance the relative priorities given to the product. When an order is scheduled against the ATP, the required quantity of the product and/or components are reserved for the particular order so that other future customer order cannot use this availability.

**System Dynamics:** In the IBM server business we study, the demand and supply planning steps are repeated every two weeks. Thus forecasts and supplier commitments are re-calibrated, and ATP information is placed in sync with supplier commitments subject to requirements of orders that have been scheduled and await fulfillment. The system maintains actual inventories at the components

level and orders are fulfilled from these when they are due. Component inventory replenishment orders might be placed at the beginning of each week and are fulfilled instantaneously, presumably because suppliers have produced a batch of components in accordance with demand forecast information available to them in the past. The orders may be only partially filled because of the constraints faced by supplier for that week. In these settings, replenishment order sizes are typically set such that they cover the expected demand for the component in the week ahead while taking into account factors like backorders and risk associated with stock-outs.

Order assembly and fulfillment processes are assumed to be instantaneous once all the required components can be reserved from inventory, though in practice a small (nearly deterministic) lead time may be involved. The fulfillment lead time can be factored into the customer's requested order ship date to present a more accurate model.

A constrained supply of components implies that the component inventory may not correspond exactly to the calculated ATP information. Other factors like the stochastic nature of demand and order configurations contribute to this mismatch. As a consequence, some orders cannot be fulfilled by the promised date and have to be rescheduled to use the fresh supply of components in the next week for its unfilled portion. Our chief interest is in determining the incidence of such cases.

### 3 SIMULATION STUDY OF MODEL

A good policy regime for the supply chain should allow the manufacturer to offer a broad range of product choices while simultaneously aiming to minimize the number of disruptions of on-time order fulfillments and the cost incurred by the manufacturer. We propose to use stochastic simulation to obtain estimates of performance of designs for the supply chain model described in Section 2. In performing these experiments, we seek specifically to compare the performance of four different policy regimes listed in Table 1.

Table 1: Policy Regimes Compared

Policy Regime	Demand Planning	ATP Generation	Order Scheduling
I	D1	A1	O3
II	D1	A1	O1
III	D1	A2	O2
IV	D2	A3	O3

Regime II corresponds to a situation where each stage of the supply chain is executed at the finished goods level, while Regime III represents a system where orders are scheduled and processed entirely at a components level. Regimes I and IV take an intermediary approach of scheduling by

checking ATPs of finished goods and re-checking ATPs of certain important components. Regime I however maintains demand planning and ATP generation at the finished goods level exclusively, while Regime IV goes through every step in the supply chain with a set of components that chosen to be independent of the finished goods information and the rest depend on finished goods.

The simulation experiment was originally conceived and conducted with an aim towards gaining insight into IBM server product lines when governed by the policy regimes listed in Table 1 under varying business conditions. These insights are perhaps better described here in the context of artificially generated data that allow clear understanding and control over the business conditions. We consider a manufacturing system that offers a product portfolio with 5 products that are constituted from 30 components. A component is considered independent if it is present in three or more products' Bill of Materials; we are thus setting highly common components aside to be forecast and/or tracked separately.

A business may be interested in evaluating the impact of various alternatives of product structures on their present supply chain with an aim towards increasing the breadth of their offerings while maintaining a high supply chain performance. We shall study four alternate scenarios that we create by changing two factors in the product structure. The first factor deals with component commonality (refer Jayaraman et al. 1998 and the reference therein) by moving from isolated product Bills of Materials towards incorporating more common parts in them. We consider two different sets of Bills of Materials, one with a lower component overlap than the other. Table 2 lists the two product compositions, specifying the index of commonality of components (the number of components that appear in one, and higher, product Bills of Materials) for each. The product Bills of Materials are distinct from each other within each scenario, for instance the number 3—common components in the Bill of Materials might vary with the products.

Table 2: Components Common to Products

Commonality Scenario	Number of components common to products				
	1	2	3	4	5
Low	15	15	5		
High	10	10	5	5	

The second business factor compares configuration options given to customers. A product requires at most one of each component listed in its Bill of Materials. We model the distribution of the component requirements with independent Bernoullis with the same probability  $p_c$  that a quantity of one is needed. In practice, significant dependencies can be expected between the component distributions of each

product (and over all products), but we feel the assumption of independence is reasonable for our purposes here. We let  $p_c = 0.95$  and  $p_c = 0.5$  represent the low and high product configuration variance scenarios respectively. This implies an increase in configuration variability by a factor of about 5.

