

# Optimizing the Profitability of Hewlett-Packard's Supply Chains

---

COREY BILLINGTON	<i>Vice President, Global Procurement Services, Hewlett-Packard Company</i>
GIANPAOLO CALLIONI	<i>Director, Strategic Planning &amp; Modeling, Hewlett-Packard Company</i>
BARRETT CRANE	<i>Project Manager, Digital Imaging Hewlett-Packard Company</i>
JOHN D. RUARK	<i>Chief Technology Officer Optiant, Inc.</i>
JULIE UNRUH RAPP	<i>Senior Business Analyst, Inkjet Supplies Hewlett-Packard Company</i>
TRACE WHITE	<i>Senior Business Analyst, Inkjet Supplies Hewlett-Packard Company</i>
SEAN P. WILLEMS	<i>Assistant Professor of Operations Management, Boston University, School of Management</i>

---

A typical supply chain modeling and analysis exercise has always been a complex undertaking, typically done by a few experts over many months. However, recent advances in supply chain inventory optimization have radically simplified and accelerated supply chain design and configuration within the Hewlett-Packard Company. These techniques, when coupled with a user-friendly front-end, allow a critical business ability that was formerly the province of the few to be deployed into the hands of the many and have saved the company over \$133 million for two product lines in less than two years.

## 1 Introduction

In 1939, Hewlett-Packard introduced their first product, an audio oscillator, which was sold to Walt Disney Studios for creating some of the special sound effects for the movie "Fantasia". In May of 2002, Hewlett-Packard merged with Compaq, with the goal of becoming the world's leading provider of IT solutions. By the end of 2002, Hewlett-Packard's revenues were \$72.3 billion. Today, Hewlett-Packard operates in over 170 countries, and is number one in fault-tolerant servers, Win-

dows servers, Linux servers, storage systems, management software, personal computers, pocket PCs, inkjet printers, multi-function printers, laserjet printers, flatbed scanners, printing services, as well as being number three in IT services.

In order to maintain its position as an industry leader, HP must keep pace with constant advances in technology in a highly competitive marketplace. To do this, HP introduces literally thousands of new products every year. The older units lose value so quickly, and the competitive market is so

fierce, that it is imperative to properly design, configure, and optimize the supply chain wherever possible.

The result of this is that cost factors such as material devaluation, scrap costs, write-offs, and fire sale discounts have become the single biggest detriment to profitability, sometimes even eclipsing the already-slender profit margins. Calculating these *inventory-driven costs* requires the consideration of service levels, demand and supplier uncertainty, and end-to-end process times across the entire supply chain. The stochastic nature of many of these factors makes this a very challenging problem.

---

*Inventory-driven costs are the single biggest detriment to profitability in the competitive computing and electronics industries.*

---

Up until now, HP tackled the problem of first determining, and then minimizing, inventory-driven costs through a combination of homegrown spreadsheets coupled with the modeling expertise of HP's internal Strategic Planning and Modeling group (SPaM). It was commonplace for managers and planners from different HP organizations to conduct supply-chain analyses on their own by building models in Excel. These employees would then share their respective models with each other, and spend considerable time attempting to influence their peers across the supply chain to buy into *their* analysis.

With 141,000 HP employees spread across 178 countries, all of whom coordinate with a vast global network of suppliers, it is easy to see that the ability to make good supply-chain analysis decisions cannot continue to rest in the hands of a few experts. For some time, HP had been searching for a methodology that would allow business units to conduct their own end-to-end supply chain analyses. What HP needed was an easy-to-use business application to provide non-PhD business users with detailed insights into the interaction of key supply-chain parameters through a controlled pro-

cess. This application would need to be broadly distributable to users in all business units and regions, so that all the users would have access to the same set of tools.

In 2000, HP began working with a Boston-based company called Optiant to bring HP's long-standing tradition of supply-chain innovation to a new and broader level. Optiant's product, called PowerChain Inventory, combined a Web-based architecture with a simple user interface. The interface was wrapped around a sophisticated optimization engine that incorporated some of HP's processes and supply-chain analysis techniques.

This paper describes both the process HP developed for determining total supply chain cost, and the supply-chain theory that was developed to solve these real-world supply-chain problems. To illustrate how broadly this technology and the underlying process has been deployed at HP, we present two radically different business cases: digital cameras and inkjet cartridges. For just these two product lines alone, the tool has helped HP reduce total supply chain costs by over \$130 million while maintaining already-high service levels.

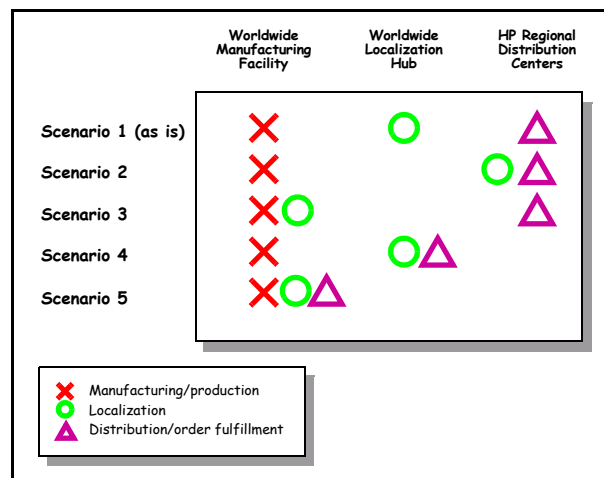
## 2 Optimizing the Digital Camera Business

Digital imaging is a key growth area for Hewlett-Packard. HP's Digital Imaging organization offers a wide array of imaging products. Consumer offerings include digital cameras, flatbed scanners, and photo inkjet printers. In the commercial arena, HP's products include Phogenix photo processing equipment and Indigo digital presses. The Digital Imaging group develops complete imaging solutions to its customers – everything from image capture to image output, and all of the software and infrastructure connecting the two.

HP entered the digital camera market in 1997, with a single camera. Today, HP offers over half a dozen cameras, ranging from aggressively priced auto-focus cameras to high optical zoom, high-megapixel products. The digital camera business is extremely competitive; HP must contend against Sony, Kodak, Nikon, Canon, and Olympus. Prod-

uct life cycles are commonly less than a year, with prices rapidly declining over this period.

In the fall of 2000, various individuals within the digital camera organization were aware that HP's digital camera supply chain was not cost-effective. In a bottom-up initiative, they formed a cross-functional team to identify ways to improve the situation. Led by supply-chain engineers, this effort was supported by resources from finance, manufacturing engineering, R&D, and logistics. After considering potential points of change, the team identified five scenarios to consider, as shown in Figure 1.



**Figure 1 Digital Camera Supply Chain Configurations**

The five scenarios considered various levels of consolidation and localization between a single worldwide manufacturing facility, a single worldwide localization hub, and several regional distribution centers.

Scenario 1 was business-as-usual. In the existing supply chain, each digital camera flowed from a dedicated factory in Asia that air-shipped to a single worldwide hub, also in Asia. Localization occurred at this hub. Regional distribution centers received build-to-stock shipments by air from the worldwide hub. From the regional DCs, the final product shipped by ground to customers.

Scenario 2 eliminated the Asia hub facility, instead opting to perform postponement at each regional distribution center. Scenario 3 also eliminated the hub, but moved localization back to the factory. Scenario 4 cut out the regional distribution centers and the subsequent ground transit phase

between those centers and the retailers. Instead, products were postponed and then shipped directly from the Asia hub by air.

Scenario 5 represented the highest level of consolidation, with a single worldwide factory performing all manufacturing and localization, and then shipping the product by air directly to retailers in the regions.

## 2.1 Determining Supply Chain Cost

In order to understand which of these configurations to adopt, the team needed to calculate the total supply chain cost for each of these configurations. HP's costing method divides the total supply chain cost into variable, fixed, and inventory-driven costs. Variable costs consist of the value added at each stage in the manufacturing process, including per-unit transit costs, direct manufacturing costs and labor costs. Fixed costs measure the cost associated with a particular supply chain configuration over a specified time frame, and include items such as the portion of a distribution center's costs that are allocated or charged against the particular product line. Inventory-driven costs are dictated by the cost of the product, holding cost rate, service level, supply variability, and demand uncertainty.

