

Accelerating the Profitability of Hewlett-Packard's Supply Chains

Corey Billington

Global Procurement Services, Hewlett-Packard Company, 3000 Hanover Street, Palo Alto, California 94304-1185,
corey.billington@hp.com

Gianpaolo Callioni

Strategic Planning and Modeling, Hewlett-Packard Company, 3000 Hanover Street, Palo Alto, California 94304-1185,
gianpaolo.callioni@hp.com

Barrett Crane

Digital Imaging, Hewlett-Packard Company, 3404 East Harmony Road, Fort Collins, Colorado 80528, barrett.crane@hp.com

John D. Ruark

Optiant, Inc., 4 Van de Graaff Drive, Burlington, Massachusetts 01803, john.ruark@optiant.com

Julie Unruh Rapp, Trace White

Inkjet Supplies, Hewlett-Packard Company, 1000 NE Circle Boulevard, Corvallis, Oregon 97330
{julie-unruh.rapp@hp.com, trace.white@hp.com}

Sean P. Willems

School of Management, Boston University, 595 Commonwealth Avenue, Boston, Massachusetts 02215, willems@bu.edu

Hewlett-Packard (HP) developed a standard and common process for analysis coupled with advancement in inventory optimization techniques to invent a new and robust way to design supply-chain networks. This new methodology piloted by HP's Digital Imaging division has received sponsorship from HP's Executive Supply-Chain Council and is now being deployed across the entire company. As of May 2003, a dozen product lines have been exposed to this methodology, with four product lines already integrating this process into both the configuration of their new-product supply chains and the improvement of existing-product supply chains. The team will highlight the application of these new capabilities within HP's Digital Camera and Inkjet Supplies businesses. The realized savings from these first two projects exceeds \$130 million.

Key words: industries: computer, electronic; manufacturing: supply-chain management.

In 1939, the Hewlett-Packard Company (HP) introduced its first product, an audio oscillator, which it sold to Walt Disney Studios for creating some of the special sound effects for the movie *Fantasia*. In May 2002, HP merged with Compaq, with the goal of becoming the world's leading provider of information technology (IT) solutions. By the end of 2002, HP's revenues were \$72.3 billion. Today, HP operates in over 170 countries and is a leading provider of fault-tolerant servers, Windows servers, Linux servers, storage systems, management software, personal computers, pocket PCs, inkjet printers, multifunction printers, laserjet printers, flatbed scanners, printing services, and IT services.

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

so fierce that it must properly design, configure, and optimize the supply chain wherever possible.

Cost factors, such as material devaluation, scrap costs, write-offs, and fire-sale discounts, have become the single biggest detriment to profitability, sometimes eclipsing the already-slender profit margins. To calculate these inventory-driven costs, HP must consider 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 problem very challenging.

Throughout the 1990s, 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). Managers and planners from different HP organizations commonly conducted supply-chain analyses on their own by building models in Excel. They would then share their models with each other

and spend considerable time trying to get their peers across the supply chain to buy into their analyses.

With 141,000 employees in 178 countries who coordinate their operations with a global network of suppliers, HP cannot rely on a few experts to make good supply-chain analysis decisions. For some time, SPaM had been searching for a robust methodology that would enhance business units' ability to conduct end-to-end modeling of the supply chain. Over the years it had developed robust methodologies and analysis techniques to solve complex supply-chain problems. However, the sophistication of the modeling often made SPaM a bottleneck in the process of making supply-chain decisions. Although the process and methodology were clearly transferable, SPaM could not find a modeling tool that it could easily transfer to the business units. More than 60 percent of the group's resources were allocated to network analyses, but demand was still not satisfied. In addition, because so many of its resources went to supply-chain analysis, SPaM could not focus on complex problems outside the supply-chain domain. HP needed an application that wouldn't require business users to become expert modelers. It had to have the supply-chain theory packaged so that users could analyze and model the key supply-chain parameters through a controlled process. It needed an application it could distribute to users in all business units and regions so that they would all have access to the same set of tools. SPaM was monitoring supply-chain advances from academia to find such a solution.

In 2000, HP began working with Optiant, a Boston-based developer of supply-chain design software, to bring HP's long-standing tradition of supply-chain innovation to a new and broader level. Optiant's product, PowerChain Inventory, combined a robust multiechelon inventory-optimization engine with an easy-to-use Web-based architecture and interface. PowerChain grew out of work originally conducted at the Massachusetts Institute of Technology (MIT) (Graves and Willems 2000), continued at MIT and Boston University (Graves and Willems 2002), and at Optiant (Humair and Willems 2003). A notable feature of the software was its drag-and-drop interface, which allowed users to quickly create and optimize any network topology.

We describe both the process HP developed for determining total supply-chain cost, and the supply-chain theory that we developed to solve these real-world supply-chain problems. We highlight the first two applications of this methodology inside HP: the initial pilot project completed in 2001 by the Digital Camera business and the initial application conducted by Inkjet Supplies. For just these two product lines alone, HP reduced total supply-chain costs by over \$130 million while maintaining already-high service levels.

Optimizing the Digital Camera Business

Digital imaging is a key growth area for HP. HP's Digital Imaging organization produces an array of imaging products. Consumer offerings include digital cameras, flatbed scanners, and photo inkjet printers. HP's commercial products include Phogenix photo-processing equipment and Indigo digital presses. The Digital Imaging group develops complete imaging solutions for 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, it sells over half a dozen cameras, ranging from aggressively priced autofocus cameras to high-optical-zoom, high-megapixel products. The digital camera business is extremely competitive; HP competes with Canon, Kodak, Nikon, Olympus, and Sony. Product life cycles commonly run less than a year, with prices rapidly declining over this period.