We shall test the performance of the four chosen policy regimes under a total of four scenarios. The IBM supply chain managers considered following performance metrics important:

- PM1 **Order Scheduling Rate:** This is defined as percentage of customer orders that are assigned and communicated scheduled ship dates. Component supply constraints might very adversely affect scheduling of certain orders by inducing very long delays from the customer requested delivery date. In such cases, managers often find that scheduling and fulfilling such orders in an expedited manner outside the regular supply chain may be the best in terms of maintaining good customer relations.
- PM2 **On-Time Fulfillment Rate:** This is the percentage of scheduled customer orders that are fulfilled on the date they were promised on. An inability in fulfilling an order on its promised ship date reflects very poorly on the supply chain and its managers, and hence maintaining a high fulfillment rate is considered crucial.
- PM3 **Order Fulfillment Delay:** This depicts the percentage of fulfilled orders that suffered a fulfillment delay (in weeks) from its scheduled ship date of more than one week, which in practice is considered quite intolerable. The full delay statistics is characterized by a distribution on integral weekly delay values.
- PM4 **Total Component Inventory:** This metric provides an idea of the cost incurred by the supply chain. We count the inventory held for all components including those that are set aside for partially filled delayed orders.

We use an IBM software offering, the WBI Modeler and Simulator, to run the simulation experiments, namely coordinating the generations of events and movements of information entity such as customer orders through various stages of the model. The simulator runs the model through a duration of simulated time, which was set to 20 quarters in this study.

The simulation model requires certain other assumptions to completely define the stochastic behaviour of the supply chain. Supplier responses are modelled as a normal random variable, with the current request as the mean and with a uniform variability around it for each week of the horizon, bounded above by the requested quantity. Supplier

flexibility farther out in the horizon can be modelled with an appropriately higher mean and variance.

Customer orders are assumed to arrive according to a Poisson Process with a mean interarrival time  $\lambda$  consistent with the forecasts generated in the demand planning stage. Each arriving order can require one of the five possible products with equal probability. In the test case results presented in this paper, the product-component recipes for each of these five products are fixed by the “High” or “Low” business scenario from Table 2. Other interesting scenarios with a mixture of recipes from both these scenarios can also be studied under this framework.

Orders that are placed at time  $t$  have to be fulfilled on or after a customer specified time in the future  $t + Y$ . The random delay  $Y$  is modeled by the discrete distribution  $G$  given in Table 3. All random variables in the simulation are generated independent of each other.

Table 3: Requested Ship Date Delay Distribution

Delay (in days)	5	10	15	20
Probability	0.55	0.10	0.15	0.20

Product forecasts have to be generated every two weeks in the demand planning phase to match the product requirements of this compound Poisson process. Component forecasts and related inventory replenishment decisions have to be made at the MRP sub-process of the supply planning phase, and the presence of the non-zero  $Y$  requires careful consideration especially in setting an initial inventory in system for a simulation of the system. In steady state, conditions under the assumptions the expected demand for components in each week can be derived from the Poisson process mean and the product-component distributions. But a simulation cannot be started off with this same inventory requirements in each week since many of the jobs that arrive by the first week will not need fulfillment by that week. This would lead to inventory accumulation in the system in the first few weeks, the effects of which might take longer to wash off.

Fortunately, the stochastic process of order requested ship dates (RSD) yields itself to straightforward analysis to determine the correct component inventory requirements. Note that the RSD process can be thought of as the output process of an  $M/G/\infty$  queue. This is a queue with an infinite number of servers that sees Poisson arrivals, where each new arrival is processed by a servers for an amount of time drawn from the distribution  $G$  introduced in Table 3. Let  $N(\cdot)$  represent the output of this  $M/G/\infty$  queue. So  $N(t) = \{\text{Number of departures from this } M/G/\infty \text{ queue till time } t\}$ . Ross (1996), on pg. 70, shows that the process  $N(\cdot)$  is also a Poisson process with mean  $\lambda p(t)$ , where

$$p(t) = \int_0^t G(t - s) ds = \int_0^t G(y) dy. \quad (1)$$

Here  $p(t)$  essentially represents the average fraction of customer orders that have arrived in the system and reached their requested ship date within time  $t$ .

Equation (1) shows how the evolution of the RSD process can be followed to choose appropriate inventory replenishment order values. For instance, the simulation can be started with orders for components to satisfy  $\lambda p(t_7)$  orders in the first week,  $\lambda(p(t_{14}) - p(t_7))$  in the second week and so on. But this might involve many cumbersome evaluations of the  $p(\cdot)$  function in (1), and in this case an easier approach can be used.