Aggregate fixed and variable costs are relatively straightforward to calculate. For fixed costs, one needs to know the cost of the fixed assets. To compute variable costs, one has to know the volume going through the facility. The uncertainty surrounding inventory-driven costs creates the dual problem that these costs are both difficult to calculate and to forecast. In particular, one year ahead in the budgeting process it is impossible to know for certain which products will be high demand or low demand, which components will be delayed, and which suppliers will have quality problems. However, these uncertainties are some of the causes for holding inventory and they are a large driver of inventory-driven cost.

The most difficult part of optimizing the inventory-driven costs in the supply chain involves identifying the location and size of buffer inventory at each point in the supply chain. Too much

inventory of the wrong type, or at the wrong place in the chain, can increase obsolescence costs; on the other hand, having too little of the models in demand can result in allocation conflicts and lost sales.

Within the digital camera worldwide supply chain, the team modeled 44 potential inventory locations. Calculating a globally optimal solution with this many variables in an Excel spreadsheet was simply not possible.

The team turned to the PowerChain application specifically to solve this inventory optimization problem and to quantify inventory-driven costs. Each of the five scenarios required the devel-

opment of a supply chain model similar to the example shown in Figure 2.

This illustration shows the supply chain map that corresponds to Scenario 1 in Figure 1. Each box on the supply chain map represents a value-added step such as product transformation or transport. A triangle within a box denotes safety stock held at that stage. PowerChain was able to optimize inventory levels and locations across the multiechelon network – which would have been next to impossible without some advanced modeling. Furthermore, the supply chain structure shown in Figure 2 was one of the first models that HP built using PowerChain, and is among the most simple to date.

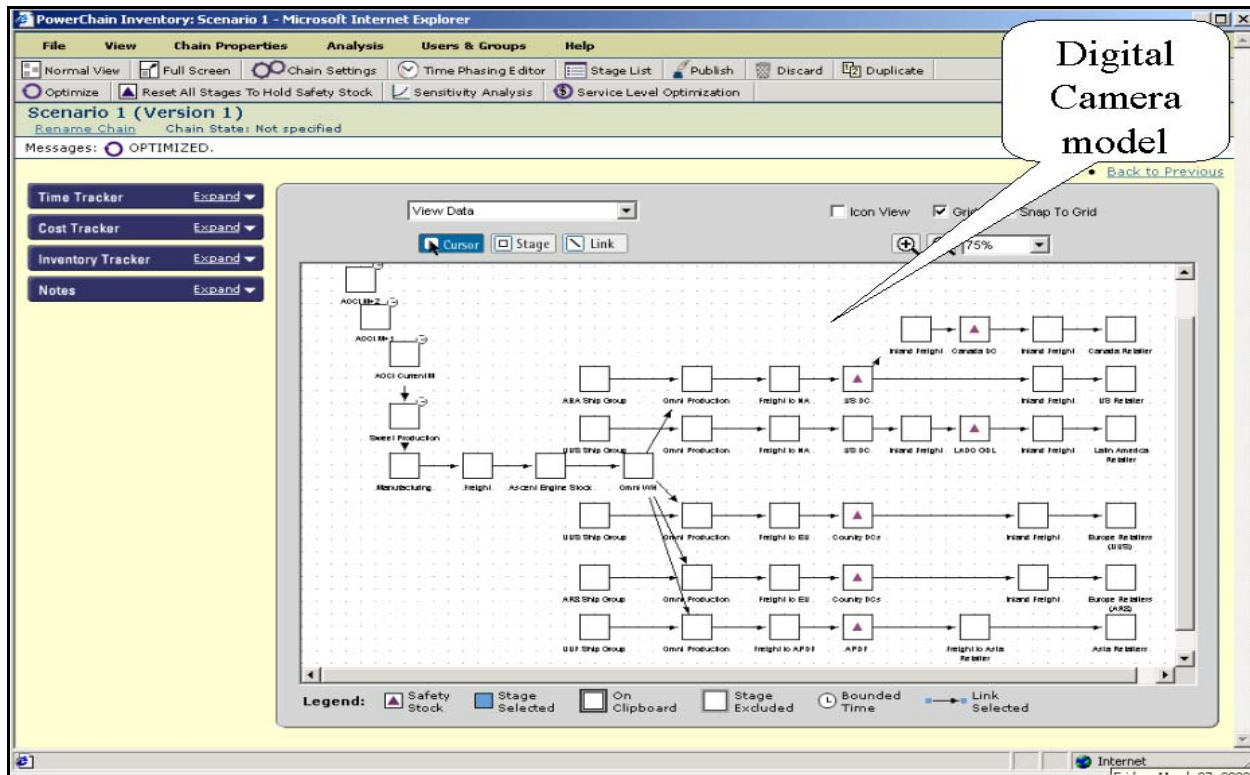



Figure 2 PowerChain Depiction of the Scenario 1 Digital Camera Supply Chain

This figure shows an example map for one of the five digital camera scenarios listed in Figure 1. This type of model differs from many other types of supply chain modeling tools in that it is process-focused rather than location-focused. In this diagram, each box represents some sort of product transformation: either a product change or a location change. Sometimes an entire vertical column of boxes can represent a single location with different product configurations, or a transit phase.

The five scenarios were built, optimized for inventory levels and locations, and then compared with each other to identify which scenario could potentially deliver the lowest total cost. Note that the team was not modeling for the least amount of inventory; rather, they were modeling for least cost. The reason is that inventory might have different per-unit costs depending on where it is in the supply chain. For example, a partially completed but un-localized unit might be cheaper to buffer than separate FGI stockpiles. Thus it is possible that a scenario might suggest holding more inventory at a particular point, and still result in lower cost.



**Figure 3 Efficient Frontier**

This illustration compares the total supply chain costs for the five digital camera scenarios, expressed as the efficient frontier for each scenario.

## 2.2 Findings

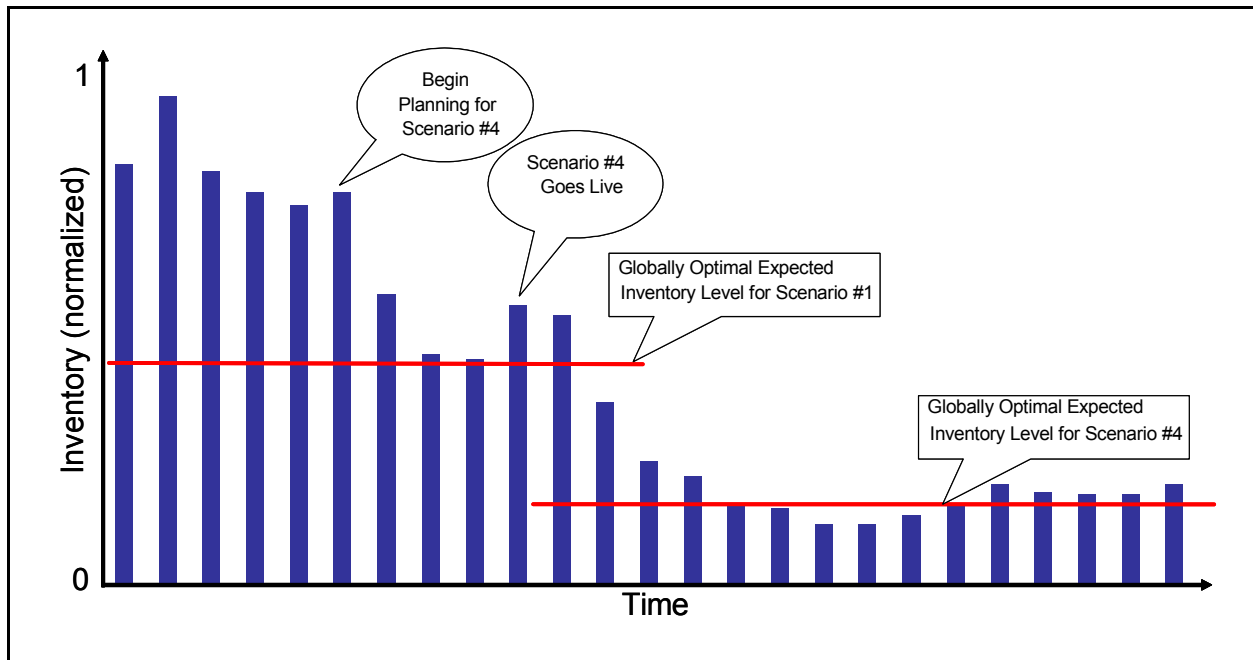
For each supply chain configuration, the cost was described by a point along an “efficient frontier.” As shown in [Figure 3](#), the fourth scenario (ship directly from Asia hub) has the lowest overall cost. Adopting this scenario achieves two benefits. First, we can create a step-function improvement in supply chain performance – which translates directly into cost – by restructuring the supply chain. Even with no other business process improvements, we benefit from the more efficient