In the fall of 2000, the manufacturing manager for the digital camera organization asked his team to improve the performance of the digital-camera supply chain. Led by supply-chain engineers, this effort was supported by resources from finance, manufacturing engineering, research and development, and logistics. After considering various potential supply-chain alternatives, the team identified five scenarios to analyze in depth (Figure 1).

Scenario 1 was business as usual. In the existing supply chain, digital cameras went from a dedicated factory in Asia to a single worldwide hub, also in Asia. This hub localized the cameras (packaged individual items for regional consumption) and sent




Figure 1: The five scenarios for digital camera supply-chain configurations considered various levels of consolidation and localization between worldwide manufacturing facilities, a worldwide localization hub, and several regional distribution centers.

build-to-stock shipments to regional distribution centers (DCs). The regional DCs shipped the final products to retailers.

Scenario 2 eliminated the Asia hub facility, instead opting to postpone localization until the cameras reached the regional DCs. Scenario 3 also eliminated the hub but moved localization back to the factory. Scenario 4 bypassed the regional DCs and delivered directly to the retailers. In this case, the Asia hub localized the products and shipped them directly to the retailers.

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

Determining Supply-Chain Cost

To decide which configuration to adopt, the team needed to calculate the total supply-chain cost for each scenario. 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 are those that are not volume sensitive and depend only on the supply-chain configuration. They include such items as the portion of a DC's costs that are allocated or charged against the particular product line. Inventory-driven costs are dictated by the cost of the product, the holding-cost rate, the service level, the supply variability, and the demand uncertainty.

Aggregate fixed and variable costs are fairly straightforward to calculate. For fixed costs, one needs to know the cost of the fixed assets. For 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 difficult to forecast. In particular, one year ahead in the budgeting process it is impossible to know for certain which products will be in high demand and which will be in low demand; which components will be delayed; and which suppliers will have quality problems. However, these uncertainties are some of the reasons for holding inventory and major causes of inventory-driven cost.

The most difficult part of minimizing the inventory-driven costs in the supply chain is 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, too little of the models in demand can create allocation conflicts and lost sales.

Within the worldwide supply chain for digital cameras, the team modeled 44 potential inventory locations. Calculating a globally optimal solution for a

problem with so many variables requires a depth of knowledge and modeling expertise beyond those of business users armed with spreadsheet-based tools.

The team turned to the PowerChain application specifically to solve this inventory-optimization problem and to quantify inventory-driven costs. The optimization engine coupled with an intuitive interface allowed us to introduce user-friendly, robust supply-chain modeling to business users. For each of the five scenarios, the team developed a supply-chain model similar to the example shown in Figure 2.

Findings

For each supply-chain configuration, the cost was described by a point along an efficient frontier. Scenario 4 (ship directly from the Asia hub) has the lowest overall cost (Figure 3). The output of the model provided the quantitative basis for the recommendation, enriched by consideration of other costs not affecting the inventory policy of the supply chain. Adopting Scenario 4 provides two benefits. First, we can create a step-function improvement in supply-chain performance, which translates directly into cost reductions, by restructuring the supply chain. Even with no other improvements in the business process, HP benefits from the efficient structure provided by Scenario 4. Second, knowing the optimal inventory policy for HP's global supply chain, we can implement it through changes to business processes.

We compared the optimized states of the five scenarios. In practice, few current supply chains are already optimized for inventory because the information required was not readily available until now. This means that, in practice, HP's 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.

Implementation Results

Based on the results of the modeling exercise, an implementation team larger than the analysis team transformed the existing supply chain to the Scenario 4 configuration (Figure 4).

The implementation had three important results. First, HP realized the optimal inventory policy in practice. Implementing the results and seeing the numbers realized is critical 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, freeing up cash for running the business. Third, this reduction reduces the downside risk in the supply chain. Given all the variabilities in a supply chain, being wrong about something is a given. By lowering the aggregate