Equation (1) also shows that  $p(t) \rightarrow 1$  as  $t \rightarrow \infty$ . Moreover, the rate of this convergence is exponential in  $t$ , and the convergence is faster with a lighter tail for the distribution  $G$  (Whitt 2002). In other words,  $M/G/\infty$  queues tend to approach steady-state behaviour fast if the  $G$  distribution is well-behaved, as our chosen  $G$  is. In steady-state, the rate of output of this  $M/G/\infty$  is equal to the rate of input  $\lambda$ . Thus, one can allow this  $M/G/\infty$ -queue based RSD process to run first for a few weeks of simulation time and then start the simulation of the entire supply chain system. The MRP process bases the inventory replenishment orders on the required order fulfillments from this starting point, which should be close to the steady state average of  $\lambda$  per our assumptions.

#### 4 RESULTS AND DISCUSSION

Table 4 provides a comprehensive summary of the results obtained from the simulation experiments. The results of all the policy regimes are grouped together for each business scenario. The table does not provide confidence intervals for these estimates in order to keep the presentation digestible, but we note that in all cases the variances of these estimates are within acceptable ranges.

Scenario 1 sees the regimes perform in a manner indistinguishable from each other, and this can be expected since finished goods and component requirements are almost deterministic factors of each other. Table 4 does clearly show that variability in the product configurations has a much larger effect on the performance of the supply chain system than does component commonality. In other words, the effect of product heterogeneity and, consequently, component commonality is muted, and this could be because the scenarios generated from Table 2 are not sufficiently differentiated from each other.

The second performance metric measures the percentage of orders that were delivered at the promised date, and is often considered most important in terms of maintaining good customer relationships. It shows remarkable changes for Regimes I, II and IV as the configuration variance is increased. Regime III is not as badly affected, and this can be attributed to its tracking the inventory positions of each

Table 4: Estimated Performance of Supply Chain Model under Various Scenarios and Policy Regimes

	Policy Regimes			
	I	II	III	IV
Scenario 1: Low Comp. Comm., Low Config Var.				
PM1 (%)	96.00	96.19	95.99	96.13
PM2 (%)	94.20	94.80	97.93	95.25
PM3 (%)	.41	.20	.03	.21
PM4 (units)	663.60	669.56	672.90	670.90
Scenario 2: Low Comp. Comm., High Config Var.				
PM1 (%)	96.80	98.60	96.01	97.20
PM2 (%)	73.83	62.28	95.34	74.67
PM3 (%)	3.77	3.42	0.08	2.68
PM4 (units)	453.04	482.59	394.18	487.16
Scenario 3: High Comp. Comm., Low Config Var.				
PM1 (%)	95.96	96.03	95.92	96.51
PM2 (%)	95.56	94.30	97.58	96.25
PM3 (%)	.35	.37	.04	.27
PM4 (units)	830.44	823.60	821.71	831.27
Scenario 4: High Comp. Comm., High Config Var.				
PM1 (%)	95.91	98.77	94.94	96.12
PM2 (%)	79.75	57.76	96.35	79.69
PM3 (%)	2.75	5.20	.09	2.92
PM4 (units)	598.19	621.14	570.14	793.26

component individually rather than taking an aggregated view at the finished goods level as the other regimes do.

Comparing Scenarios 1 and 3 shows that while an increase in component commonality does not have much effect on its own when the variability is held low, it does compound the effect of high configuration variability (Scenarios 2 and 3) on system performance.

Overall, it is reasonable to conclude that within a given scenario, the performance numbers of Regime III are often more favorable than the others. This is not very noticeable in the low product configuration variability scenarios, but the advantage is distinctive under high variability. While Regime II does manage to schedule more orders within the system (PM1) and thus requires less external expediting, its schedules are often overly optimistic (PM2). This can be attributed to the fact that Implosions in Regime I (and Regimes II and IV) provide an optimistic view of future finished goods capacity, while Regime III, which checks each components availability, schedules conservatively. It however performs in an exemplary fashion when called on keeping its promised delivery dates, and this behaviour is consistent over all scenarios, while Regime II for instance can see a huge drop in its ability to keep its promises as the configuration variability increases.

The inventory numbers in Table 4 cannot be compared across various scenarios since the average quantity of each component needed per order changes with the scenario. Within each scenario, Regime II has a better overall average

inventory carrying requirements, and this is especially true under high uncertainty.

Regimes I and IV clock in somewhere between the Regimes II and III; hence, they can be considered as a good balance between these. Regime III might be cumbersome in implementation since it needs the capability to track the evolution of each component separately, a daunting task in a real world system with thousands of components. Thus, Regime I or IV might prove to strike a better balance between performance and ease of implementation.