structure provided by Scenario 4. Second, we now know the optimal inventory policy for our global supply chain, and can implement it through changes to business processes.

Note that we compared these scenarios to each other in their optimized state. In practice, few current supply chains are already optimized for inventory, because this information was not readily available until now. This means that, in practice, savings from moving to a new supply chain are likely to be even greater than projected in this comparison because even keeping the “as-is” structure would result in some savings if the inventory were optimized.

## 2.3 Implementation Results

Based on the results generated by the modeling exercise and the ability of PowerChain to make these results transparent to stakeholders across the supply chain, a larger implementation team undertook the transformation of the current supply chain to the configuration supported by Scenario 4. [Figure 4](#) shows actual inventory levels before and after implementation.



**Figure 4 Implementation Results**

This figure shows 1) actual inventory levels for Scenario 1 up until the implementation of Scenario 4, and 2) the resulting inventory levels after Scenario 4 was implemented. Note that the post-implementation inventory levels drop below the optimized levels that would have occurred had the group stayed with Scenario 1, and that the expected and actual inventory levels for Scenario 4 are closely matched. The expected/actuals for Scenario 1 pre-implementation differ far more than expected/actuals post-implementation, because the actual Scenario 1 was not optimized for inventory.

Figure 4 demonstrates two important results. First, the expected safety stock targets from PowerChain were realized in practice. The ability to implement a model's results in practice and see the numbers actually realized is critically important to gaining the user community's confidence in the tool. Second, the move from Scenario 1 to Scenario 4 has reduced not only inventory-driven cost but also the aggregate amount of inventory in the supply chain. This reduction acts to reduce the overall risk in the supply chain if demand comes in lower than expected. Given all the variabilities that face a supply chain, being wrong about something is a given. By lowering the aggregate inventory level, the cost of being wrong has been minimized. Because the cost of being wrong has been minimized, the overall risk has been minimized.

The camera business achieved over 30% reduction in inventory levels, as well as over 5% reduction in total supply chain costs. Furthermore, this new supply chain was able to maintain the existing

high service levels while at the same time reducing new products' time-to-market by 2-3 weeks.<sup>1</sup>

The project is 100% implemented, and has been for over a year. In fact, three generations of digital cameras have been sourced using scenario 4's supply chain configuration. To date, more than \$35M in savings has already been realized. The 5-year NPV of savings is well over \$50M. This is calculated using a one-time reduction in inventory, plus annual savings in variable and inventory-driven costs achieved through the adoption of scenario 4. At this point, what remains is the accrual of savings generated based on the implemented changes.

## 2.4 Disseminating Proven Technology

The success of these optimization techniques within the digital camera business prompted a wider dissemination of the tools across HP. The

1. We did not calculate the benefit from reduced time-to-market, but getting product to market faster is a significant benefit of this work.

Strategic Planning and Modeling Group (SPaM) has taken the lead in sponsoring and deploying the processes and the accompanying technology to other business units through a central server supported by HP IT resources. SPaM provides consulting support to help business units build their models. HP's Supply Chain Council, containing manufacturing Vice Presidents from across the company, has also endorsed this methodology.

In short, the digital camera business case had showed very promising results. We now move on to another business case, which illustrates the portability of these optimization techniques into one of HP's largest and most important business units: the inkjet supplies business. In this second case, the tool enabled migration from a traditional, localized inventory modeling approach, to a holistic and systems-oriented one, yielding new insights along the way that led to cost savings in otherwise unanticipated areas. In this next case, what started out as an inventory problem resulted in changes to transportation policy, in a journey that included the discovery of counter-intuitive findings that would not have been possible without a systems approach. Furthermore, the visual approach provided by the maps was a powerful tool for change, because to communicate these findings through complex analysis alone would have been a much harder sell.

### 3 Inkjet Supplies: Modeling Complex Vertical Networks

Hewlett-Packard pioneered inkjet printing technology in 1984, and continues to lead the world market in this area. Today, Hewlett-Packard's Inkjet Supplies organization manufactures and distributes ink cartridges, known internally as "pens", and distributes them through many commercial and consumer channels. These cartridges are used in HP printers, plotters, copiers, and fax machines around the world.

The inkjet business includes approximately 15 product families and over 250 manufacturing SKUs. Unlike digital cameras, inkjet supplies have a product life cycle in the tens of years, because the

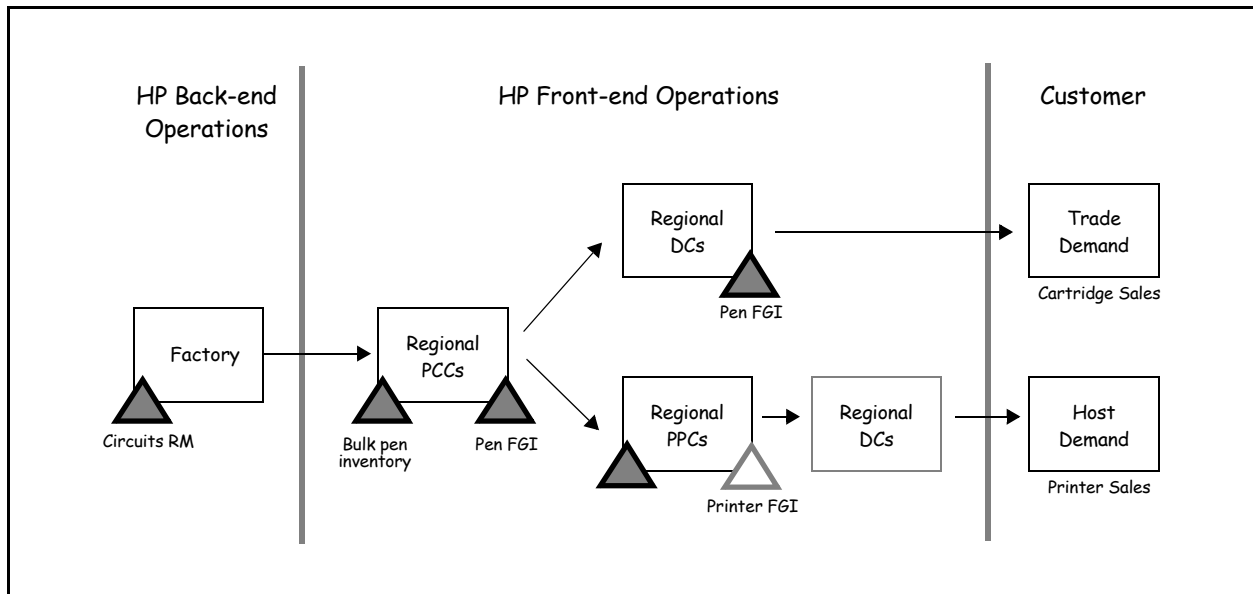
supplies are manufactured to support a large installed printer base (over 100 million printers sold around the world since 1987).<sup>1</sup>

#### 3.1 The Inkjet Supplies Challenge

The Inkjet Supplies network is very complex. Five internal factories ship inkjet cartridges in bulk to three regional completion centers for localization and packaging. These centers, called Pen Completion Centers (PCCs), ship the finished goods to tier-1 channel customers and numerous HP distribution centers. The PCCs also ship material to HP-owned Printer Postponement Centers, which bundle printers and cartridges for sale. The material flow supports two different demand streams: demand from customers who already have printers, served by the tier-1 channel partners and the HP distribution centers, and demand from printer sales, since all printers come bundled with a "starter kit" that typically includes an initial set of cartridges. [Figure 5](#) shows a high-level view of this network.