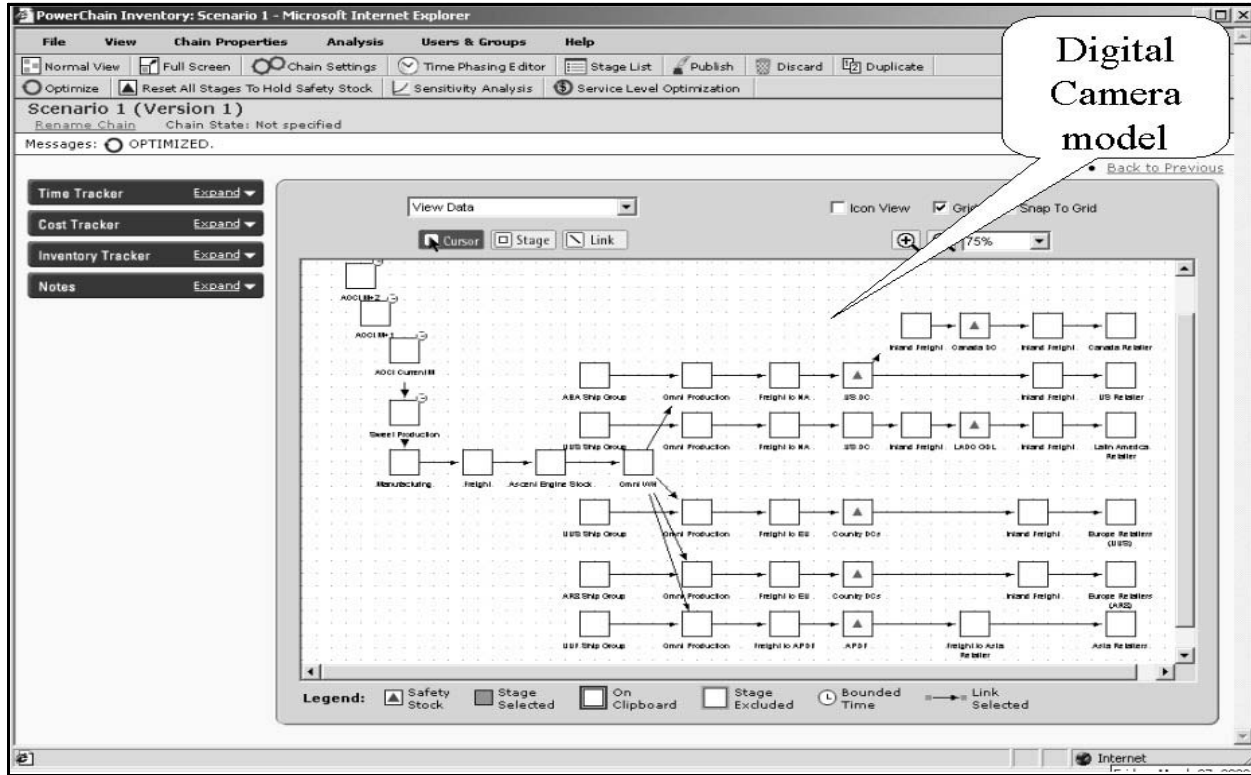


Figure 2: The PowerChain depiction of the Scenario 1 digital-camera supply chain consists of 44 stages representing the raw materials, manufacturing, logistics, and customer demands corresponding to the specifics of the scenario's supply-chain configuration. Each box on the supply-chain map represents a value-added step, such as product transformation or transport. A triangle within a box denotes safety stock held at that stage. The team optimized inventory levels and locations across the multiechelon network. This supply-chain structure was one of the first models HP built, and it is among the simplest.

Note. The text associated with each stage is intentionally blurred to protect confidential information.

inventory level, HP minimized the cost of being wrong and consequently the overall inventory risk.

The digital-camera supply chain reduced its inventory levels by over 30 percent and reduced total supply-chain costs by over five percent. Furthermore, this new supply chain maintained the existing high levels of service while reducing new products' time to market by two to three weeks. We did not calculate the revenue and cost benefit from reducing the time to market, but getting product to market faster is a significant benefit.

The project is 100 percent implemented. In fact, HP has produced three generations of digital cameras using Scenario 4's supply-chain configuration. That is, after its initial success in converting the existing digital-camera product lines from Scenario 1 to Scenario 4, HP has built all subsequent generations using Scenario 4's structure and inventory targets that depend on the new product line's demand characterization. The five-year net present value (NPV) of savings is well over \$50 million calculated using a one-time reduction in inventory plus annual sav-




Figure 3: The total supply-chain costs for the five digital camera scenarios are expressed as efficient frontiers.




Figure 4: The postimplementation inventory levels drop below the optimized levels that would have occurred had the group stayed with Scenario 1, and the expected and actual inventory levels for Scenario 4 are closely matched. The preimplementation actual levels for Scenario 1 differ from the expected levels because the expected level for Scenario 1 was only a baseline and never implemented.

ings in variable and inventory-driven costs achieved through the adoption of Scenario 4. By the summer of 2003, HP realized more than \$35 million in savings.

Disseminating Proven Technology

The success of these optimization techniques within the Digital Camera business prompted a wider dissemination across HP. SPaM sponsored and deployed 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. By utilizing SPaM to disseminate these capabilities across HP, HP obtained three clear benefits. First, it established a shared horizontal process that spans various business units, and that allows them to leverage each other's work in supply-chain analysis using a common language. This is important because the supply chain of any product crosses multiple organizations inside and outside of HP. Second, the new offering still leverages SPaM knowledge and expertise, and the group is positioned to act as a central broker for knowledge the business units can tap into, helping them to share best practices. Third, SPaM was able to shift resources to other domains of expertise outside of the supply chain.

HP's supply-chain council sponsored and supported these activities. The council consists of the supply chain vice presidents across the company, and it sets and manages HP's supply-chain strategy and policies.

Inkjet Supplies: Modeling a Vertically Integrated Supply Chain

HP pioneered inkjet printing technology in 1984 and continues to lead the world market in this area. Today, HP'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.

The Inkjet Supplies business includes approximately 15 product families and over 250 manufacturing stock keeping units (SKUs). Unlike digital cameras, inkjet supplies have product life cycles in the tens of years, because they are manufactured to support a large installed printer base (over 150 million printers sold around the world since 1987).

To protect confidential information, we have altered the supply-chain maps in this section. We have not changed the qualitative and quantitative results from those results achieved in practice.

The Inkjet Supplies Challenge

The Inkjet Supplies network is very complex. HP-owned factories ship inkjet cartridges in bulk to regional completion centers for localization and packaging. These centers, called pen completion centers (PCCs), ship finished product to channel partners and to numerous HP distribution centers. The PCCs also ship cartridges to HP-owned printer postponement centers, which bundle printers and cartridges for sale. The supply chain supports two demand