## 5 CONCLUSIONS AND ENHANCEMENTS

Based on the preliminary evidence provided in Section 4, a policy regime that tracks demand and availability information of each component within the product portfolio of the supply chain system as it evolves over time provides the most consistent performance when faced with changing variabilities in product compositions in a CTO supply chain system as described in Section 2. Practical implementation considerations might often indicate that it would be better to implement a regime that tracks only some of the more important components while simultaneously tracking forecasts and availability at a finished goods level.

We emphasize here that this study is very preliminary in nature in that several important factors need to be included to make the model more realistic. Of primary important is the fact that the model described in this article leaves out all cost factors. Inclusion of cost information might change the behaviour of the model in significant ways. First, an explicit cost metric will provide a single scalar scale to compare the performance of each of these regimes on, which helps better understand the value of the trade-offs each policy regime engages in.

Second, knowledge of costs of each component and its marginal value to each product may help better determine the set of components that need be tracked independently in Regimes I and IV, thus possibly bringing their performance closer to Regime III in the high variability cases. This will also help determine the true cost of carrying inventory within each system.

Cost and revenue information on the products might help device better ATP generation, order scheduling and fulfillment procedures that could conceivably bridge the performance gap between Regimes II and III. This can be done for instance by introducing relative priorities between the products ordered. Thus higher priority orders might be better satisfied, albeit at a possible loss to the lower priority customers.

Orders often tend to arrive more frequently towards the end of a business quarter influenced by many factors, for instance budgetary cycles of customers, and thus tend to be more skewed than a homogeneous Poisson process. An order arrival process that recognizes various time-of-week

and -month effects might be a better suited model of real world arrival processes.

Lastly, the scenarios chosen from Table 2 may not have exhibited sufficient product homogeneity, which might account for the fact that the effect of component commonality on the results in Table 4 is muted. The conclusions of this study can be strengthened by considering more heterogeneous product recipes.

## REFERENCES

- Bagozzi, R. P. 1986. *Principles of marketing management*. Chicago, IL: Science Research Associates.
- Jayaraman, R., J. Rappold, R. Roundy, R. Srinivasan, and S. R. Tayur. 1998. Component procurement under uncertainty. Technical Report 1227, School of O.R. and I.E., Cornell University.
- Lu, Y., and J. S. Song. 2005. Order-based cost optimization in assemble-to-order systems. *Operations Research* 53:151–169.
- Ross, S. M. 1996. *Stochastic processes*. 2nd ed. New York: Wiley.
- Song, J. S., and P. Zipkin. 2003. Supply chain operations: Assemble-to-order systems. In *Supply Chain Management*, ed. T. De Kok and S. Graves, Volume 30 of *Handbooks in Operations Research and Management Science*. North-Holland, Amsterdam: Elsevier Science.
- Swaminathan, J. M., and S. R. Tayur. 1998. Managing broader product lines through delayed differentiation using vanilla boxes. *Management Science* 44:163–172.
- Whitt, W. 2002. *Stochastic process limits*. New York: Springer-Verlag.

## AUTHOR BIOGRAPHIES

**SOUMYADIP GHOSH** is a Research Staff Member of the Mathematical Sciences Department at IBM T.J. Watson Research Center, Yorktown Heights, NY. His research interests encompass simulation theory and practice. He can be contacted at <ghoshs@us.ibm.com>.

**TOM ERVOLINA** is a Research Staff Member of the Mathematical Sciences Department at IBM T.J. Watson Research Center, Yorktown Heights, NY. His research interests include invention, design, and development of algorithms and software for supply chain optimization, particularly in the areas of supply/demand rationalization and Available to Promise (ATP) information generation. He can be contacted at <ervolina@us.ibm.com>.

**BARUN GUPTA** is a Senior Technical Staff Member at the IBM Integrated Supply Chain Group, IBM, Hopewell Junction, NY. His area of specialization includes end-to-end integrated demand and supply planning processes and

applications. He responds to queries at <barun@us.ibm.com>.

**YOUNG M. LEE** has been working in the mathematical science department of the IBM T.J. Watson Research Center since 2002 in the areas of supply chain simulation and optimization. Prior to joining IBM, he had worked for BASF Corporation for 14 years, where he had founded and managed the Mathematical Modeling Group, and led development of numerous optimization and simulation models for various logistics and manufacturing processes. He has extensive experience in chemical, pharmaceutical, process industries, and managed and modeled various manufacturing and logistics processes using mathematical programming, artificial intelligence and simulation. Recently at IBM, he developed several complex simulation models that are instrumental in analyzing and improving business processes, supply chain and IT solutions. He has a B.S., a M.S., and a Ph.D. degree in Chemical Engineering from Columbia University, City of New York. His research interest includes simulation and optimization of supply chain, manufacturing and business processes. His email address is <ymlee@us.ibm.com>.