---

1. Source: HP presentation to analysts, December 1999.



**Figure 5 Inkjet Supplies Network**

This figure shows the vertically integrated Inkjet Supplies business. HP Back-end Operations includes 5 factories, which ship bulk pens to three regional Pen Completion Centers. After undergoing localization and packaging at the PCCs, the flow of finished products splits into two streams. The “trade demand” for replacement supplies for previously purchased printers goes to regional distribution centers as well as to Tier 1 channel customers. The “host demand” for cartridges bundled with new printers goes to regional Printer Postponement Centers. The bundled printer/cartridge products then go to regional distribution centers.

The above diagram also notes where safety stock is currently held to ensure high service levels. Bulk pen inventory at the three PCCs plays a key role in de-coupling customer-facing operations from the long lead times of the factories, and in buffering the system from the demand variability associated with the dozens of pen packaging configurations offered by HP. Bulk pen inventory is the largest safety stock buffer in the system; the concept of risk pooling raw inventory at the PCCs was to enable postponement of packaging decisions in order to respond to shifts in regional demand.

The front-end supply chain planning community knew that bulk demand variability was far less than end-SKU demand variability. While the current inventory buffer policy relied heavily on the common inventory at this holding location to respond to mix swings and end demand not covered by FGI safety stocks, the front-end organization wanted to know if there was any benefit to moving the bulk buffer, or some portion of it, further upstream to the factories. The rationale behind the question was that since each factory supplied

multiple PCCs, perhaps the supply chain could benefit from pooling demand uncertainty at the factory stage, before shipment to the three regional PCCs.

Unfortunately, the inventory modeling tools available to the front-end Inkjet Supplies planning community could not evaluate this opportunity. Utilizing high-level cycle time and variability assumptions, these applications were useful in establishing buffer strategy between any two locations (local optimization) but were incapable of identifying network-wide opportunities (global optimization).

This project, called the Supplies Inventory Modeling (SIM) project, was initiated with a two-fold mission: 1) to investigate the benefits of holding bulk safety stock at the factories, and 2) identify other areas of opportunity as time allowed.

### 3.2 Team Composition

Previous modeling and analysis projects had failed to result in major change within the Inkjet Supplies business, because the organization had

relied on consultants to perform and present the work. At the completion of the engagement, the consultants would leave, taking with them the modeling and implementation skills that were necessary to transition from modeling to pilot stages and beyond. Additionally, the business stakeholders had more leeway to dismiss findings that didn't fit their desired outcomes when the models did not have an internal champion to defend them.

To address this common problem, the SIM project assigned collective responsibility to roughly a dozen supply chain planners and financial analysts from the factories and regions across three continents. These people would build a common set of analysis skills, develop the supply chain models, complete the analysis, and communicate the conclusions themselves. The day-to-day experts on running the Inkjet Supplies business were the same resources who were analyzing the problem, building their own modeling proficiency, and answering detailed questions posed by management.

PowerChain was instrumental in enabling this team's work, because it provided an intuitive means of accessing some very advanced mathematical programming techniques. This would have been unheard-of had the team been using more traditional modeling methods. Furthermore, once they were trained, the globally distributed team members could collaboratively develop and manage models. Thus, their outputs were truly team deliverables, and represented a common understanding among all the team members.

### 3.3 Marrying Process and Technology

This section outlines the process HP followed to successfully conduct the SIM project. (REFERENCE TBD PICTURE)

**Training:** The initial training session required only two days of formal instruction. Over 50% of this instruction involved hands-on learning building a model, with the remainder of the time spent in an interactive lecture environment. For the first time in their careers, business analysts were able to globally optimize a supply-chain they drew on the screen of a computer.

**Mapping:** By the third day, the team was able to begin building the supply chain map for inkjet cartridges. The ability to draw supply chains rapidly, combined with a favorable learning curve, meant that after two days of formal training and a third day of project-focused model construction, the team had much of the supply chain structure built. Two factors contributed to this achievement. First, the optimization algorithms were flexible enough to "solve what you can draw," making rapid modeling possible. Second, and equally as important, the SIM team was able to use a "divide and conquer" approach to global supply-chain modeling. Individual team resources undertook to build those portions of the supply chain for which they were responsible on a day-to-day basis. This is vastly different from a traditional modeling effort, where one or two expert modelers construct the entire model based on interviews with those very individual resources. The visual nature of PowerChain enabled the team to quickly address logic issues and flow/data questions between representative portions of the supply chain.

**Data Collection:** Next, the team had to populate the models with data. In the overall two months of the project, three-quarters of that time was spent gathering data. If that data had been readily available, model development could have been completed in less than two weeks. Much of the effort involved taking existing data from HP systems, and extracting out enough detail to meet the granularity of the process steps mapped in the supply-chain.

**Validation:** Model validation occurred in two phases. First, modelers verified the flow of logic and input data. Second, HP's Finance group verified the current-state models against actual costs by comparing supply chain cost and inventory performance with actual booked values. Model validation required a day or two.

**Optimization:** The team optimized the model. This process took less than 10 seconds using an HP server.

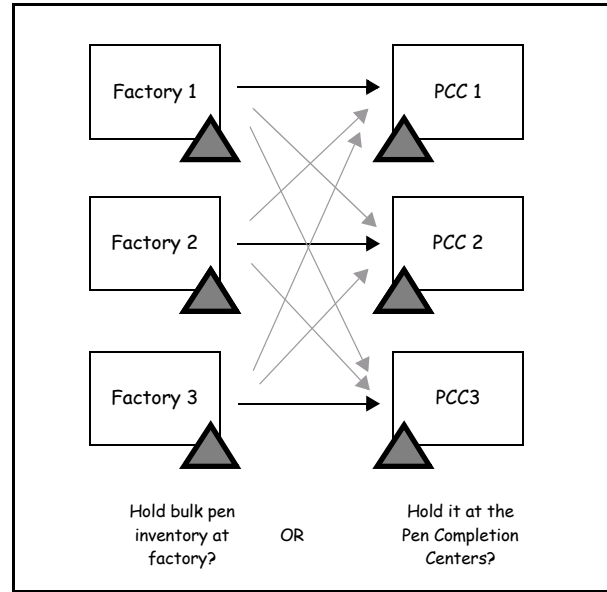
**Constraining Stages:** This can be thought of as the "reality" step in the process. Once the team identified the global unconstrained inventory solu-

tion, specific process steps were constrained to match real-world conditions. The most common constraint involved forcing a stage to hold no inventory, in other words setting its net replenishment lead time to zero. In a few cases, a stage was forced to hold inventory to match a business process requirement, by setting its service time to zero enabling immediate availability. Pipeline stock costs were also constrained, to accurately reflect the portions of those inventories on HP's books. During this phase, the team also added existing supply chain inventory policies into the model to provide a baseline comparison. Modifying the supply-chain model to accurately reflect reality required approximately one week.

**Second Validation:** The team re-visited the assumptions in the model and compared the overall supply-chain costs with aggregate financial data from HP systems. This validation landed the model's supply-chain costs within a penny of Finance's projected costs! The similarity in outputs between PowerChain's and Finance's results was another major success factor in the project; the team believed the one-penny differential was below the noise or measurement limit of their efforts.