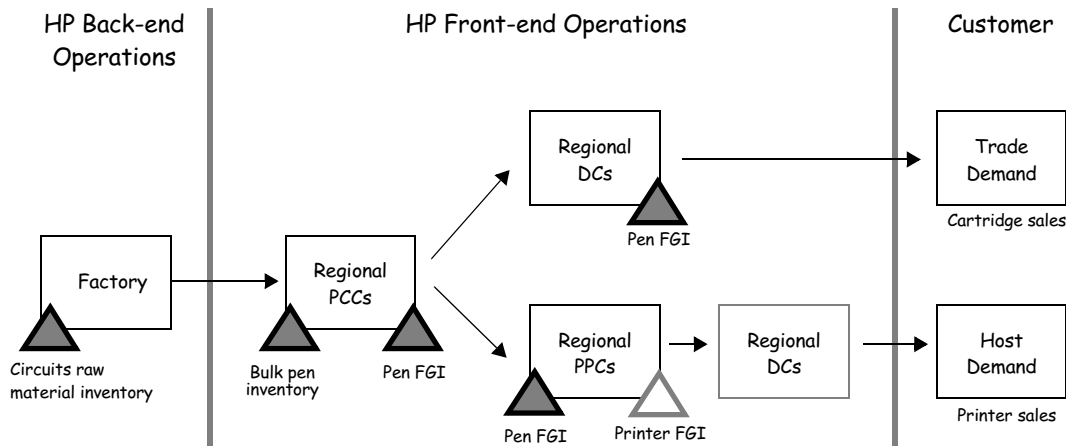


Figure 5: In the vertically integrated Inkjet Supplies business, HP's back-end operations ship bulk pens (inkjet cartridges) to regional pen completion centers (PCCs). After the PCCs localize and package the pens, they ship finished products in two streams. Replacement supplies (trade demand) for previously purchased printers go to regional distribution centers (DCs) and channel partners. Cartridges to be bundled with new printers (host demand) go to regional printer postponement centers (PPCs), which ship the bundled printer-cartridge products to regional DCs.

streams: demand from customers who already have printers, served by HP's channel partners and the HP distribution centers, and demand from printer sales, because all printers come bundled with a starter kit that typically includes an initial set of cartridges (Figure 5).

HP holds safety stock at various points in the supply chain to ensure high service levels. Bulk pen inventory at the PCCs helps decouple customer-facing operations from the long lead times of the factories and to buffer the system from the demand variability associated with the dozens of pen-packaging configurations HP offers. Bulk pen inventory is the largest safety-stock buffer in the system; by pooling unpackaged, bulk inventory at the PCCs, HP postpones packaging decisions so that it can respond to shifts in regional demand. After packaging at the regional PCCs, the pens are stored as finished goods inventory (FGI) and ready to distribute for sale.

Planners for the front-end supply chain knew that demand for bulk pens varied far less than demand for individual SKUs. While HP's existing inventory policy relied heavily on bulk pen inventory to respond to swings in demand for individual SKUs, the front-end organization wanted to know if HP would benefit by moving the bulk pen buffer, or some portion of it, further upstream to the factories. Because each factory supplied multiple PCCs, they reasoned that the supply chain might benefit from pooling demand uncertainty at the factory stage, before shipment to the regional PCCs.

Unfortunately, the inventory-modeling tools available to the front-end planning community could not

perform this evaluation. Using high-level assumptions concerning cycle time and variability, the tools could establish buffer strategy between any two locations (local optimization) but could not identify networkwide opportunities (global optimization).

To consider the issue of bulk-inventory placement, the Inkjet Supplies organization initiated the supplies inventory modeling (SIM) project with a twofold mission: (1) to investigate the benefits of holding bulk safety stock at the factories and (2) to identify other areas of opportunity as time allowed.

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 their analyses. At the end 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. The business stakeholders also could dismiss findings they disagree with when the models had no internal champion to defend them.

To address this common problem, the Inkjet Supplies organization assigned collective responsibility for the SIM project 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 who run the Inkjet Supplies business were the people analyzing the problem, building their own

modeling proficiency, and answering detailed questions posed by management.

PowerChain was instrumental in enabling this team's work, providing it with an intuitive means of accessing some advanced mathematical programming techniques it could not have used with traditional modeling tools. Furthermore, once they were trained, the geographically dispersed team members could collaboratively develop and manage models over the Web. Thus, their outputs were truly team deliverables and represented a common understanding among all the team members.

Marrying Process and Technology

The initial training session consisted of two days of formal instruction (Figure 6). Over half of this instruction was hands-on learning to build a model, with the remainder of the time spent on interactive lectures. For the first time in their careers, business analysts globally optimized supply chains they drew on computer screens.

On the third day, the team began 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 built much of the supply-chain structure. 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 important, the SIM team used a divide-and-conquer approach to global supply-chain modeling. Individuals built those portions of the supply chain for which they were responsible day to day. This approach differed from a traditional modeling effort in which one

or two expert modelers construct the entire model based on interviews with individuals. The visual nature of the tool helped the team to quickly address logic issues and flow and data questions between portions of the supply chain.

Next, the team had to populate the models with data. Of the overall two months of the project, three-quarters was devoted to gathering data. If that data had been readily available, model development could have been completed in less than two weeks. The team spent much of its effort taking existing data from HP systems and extracting enough detail to meet the granularity requirements of the process steps mapped in the supply chain.

To validate the model, the team verified the flow of logic and input data, such as stage costs, lead times, and demand. The team wanted to be sure that it correctly captured the process details, included the relevant cost drivers, and used the correct data. Model validation took a day or two.

Once the team finished its initial validation of the model, it immediately ran the model to optimize inventory levels and locations. This laid the foundation for the next step of comparing the initial model's solution against the business constraints.

Determining and introducing constraints at the various stages was the reality step in the process. The initial optimization run was only the starting point for the team's analysis. The SIM team was experienced enough to know that a model's optimal solution need not be the best to implement in practice. In reality, many near-optimal solutions that have slightly higher costs may either be far easier to implement or reflect operating constraints not factored into the original model. In this step, the team revisited the model's

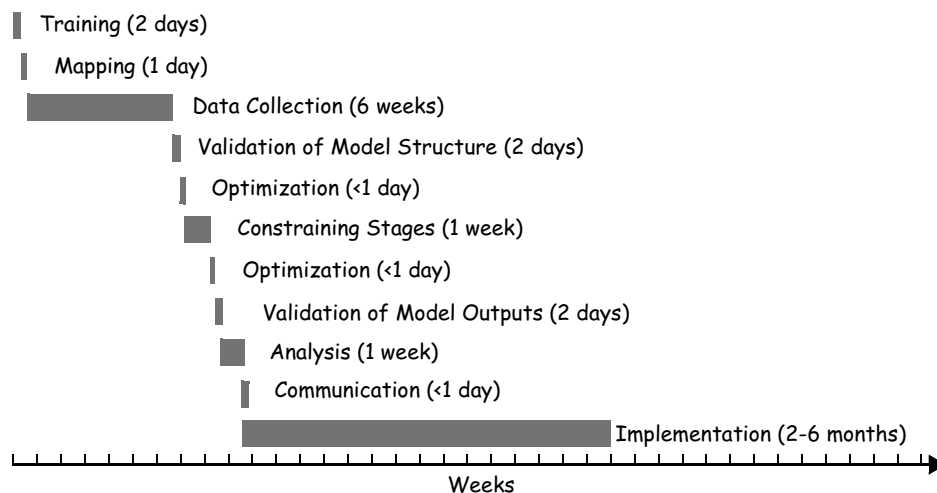


Figure 6: In the timeline for the process of designing a supply chain, data collection and implementation consume most of the time.

recommendations for the inventory policy and the supply chain to make sure they matched real-world conditions. For example, contract terms may or may not require a supplier to carry inventory or a certain work center may not actually be an ideal candidate for postponement. During this phase, the team also added existing supply-chain inventory policies to the model to provide a baseline comparison. The team did not have to go outside the model to capture these real-world constraints. It could just enter the relevant constraints for each stage and let the model figure out the best solution given these new constraints. Modifying the supply-chain model to accurately reflect reality took roughly one week.

With the new capability for optimizing across an entire supply chain, the SIM project went far beyond localized inventory calculations. For the first time, analysts could evaluate the best holding locations and the optimal quantities for those locations to meet customers' predefined requirements for service. Instead of performing a localized analysis that would likely suboptimize total system performance, the team could evaluate an entire system and, within that global context, determine the optimal inventory strategy.

The team performed a second validation step in two phases. First, it revisited the assumptions in the model and compared the overall supply-chain costs with aggregate financial data from HP systems. Second, HP's Finance group verified the current-state model against actual costs by comparing the model's outputs with actual booked values. This crucial validation step showed that the model's supply-chain costs were very close to projected costs, not only as an aggregate number but at the per-unit SKU level. The modeled and projected per-unit costs were within a penny of one another. The similarity in cost outputs 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.

The team then set out to interpret the results of its global inventory modeling effort and to do this within a week. In the past, it would have treated each stage in isolation, assuming that each stage held a decoupling safety stock. Because of the SIM project, the team could model the entire end-to-end supply chain (Figure 7).

The results were very surprising. Contrary to the team's expectations, the optimization did not show a benefit from holding bulk pen inventory at the factory. Instead, it suggested that the system would be better served if HP held safety stock in forms other than bulk pens altogether. Specifically, it suggested increasing the amount of circuits raw-material inventory in front of the factory and increasing the amount

of packaged finished goods inventory near the customers. The team might never have made their finding had it been forced to model only a small number of supply-chain stages explicitly.

With the original bulk-inventory question answered, the team used its remaining time to evaluate other possibilities. In fact, the team had to restrain itself 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 people who were responsible for day-to-day operations of the supply chain made a discovery. To construct a global model, they had to detail the costs and times of individual freight lanes within the model. If the team were to move transoceanic freight lanes from air to sea, they calculated that HP could realize an NPV greater than \$80 million. 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.

Moving shipments from air to sea seemed completely counterintuitive. Traditional rules of thumb had focused on responsiveness through speed. HP saw air shipment as a way to help minimize both cycle times and inventory levels. Also, the planning granularity of the current inventory-management system and a highly variable demand signal made increasing transportation time seem risky. Without such countermeasures as holding inventory, HP risked harming end-customer availability. The model showed that by sizing the inventory at the DCs appropriately, HP could avoid expensive air-freight charges without changing its service to customers.

The project team would have missed many of these insights if it had focused on a localized modeling approach centered around bulk inventory of pens with outcomes constrained to only two choices (hold bulk at PCCs or hold bulk at factories). In this exercise, the team answered the bulk-inventory question and, in addition, identified other opportunities based upon insights gathered from using systemwide optimization and the team's detailed knowledge of the day-to-day supply-chain operations.

The purpose of communicating results is to influence change. From a change-management perspective, the analysis stage is the easy part. Even a recommendation that sounds simple, switching from air to sea, can take tremendous effort to implement, particularly when the business is large. When the arguments in favor of the change are largely based on arcane-sounding modeling exercises performed by




Figure 7: The Inkjet Supplies organization initially asked the team to determine whether it was optimal to shift inventory of bulk pens from the PCCs back to the factories. With the end-to-end formulation of the supply chain, the team quickly determined this was not a cost-effective change. Instead, it was optimal to increase the amount of circuits raw-material inventory and the finished goods inventory near the customer. Next, the team identified significant savings from shipping by sea versus air between the factories and the PCCs.

Note. The text associated with each stage is intentionally blurred to protect confidential information.

outsiders using black-box tools, selling the idea can be a Herculean task.

A graphical modeling tool that the change agents themselves could use and understand helped them to accept and implement the changes. In this case, the team was trying to persuade the organization to change processes that had been in place for years and were ingrained conventional wisdom. Their work also challenged the notion that the Inkjet Supplies business was too large for the application of management science. The graphical nature of the tool was a key communication aid during presenta-

tions and working sessions with various management tiers. Managers could see the structures of their supply chains clearly in a way that hadn't been possible before. With a picture to look at, stakeholders felt 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.

In an environment where senior management had requested a reduction in overall inventory, the Inkjet Supplies team instead successfully influenced stake-

holders to increase inventory in the supply chain to reduce total supply-chain costs!