**Optimization:** The team next optimized the constrained model. This process took less than 10 seconds using an HP server.

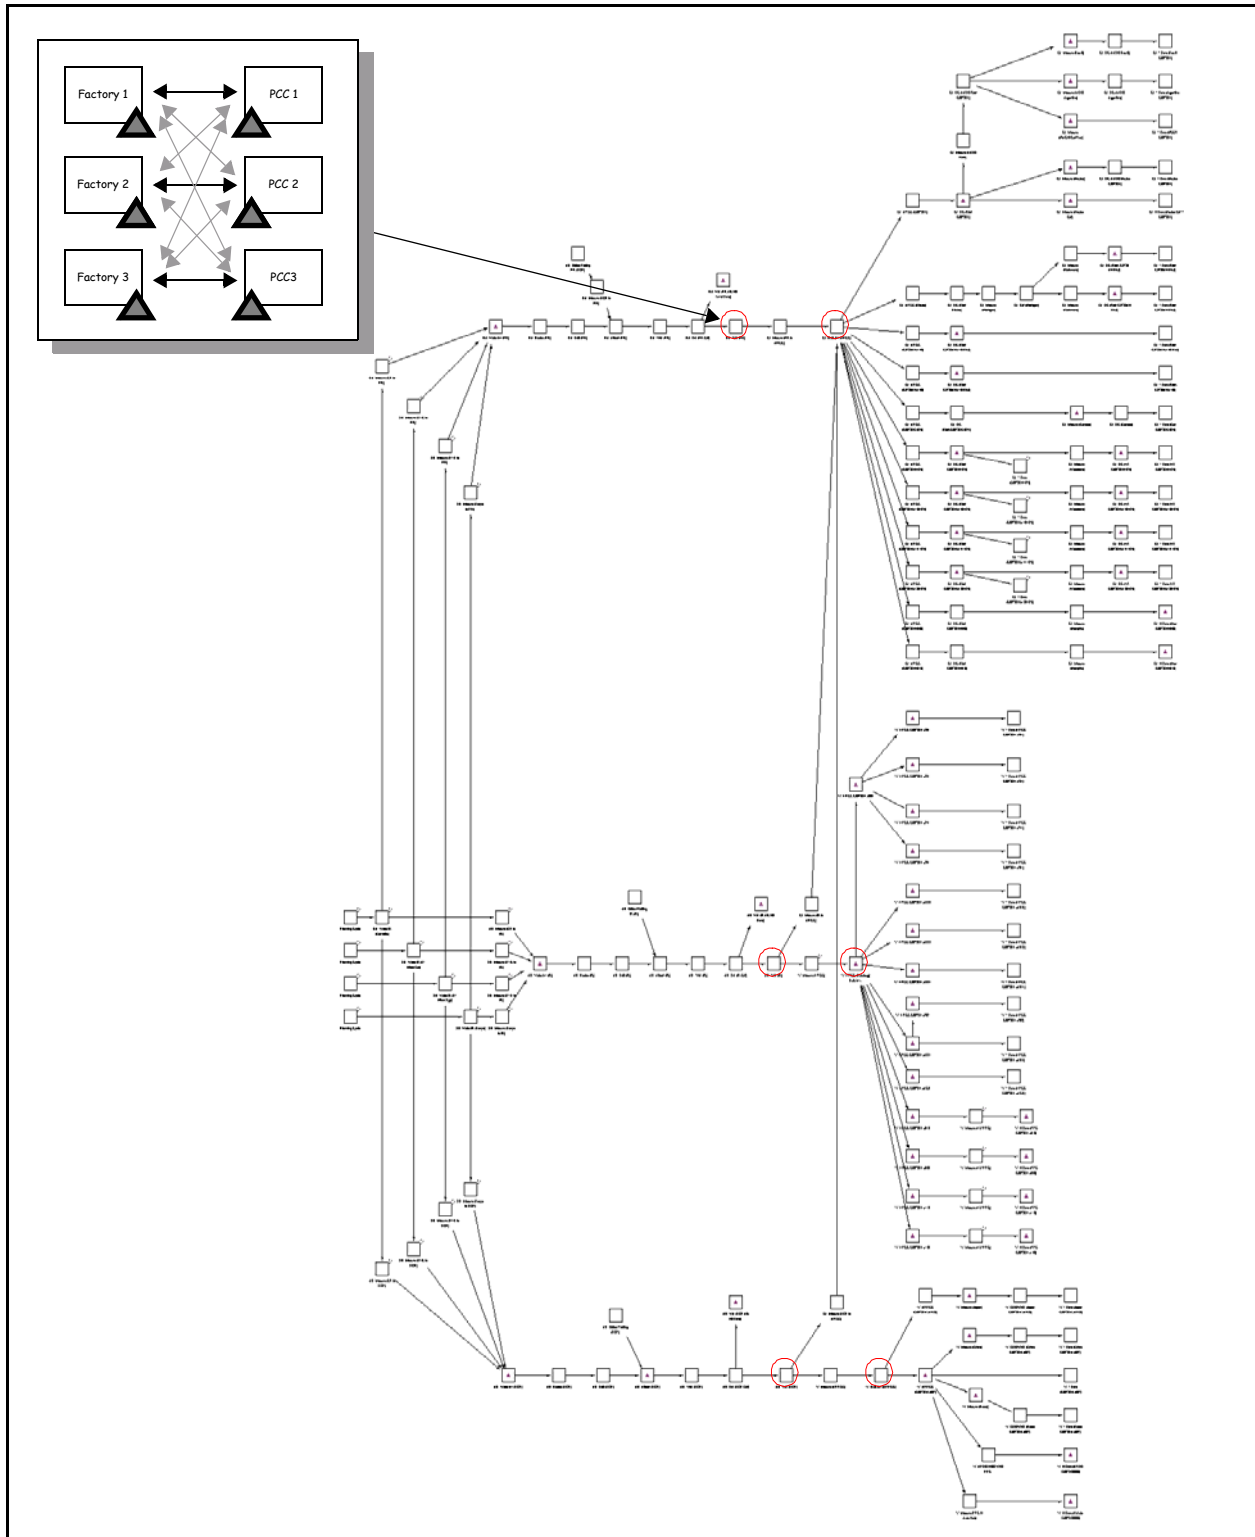
**Analysis:** The team now set out to interpret the results from their global inventory modeling effort, and to do this within a week. Figure 6 illustrates the scope of the original question posed to the team by sponsors, namely to focus on six inventory locations within the supply-chain. Traditionally, the approach for solving inventory problems is to treat each stage in isolation, and assume that they will all hold safety stock. In our experience, this decoupled approach is imposed by the limited capabilities of the current generation of commercial supply-chain optimization tools.



**Figure 6 Localized Inventory Optimization**

This figure shows the localized inventory optimization as performed by traditional inventory-optimization tools. Each pair of facilities is evaluated independently of all the others, which could result in overall sub-optimization.

The existence of a new capability for optimizing across an entire supply chain allowed the SIM project to go far beyond these localized inventory calculations. For the first time, analysts could evaluate both the best holding locations, and the optimal quantities for those locations, in order to meet pre-defined customer service level requirements. Instead of performing a localized analysis that would likely sub-optimize total system performance, the team could evaluate an entire system, and, within that global context, determine the best strategy for bulk pen inventory management. See [Figure 7](#) for an example of resulting supply chain map developed in PowerChain.



**Figure 7 Localized vs. Systems Approach**

This illustration shows a typical PowerChain supply chain model, to compare with the localized approach shown in [Figure 6](#). The 6 locations shown in [Figure 6](#) are circled in this diagram.

The results were very surprising. Contrary to expectations, the optimization did not show a benefit to holding bulk safety stock at the factory level. Instead, it suggested that the system would be better served if safety stock were held at locations other than bulk all together. This finding might never have occurred to the team, had they been forced to explicitly model only a small number of supply-chain stages - another benefit to a system-wide optimizer.