For implementation to succeed, all stakeholders must understand the analysis, and they did because of the portability and standardization of the analysis techniques. Furthermore, because the people who would do the implementation did the modeling, they did not need to be convinced of the results. The stakeholders could all agree on what the supply-chain characteristics were without resorting 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 team finalized and implemented the plan two months after the SIM project's completion. HP targeted several products to move to sea shipment. For others, switching to sea transport did not make sense for such reasons as product economics or a short shelf life. For those that can be shipped by sea (the population of interest), the team took several considerations into account. Engineering had to confirm that the products' quality would not be harmed by any aspect of sea transportation. Manufacturing needed enough capacity to produce the extra units required to fill the new pipeline and still keep regional buffer levels at target. To address manufacturing's concern, HP modified factory schedules with a nine-month transition plan.

Within six months of the implementation's start date, more than 90 percent of the units targeted for ocean freight were indeed on the ocean, giving the business an annuity stream in excess of \$20 million per year, which goes directly to the bottom line. With our systems approach, we did not implement localized solutions that had unreported negative consequences for HP's partners who might eventually push back or be forced to end the business relationship. The model considered the overall impact on the end-to-end supply chain, including partners, thus reducing the possibility of consequences that would later undermine any reported gains.

For the Inkjet Supplies exercise, the team quickly discovered and implemented opportunities beyond the one it was asked to evaluate. After deciding where to hold inventory, it quickly discovered the air-versus-sea transportation opportunity and began to implement it almost immediately.

This project highlighted the benefits of coupling a global inventory optimization algorithm with an intuitive interface that allowed the real-life business users to conduct modeling exercises rapidly. The in-house team identified opportunities in their field of expertise that an external consultant might never have uncovered.

Lessons

We discussed two business cases to illustrate the application of a methodology in two very different business environments and to demonstrate the portability of these techniques within diverse organizations. However, HP now uses these techniques in areas beyond these two business units. It uses this modeling method for such product lines as flatbed scanners, inkjet printers, servers, and laser printers. The user community includes analysts in HP regions all over the world (Figure 8).

The techniques are currently disseminated on an as-requested basis. Through 2002, we had conducted more than 10 training sessions around the globe for over 120 people. Clearly, companies need standardized, portable inventory-optimization techniques, and HP now has a portable tool that business users can use in concert with HP's internally developed process for decision making in the supply chain.

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.

Between January 2001 and January 2003, HP used PowerChain to model over 1,500 supply chains. In early 2001, the pilot study in the Digital Camera business and the initial application in Inkjet Supplies were each creating multiple supply chains to represent various supply-chain configurations. Because these early teams were so successful, the adoption process was essentially viral. During 2001 the use grew on an annual basis of 200 unique chains modeled per year (this is the number of original maps created, not including revised versions of those maps). While even this growth rate was unprecedented within HP, the catalyst for exponential growth, now 1,400 chains per year, was the implementation of a formal training class sponsored by HP Corporate Education. This



Figure 8: The PowerChain user community includes HP facilities in Boise, Idaho; Corvallis, Oregon; Fort Collins, Colorado; Houston, Texas; Palo Alto, California; Roseville, California; San Diego, California; Vancouver, Washington; Guadalajara, Mexico; San Juan, Puerto Rico; Grenoble, France; Boeblingen, Germany; and Singapore.

training class exposed new users to the business process and tools and equipped teams with the process knowledge to improve their supply chains.

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 do not report directly to them. Achieving buy-in across the supply chain from a position of nonauthority is very difficult.

To make changes and obtain the cooperation of diverse groups and individuals, all the players must understand the details and trade-offs. Without the ability to determine and validate the total supply-chain cost, one cannot possibly achieve a majority buy-in. While sophisticated business people recognize that optimizing safety-stock levels can reduce total supply-chain costs, they also know that determining an 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 between analysis and implementation. The entire user community can determine global inventory levels within a structured framework. This translates into the following benefits:

—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 demonstrate the behavior of the supply chain to nonmodelers.

—Stakeholders from across the organization can use the same framework in discussing the supply chain, which is critical for holistic analysis. Users can see the trade-offs between local optimization and global optimization and persuade others within their local divisions to buy into their recommendations. By using an algorithm that is flexible enough to solve what they can draw, users get the benefit of a picture that is worth a thousand spreadsheets in communicating their findings to a broader audience.

—The ability for each user to add his or her part of the supply chain greatly aids the decision-making process. It is difficult for one analyst from a particular business operation to capture the entire network; it is equally difficult for analysts from two business operations to combine their respective custom-built models. By creating a modeling framework that allows for decentralized management and control, the team removed a barrier to creating, maintaining, and interpreting an end-to-end model. This framework greatly facilitates the implementation process.

—One of the unforeseen benefits of using this tool is that each business user can input his or her own

Cumulative number of unique chains modeled




Figure 9: For 2001, supply chains were being modeled at a creation rate of 200 supply-chain maps per year, which was unprecedented adoption by business users at HP. At the start of 2002, HP Corporate Education began sponsoring training on the business processes outlined in this paper, creating an order-of-magnitude utilization increase to a new rate of 1,400 supply chains per year.

portion of the data. In other situations, modelers tend to model supply chains using data that is readily available to them. Many times, they miss or cannot get pertinent data that might be much more available to users in the field.

Conclusion

HP has benefited from a standardized process to optimize total supply-chain cost. HP uses a single process and optimization technology coupled together to obtain benefits in radically different business situations and to identify different opportunities.

Large numbers of business analysts can now model a supply chain together, make global decisions together, conduct sensitivity analyses together, and make unified recommendations to a cross-organizational management team based on their efforts. Seeing these unified recommendations, managers throughout HP tend to quickly authorize implementation. Stakeholders from across the supply chain understand the global impact of their inventory decisions, and HP's implementation teams 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. They no longer reside 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 process and a framework that helps people to communicate, collaborate, model, and understand their supply chains. This new combination of technical innovation enabling broad-based organizational decision making is noteworthy, and we believe that it is exactly what Bill Hewlett and David Packard would expect from the employees of HP.