By adjusting the quantity of raw-material inventory in front of the upstream factory processes and the amount of packaged FGI inventory near the customers, an optimal buffer strategy emerged that defied intuition.

Figure 8 and Figure 9 show the back-end and front-end supply chain characteristics of the resulting inkjet supply chain.

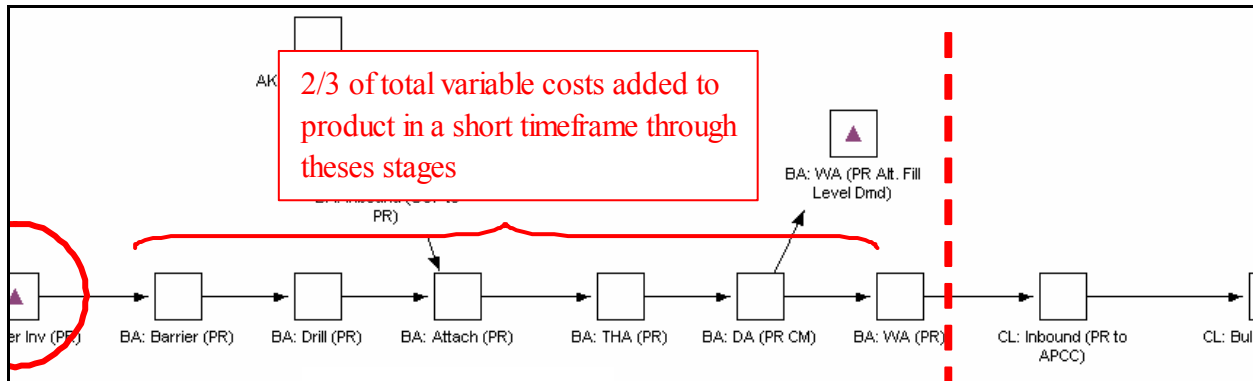


Figure 8 System-Optimized Inventory Strategy: Factory

In the factories, PowerChain recommends holding wafer inventory to minimize the system cost, recognizing that 1) it can be quickly converted to bulk as the system requires, and 2) holding wafer safety stock de-couples the rest of the system from the long lead times associated with wafer fabrication. The upright triangles represent the recommended (optimized) locations for holding safety stock.

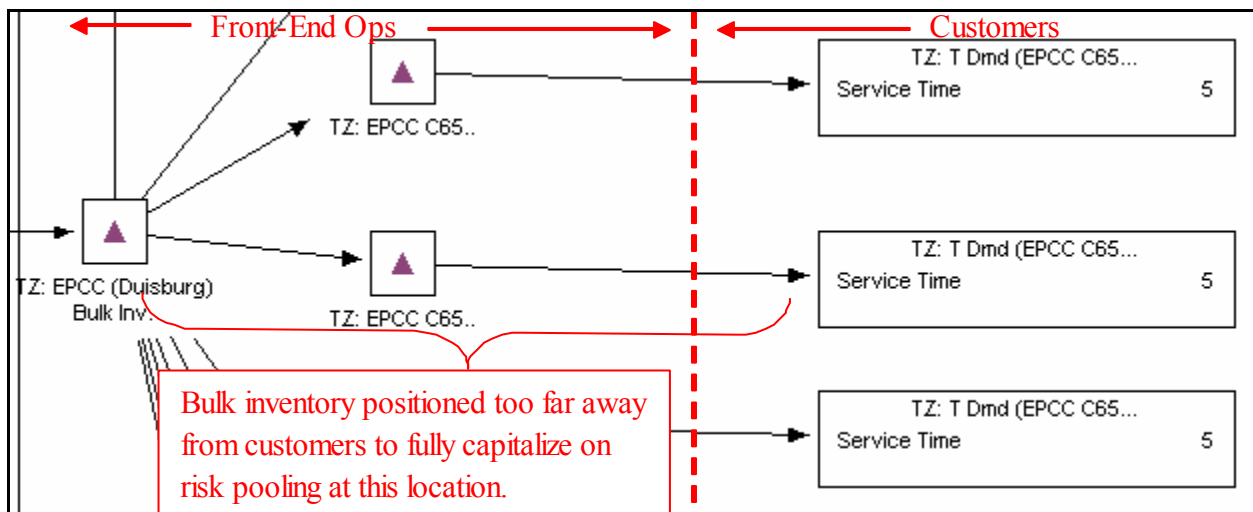


Figure 9 System-Optimized Inventory Strategy: Front-End

In the regions, analysis shows that the optimal place to hold the majority of safety stock to buffer against demand variability, given the customer-requested lead times, is near the customer. To capitalize on build-to-order operations, which would pull directly from bulk inventory, pipeline would need to be removed from the system. The upright triangles represent the best locations for holding safety stock.

When the original bulk inventory question turned out to be a red herring, the team used its remaining time to evaluate other opportunities. The team was originally looking to optimize inventory within the existing supply-chain. However, PowerChain provided an unsurpassed ability to analyze the end-end supply-chain in detail, thus leading to exploration of other questions as well. The challenge for the team became the need to restrain themselves from exploring too many possible avenues, because it was so easy to duplicate, edit, and compare the initial supply-chain against "what-if" improvements.

Within a few days, the very resources that were responsible for day-to-day operations of the supply-chain made a discovery. By constructing a global model, they were forced to detail the cost and time of individual freight lanes within the model. If the team were to move transoceanic freight lanes from air to sea, they calculated a positive NPV of \$83M could be realized. But this would require an increase in supply-chain inventory, a challenging sell in an environment focused on inventory reduction. The modeling team needed to demonstrate that the product cost savings far outweighed the additional inventory-driven costs incurred by the longer transit times.

The movement from air to sea seemed completely counter-intuitive. Traditional rules of thumb had focused on responsiveness through speed. Air shipment was seen as a way to help minimize both cycle times and inventory levels. Additionally, the planning granularity of the current inventory management system and a highly variable demand signal made the notion of increasing response time seem counter-intuitive, because it risked affecting end-customer availability. The model showed that by appropriately sizing the inventory at the DCs, expensive air freight charges could be avoided without any change in service seen by the customer.

Many of these insights would have been missed if the project team had focused on a localized modeling approach centered around bulk pen inventory, with outcomes constrained to consider only two possible choices (hold bulk in regions or

hold bulk at factories). In this exercise, the bulk inventory question was indeed answered - and, in addition, other potential opportunities were identified based upon insights gathered from use of a system-wide optimization approach by the people who had detailed knowledge of the day-day supply-chain operations.

**Communication:** The purpose of communicating results is to influence change. Of course, from a change-management perspective, the analysis stage is the easy part. Even a simple-sounding recommendation such as the air-to-sea switch can require tremendous effort during implementation, because of the size of the business involved. For this reason, large organizations are often highly resistant to change, and when the arguments in favor of the change are largely expressed as arcane-sounding modeling exercises performed by outsiders, even selling the idea can be a Herculean task.

Having a graphical modeling tool that the change agents themselves could use and understand facilitated both change acceptance and implementation. In this case, the team was trying to persuade the organization to change processes that had taken years to put into place and were ingrained conventional wisdom. This work also challenged the notion that the size of the inkjet cartridge business was a barrier to the application of management science. The graphical nature of the tool became a key communication aid during presentations and working sessions to multiple management tiers. Managers could clearly see the structure of their supply chains in a way that wasn't even possible before. Having a picture to look at helped the stakeholders feel that they were in control, because they could understand what they were seeing. The "black box" of traditional supply-chain modeling had been replaced by an open window into the cause and effect of supply chain performance.

So, in an environment where senior management had requested a reduction in overall inventory, the Inkjet Supplies team instead successfully influenced stakeholders to increase inventory in

the supply chain to reduce overall supply chain costs!

**Implementation:** For successful implementation, it is imperative that all stakeholders understand the analysis. Portability and standardization of the analysis techniques made this possible. Furthermore, since the actual people that would need to do the implementation conducted the modeling themselves, they did not need to be convinced of the results. The stakeholders could all agree on what the supply chain characteristics really were, without having to resort to higher math. For example, a question about postponement at a particular point could be quickly answered by looking at the supply-chain maps.


The implementation plan was finalized and implemented two months after the completion of the SIM project. There were several products that were targeted to move to sea shipment. For others, however, switching to sea transport did not make sense, for reasons such as product economics or a short shelf-life. Of those that can be shipped by sea (the population of interest), there were several considerations taken into account. Engineering needed to validate that each individual product's quality would not be negatively affected by any aspect of the sea transportation mode. Manufacturing needed to have sufficient manufacturing capacity to produce the extra units required to fill the new pipeline and still keep regional buffer levels at target. To address this last concern, factory schedules were modified with a 9-month transition plan.

Within six months of the implementation's start date more than 90% of the units targeted to ocean freight were indeed on the ocean. This is providing the business an annuity stream of over \$2M per month, which goes directly to the bottom line. Although this is respectable enough, there is one other point to highlight about the importance of these numbers in the context of our systems approach. For example, there is no point in touting \$100M in savings if it costs an upstream supplier \$120M on their end. It can happen that localized solutions have unreported negative consequences on other partners; the long-term result is, those

partners will eventually push back, or be forced to end the business relationship. With a systems-wide approach, the model considers the overall impact on the end-to-end supply chain - including partners - thus reducing the possibility of consequences that later undermine any reported gains.

With the Inkjet Supplies exercise, it is important to highlight the speed with which the team was able to discover that they hadn't been investigating the right question to begin with. The original question - whether to move inventory from one location to another - was quickly supplanted with another opportunity, namely, the air-versus-sea transportation question. Once the team found the right question to ask, they found this opportunity and implemented it almost immediately.

This project highlighted the benefits of coupling the rapid analysis capabilities provided by a global inventory optimization algorithm with an intuitive interface that allowed the real-life business users to conduct the modeling exercises themselves. The in-house team identified opportunities in flexibilities which they knew about, but that a consultant might have taken for granted.



**Figure 10** Some Kind of Picture would be nice

Maybe a high-level depiction of cost savings over time?

## 4 Learnings

We chose to discuss two particular business cases in order to illustrate the application of a methodology in two very different business problems, and to demonstrate the portability of these techniques within diverse organizations. However, the use of these techniques currently extends far beyond these two business units within HP. Product lines currently using this modeling method include flatbed scanners, inkjet printers, servers, and laser printers. The user community includes analysts in HP regions all over the world, as shown in [Figure 11](#).

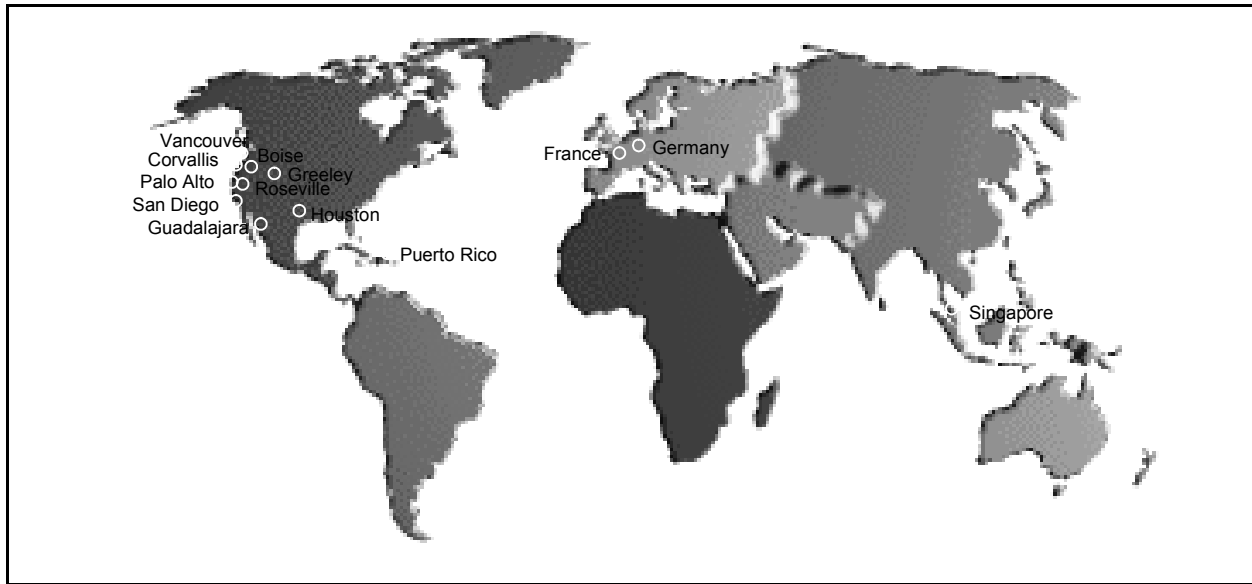


Figure 11 HP User Community

This map shows user communities, with different size dots representing different sized groups of users. [<map prints poorly b/w. make circles black outline white interior and if possible labels should be off the map to avoid contrast problems>](#)

The techniques are currently disseminated on an as-requested basis. To date, we have conducted over 9 training sessions around the globe for over 120 people. In addition to these trainees, there are currently at least 80 more who have expressed an interest in adopting these techniques. Clearly, the need for standardized, portable inventory optimization techniques is significant. For the first time in HP's history, we have a tool that can do this.

This deployment to a broad distributed user base goes beyond reaching the people we have actually trained. These very people influence a broader group of people who are not interested in the deeper analysis but are interested in the broad implications of the analysis. With a common language and framework, the efficiency of supply-chain analysis for inventory optimization has dramatically increased within HP.

Supply chain engineers have the task of re-engineering business flows without having any real authority to do so. The people they depend on for implementation are not directly reporting to them. Achieving buy-in across the supply chain from a position of non-authority is very difficult.

In order to create change on a broad scale and obtain the cooperation of diverse groups and indi-

viduals, all the players must understand the details and trade-offs. Without the ability to determine and validate the inventory costs, achieving a majority buy-in is simply not possible. While more sophisticated business people recognized that optimizing safety stock levels can reduce total supply-chain costs, they also know that determining any sort of optimal inventory position is very difficult. Sometimes, the business users spend quite a bit of time developing their own models. Others rely on gut feel. Either way, they commonly fail to persuade others to share their views.

Having an intuitive and proven analysis tool removes this major barrier that stands between analysis and implementation. Now, an entire user community can determine global inventory levels within a structured framework. This translates into the following benefits:

1. **Education of the business community.** With a structured method for determining supply-chain costs, business users learn which parameters are important and how the parameters interact. Best of all, the modeling team can show the behavior of the supply chain to non-modelers.

2. **Communication to a wider audience.** Stakeholders from across the organization can use the same framework for discussing the supply chain. This is critical for holistic analysis. Users can see the trade-offs between local sub-optimization and global optimization, thus creating buy-in within their local divisions. By using an algorithm that is flexible enough to solve what users can draw, users get the benefit of a picture that is worth a thousand spreadsheets when it comes to communicating their findings to a broader audience.
3. **Collaboration across supply-chain stakeholders.** Having an algorithm flexible enough for each user to add their part of the supply chain greatly aids the data gathering process. It is extremely difficult for one advanced modeler to capture the entire network; it is equally difficult for two advanced modelers to combine their respective custom-built models. It is sometimes next to impossible for them to get the results of their models implemented, since the people responsible for implementation often do not understand the modeling effort.
4. **Empowerment of business users.** By providing sophisticated supply-chain analysis to a broad user base, we bring many more ideas, insight, and understanding into a formerly elite modeling community. The people who operate the supply chain can now understand its behavior on a whole new level. Understanding through hands-on modeling is mandatory to drive, and lead, effective change.
5. **Better data collection.** One of the unforeseen benefits of using this tool is that each business user can input their own portion of the data. In other situations, supply chains tend to be modeled using data that is readily available to the modelers. Many times, they miss or cannot get at pertinent

data that might be much more available to the field users.

## 5 Conclusion

Imagine an environment where large numbers of business analysts can now model a supply chain together, make global decisions together, conduct sensitivity analyses together, and make a unified recommendation to a cross-organizational management team based on their efforts. Imagine management from across HP seeing these unified recommendations and quickly authorizing implementation. Imagine stakeholders from across the supply chain understanding the global impact of optimized inventory tradeoffs. Imagine an implementation team that can respond to on-the-fly changes in the supply chain and model them rapidly.

For the first time within HP, and, we believe, in the world – we have brought inventory optimization techniques to a broad audience. It no longer resides only in the hands of highly trained OR individuals. HP is staying true to its heritage of innovation. Not only have we introduced a better technical analysis method, but we have also introduced a framework for people to communicate, collaborate, model, and understand their supply chains better. This new combination of technical innovation enabling broad-based organizational decisionmaking is noteworthy, and we believe that it is exactly what Bill Hewlett and David Packard would expect from the employees of HP.

The business cases presented in this paper highlight how HP has benefited from a standardized process to optimize total supply chain cost. We show an example of a highly cyclical, short life-cycle business deriving significant benefit and prove the portability of this approach to a vertically integrated, long life-cycle business. The significance is one set of process and technology could yield benefit to radically different business situations and identify completely different opportunities. This highlights the power and robustness that can be derived from marrying optimization and process excellence.

## 6 References

- 
- Arntzen, B. C., G. G. Brown, T. P. Harrison, and T. L. Linda (1995), "Global Supply Chain Management at Digital Equipment Corporation," *Interfaces*. 25, 69-93.
- Camm, J. D., T. E. Chorman, F. A. Dill, J. R. Evans, D. J. Sweeney, and G. W. Wegryn (1997), "Blending OR/MS, Judgement, and GIS: Restructuring P&G's Supply Chain," *Interfaces*. 1, 69-93
- Ettl, M., G. E. Feigin, G. Y. Lin and D. D. Yao (2000), "A Supply Network Model with Base-Stock Control and Service Requirements," *Operations Research*. 48, 216-232.
- Graves, S. C. and S. P. Willems (2000), "Optimizing Strategic Safety Stock Placement in Supply Chains," *Manufacturing & Service Operations Management*, 2, 68-83.
- Humair, S. and S. P. Willems (2003), "Optimal Inventory Placement in Networks with Clusters of Commonality," Working Paper, January, 19 pages.
- Lee, H. L. and C. Billington (1993), "Material Management in Decentralized Supply Chains," *Operations Research*, 41, 835-847.
- G. Lin, M. Ettl, S. Buckley, S. Bagchi, D. Yao, B. Naccarato, R. Allan, K. Kim, and L. Koenig (2000), "Extended Enterprise Supply Chain Management at IBM Personal Systems Group and Other Divisions," *Interfaces*, 30, 7-25.
- 

### The Algorithm and PowerChain Architecture

[<I didn't touch this section>](#)

PowerChain is a collaborative application for the interactive visualization, design, and optimization of supply chains. Using web-based technologies in a multi-user environment, the software encompasses the entire modeling lifecycle, including model creation, supply-chain structure mapping, data loading, optimization and sensitivity analysis, and model versioning and workflow. A typical project will analyze one or more supply-chain scenarios (such as the five digital camera options presented in the paper). Within the PowerChain software, a scenario is modeled by a construct called the "chain object." A chain object has its own version control, similar to what is found in document management or source-code control systems: a chain can be published to freeze its state; checked-out as a new version to make changes; and shared among different users who are stakeholders of the analysis.

A chain object comprises global settings, stages, and links. Global settings include such

parameters as the name of the chain, the modeling horizon, and whether advanced features such as batching and capacity are enabled. Stages identify and encapsulate the meaningful activities that describe the supply chain structure. It is at the stages that costs and time are incurred, demand is calculated and satisfied, and where safety stocks are potentially held. Most of the data requirements for PowerChain relate to stages, including demand, target service levels, holding cost rates, yields, capacities, cost structures, and time delays. Links relate stages to each other and represent the flow of materials or information, but by themselves contain very little data, mainly a multiplier that indicates a "goes into" factor from the source stage to the destination stage. Stages with no outgoing links are demand points, while those with no incoming links are raw material suppliers.


Because stages contain most of the data, in particular cost and time data, links are not transportation activities per se. Instead, transportation activities are typically modeled as stages themselves, often with names such as "Ship to DC." In fact, the concept of the "stage object" is sufficiently abstract that a stage is often characterized as, "any

activity that has some cost or takes some time, and after which you might want to hold safety stock.” A stage can but need not have a defined spatial relation to other stages, such as a geographical location. As such, within the same supply chain model, ten stages might be used to model the behavior within an assembly plant while one stage is used to model a distribution facility. Conversely, if the user is analyzing a distribution network, the “raw materials” for a supply chain model might be the finished goods originating at a single stage (the

factory) while the distribution network is modeled by hundreds or thousands of stages.

The flexibility of the individual stage coupled with the optimization and calculation techniques embedded within the software are what enable PowerChain to solve “what the user draws.” However, because of the web-based, multi-user, distributed nature of the application, the path from drawing a supply chain to optimizing inventory policies is hardly direct. Figure 12 depicts the system component architecture of the PowerChain software:

**Figure 12 PowerChain Architecture**



Caption: A user connects to the PowerChain application using a web browser and draws a supply-chain using a Java applet. The Web Server generates requests for performing supply-chain actions using an XML schema specifically designed for supply-chain modeling. These XML requests are dispatched and turned into commands such as “Add Stage” or “Optimize.” An asynchronous job manager handles complicated

and time-consuming actions like optimization, while business objects, reporting, and security layers handle traditional application-tier activities. All of these operations rely on a relational database.

The use of web-technologies such as a web browser, HTTP, HTML, and XML enables HP to quickly add new users to the PowerChain system at almost zero cost. A system administrator only needs to add the user to the PowerChain database

(via a simple registration screen) and then send the user the URL to the PowerChain login screen. Furthermore, after initial installation and creation of backup procedures, the server system has very low maintenance requirements. Therefore, the total IT involvement and commitment is minimal. Having a low total cost of ownership is critical to enabling wide-spread deployment within a company like HP.

While initially deployed at HP on a single 4-processor Intel-processor HP server, PowerChain can, because of its modular component architecture, be deployed across a farm of servers. In particular, a set of servers can be configured for each of the primary tiers of operation: database, job, application services, and web. Many traditional  $n$ -tier systems have database, application, and web farms. To support the complexity and system requirements of the optimization process, PowerChain separates optimization from normal business logic. PowerChain includes a “job” mechanism that allows the complicated activities to be enqueued within an array of job servers which process these jobs asynchronously. This enables a client to scale the number of jobs they can process in parallel as easily as scaling a web farm or database cluster.

Each job is either an optimization plan or a sensitivity analysis plan. In either case, the process involves initialization, validation, algorithm execution, output generation, and termination. Initialization comprises loading supply-chain data from the PowerChain database, allocating memory necessary for the algorithm, and preparing logging activities. Validation checks that the supply-chain data satisfies the job requirements, such as checking that the supply-chain map is a directed acyclic graph. The algorithm execution for optimization includes two main steps: algorithm selection and algorithm processing. In algorithm selection, the supply-chain map network structure is used to determine which of a handful of algorithms is most suitable to optimize the chain. The selected algorithm is then run. Outputs are generated and stored back into the database, and finally termination cleans up the allocated memory and signals to

the server that the job is complete. Sensitivity analysis varies from optimization in that it iterates over the range of sensitivity inputs and either calculates or optimizes multiple times.