Appendix: PowerChain Intuition and Algorithm

The authors of several excellent supply-chain papers have competed for the Edelman award in previous years. In particular, Arntzen et al. (1995) and Camm et al. (1997) considered network design problems. Lin et al. (2000) considered multiechelon inventory optimization. In our work, we continued this stream of work in focusing on developing a pragmatic solution to optimizing inventory levels and locations in a supply chain. The success of PowerChain rests on its ability to marry supply-chain optimization algorithms and an intuitive application environment to support the decisions planners make every day.

Problem Complexity

In even the simplest multiechelon supply chain, the problem of choosing inventory-stocking locations and the target levels at those locations has exponential complexity. Consider a supply chain with five potential stocking locations in series, and suppose we are to decide only where to buffer safety stock. Assume buffering at a location means the location fully satisfies demand from inventory, whereas not buffering at a location means the location never satisfies demand from inventory. In a five-stage serial supply chain, there are 32 possible buffering policies. While complete enumeration is feasible for this stylized example, real-world supply chains have hundreds to thousands of stages; furthermore, once a stocking location is known, the actual inventory level to maintain at that location must still be determined.

Traditionally and with pre-existing tools, analysts leverage experience and intuition to reduce a priori the possible sets of locations to make their analyses manageable. In our five-stage supply chain (Figure 10), a typical intuitive reduction might result in three scenarios for consideration: stock everywhere, stock at the final stage only, and stock at the final stage and at some upstream stage.

In the stock-everywhere scenario, each stage fully buffers. This is a common scenario and is typically




Figure 10: For a five-stage serial supply chain, there are 32 possible decoupling inventory policies. This figure shows three of them.

the output of single-stage solutions, because buffering effectively decouples the echelons. The common downside of buffering everywhere is that variability of lead times cannot be pooled across stages, and we have found in practice that maintaining many stocking locations increases the risk of mismanaged (in particular too much) inventory throughout the supply chain.

In the stock-at-final-stage-only scenario, the final stage buffers the entire supply chain, while upstream stages have no buffering. The problem is that, at the final stage, the product held in inventory is fully differentiated, so inventory is its most expensive at this point, making the cost of maintaining high service prohibitively expensive.

In the stock-at-final-stage-and-some-upstream-stage scenario, the final stage holds safety stock to satisfy customer demand, but this inventory is not buffering for the entire supply chain, only part of it. One of the upstream stages also buffers, providing an intermediate decoupling point. This scenario might be an appropriate solution, and would certainly be considered by an analyst, if one of the stages has a very high processing cost. Buffering before this high-cost stage and not after it might perform better than either of the other two scenarios because the expensive, processed inventory would be minimized.

The challenges of relying on experience and intuition are manifold. It is unlikely that the necessarily partial analysis yields a true optimal solution. Without the ability to justify a solution in financial terms, one would have difficulty obtaining support from key stakeholders, especially when they must make personal sacrifices for a global improvement. More thorough analyses require substantially more time. Therefore, what we need is a method of determining globally optimal stocking policies in a reasonable amount of time, while providing visibility to stakeholders.

Algorithmic Formulation

In this section, we outline the algorithmic features of the model. Graves and Willems (2000) and Humair and Willems (2003) provide more detailed descriptions.

We model a supply chain as a network where nodes are stages in the supply chain and arcs denote that an upstream stage supplies a downstream stage. A stage represents a major processing function in the supply chain, such as procurement of a raw material, production of a component, manufacture of a subassembly, assembly and test of a finished good, or transportation of a finished product from a central DC to a regional warehouse. Each stage is a potential location for holding a safety-stock inventory of the item processed at that stage. Let the directed, acyclic, connected network $G = (V, A)$ represent the stages and

arcs of the supply chain, where V is the set of stages, indexed by i , and A is the set of all arcs (i, j) . The set V_d is the set of demand stages, such that stage k is in V_d if there are no arcs (k, j) in A .

Each stage has a stage time, T_i , which is the time from when all of the inputs are available until production is completed and available to serve demand, including waiting time, processing time, and transportation time to put the item into inventory.

The decision variables are, for each stage, S_i and SI_i . S_i is the service time quoted by stage i to each of its downstream stages, while SI_i is the maximum service time being quoted to stage i by its upstream stages. For each of the demand stages in V_d , there is a maximum service time bound that can be quoted to customers, represented by b_i .

Cost is captured arbitrarily by a real function $c_i(S_i, SI_i)$ for each stage i , representing the holding cost of inventory at each stage, given the outgoing and maximum incoming service times.

The objective function is therefore

$$\min \sum_{i=1}^N c_i(S_i, SI_i).$$

The feasible region is constrained as follows. First, all service times are nonnegative and integral:

$$S_j, SI_j \geq 0 \quad \text{and integer for } j = 1, 2, \dots, N.$$

The maximum incoming service time to a stage j is greater than or equal to all of the outgoing service times of that stage's upstream stages:

$$SI_j - S_i \geq 0 \quad \text{for all } (i, j) \in A.$$

The service times may be bounded arbitrarily by the user at every stage (L_j and U_j are the lower and upper bounds the user specifies):

$$0 \leq L_j \leq S_j \leq U_j \quad \text{for } j = 1, 2, \dots, N.$$

The expression $SI_j + T_j - S_j$, named the *net replenishment lead time*, is the time of exposure over which the base stock must cover demand. It is the time required to replenish an item at a stage (the maximum incoming service time plus the stage's processing time), netted by the service time being quoted by the stage. (If the stage quotes a service time equal to its replenishment time, it needs no safety stock because all orders will be satisfied after one full replenishment cycle; this corresponds to a net replenishment lead time of zero.) The net replenishment lead time can be arbitrarily bounded by the user at every stage to reflect situations in which stages are either not allowed to hold inventory or, by contractual terms, are required to hold a certain amount of inventory:

$$NL_j \leq SI_j + T_j - S_j \leq NU_j \quad \text{for } j = 1, 2, \dots, N.$$

The cost functions $c_i(S_i, SI_i)$ are obtained from the inventory dynamics. Graves and Willems (2000) show the cost function to be

$$c_i(S_i, SI_i) = h_i \{ D_i(SI_i + T_i - S_i) - (SI_i + T_i - S_i) \mu_i \},$$

where h_j is the holding cost of inventory at each stage, $D_j(x)$ is the maximum possible demand seen by the stage over the time interval of length x , and μ_j is the average demand in one time period. Extensions to the formulation to account for stage-time variability and production-capacity constraints can be captured by more complex treatments of the $c_i(S_i, SI_i)$ functions. In even the simplest forms, these are general nonlinear functions, so that this problem is nonlinear with integral decision variables. We have developed dynamic programming techniques, described by Graves and Willems (2000) and Humair and Willems (2003) for solving this problem.

Architecture

PowerChain is a collaborative application for the interactive visualization, design, and optimization of supply chains. Using Web-based technologies in a multiuser environment, the software encompasses the entire modeling life cycle, including model creation, supply-chain structure mapping, data loading, optimization sensitivity analysis, model versioning, and workflow. A typical project will analyze one or more supply-chain scenarios (such as the five digital camera options we presented). Within the 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. 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 go-to factor from the source stage to the destination stage.

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, multiuser, distributed nature of

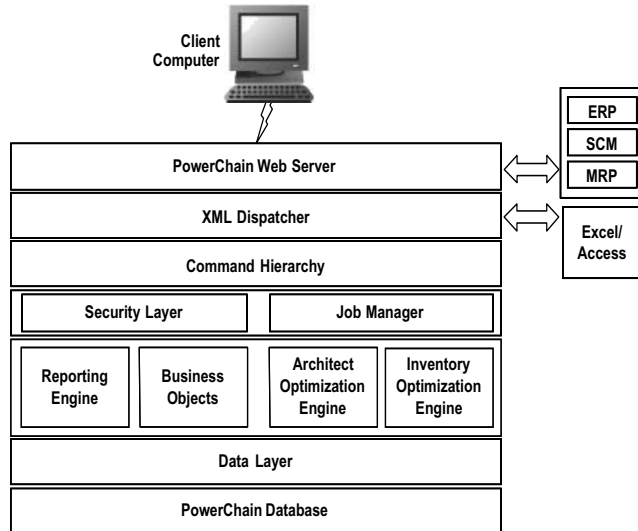


Figure 11: A user connects to the PowerChain application using a Web browser and draws a supply chain using a Java applet. The Web server generates requests for performing supply-chain actions using an XML schema specifically designed for supply-chain modeling. These XML requests are dispatched and turned into commands, such as “Add Stage” or “Optimize.” An asynchronous job manager handles complicated and time-consuming actions, such as optimization, while business objects, reporting, and security layers handle traditional application-tier activities. All of these operations rely on a relational database.

the application, the path from drawing a supply chain to optimizing inventory policies is hardly direct (Figure 11).

Web technologies enable users to come up to speed quickly while minimizing the total cost of ownership. New users can be added to the system and brought online with minimal IT investment as soon as they are trained on the software.

References

- Arntzen, B. C., G. G. Brown, T. P. Harrison, T. L. Linda. 1995. Global supply-chain management at Digital Equipment Corporation. *Interfaces* 25(1) 69–93.
- Camm, J. D., T. E. Chorman, F. A. Dill, J. R. Evans, D. J. Sweeney, G. W. Wegryn. 1997. Blending OR/MS, judgment, and GIS: Restructuring P&G's supply chain. *Interfaces* 27(1) 69–93.
- Graves, S. C., S. P. Willems. 2000. Optimizing strategic safety stock placement in supply chains. *Manufacturing Service Oper. Management* 2(1) 68–83.
- Graves, S. C., S. P. Willems. 2002. Strategic inventory placement in supply chains: Nonstationary demand. Working paper, Boston University, Boston, MA.
- Humair, S., S. P. Willems. 2003. Optimizing strategic safety stock placement in supply chains with clusters of commonality. Working paper, Boston University, Boston, MA.
- Lin, G., M. Ettl, S. Buckley, S. Bagchi, D. Yao, B. Naccarato, R. Allan, K. Kim, L. Koenig. 2000. Extended enterprise supply-chain management at IBM personal systems group and other divisions. *Interfaces* 30(1) 7–25.

Jeff Clarke, Executive Vice President, Global Operations, Hewlett-Packard, writes: “HP’s supply chains cross multiple organizations and companies. The implementation of these supply-chain optimization techniques and tools, leveraged across multiple organizations, delivers on this need for the supply-chain development community. This work highlights the power of standardization and knowledge leverage, and shows HP’s continual leadership in supply-chain innovation.”

Vyomesh Joshi, Executive Vice President, Imaging and Printing Group, Hewlett-Packard, writes: “IPG’s goals in operational excellence require us to optimize costs, accelerate time-to-market, and maximize flexibility while delivering industry leading customer experience. This initiative successfully defined and implemented tools that allow IPG businesses to model and analyze the end-to-end supply chain tackling these goals. Through global collaboration, the Digital Imaging and Inkjet Supplies businesses were able to implement supply-chain improvements yielding over \$130 million in benefits to IPG.”

Mary Peery, Senior Vice President, IPG Digital Imaging and Publishing Organization, Hewlett-Packard, writes: “Digital Imaging is a significant growth business for HP, with a goal of providing image rich conversation seamlessly between people. Digital cameras are a core element of this strategy to bring affordable solutions to market. By reducing the camera total supply-chain costs by more than 5% and improving our time to market by more than two weeks on every camera, our customers have access to the latest technologies at an affordable price. This is not only a win for Digital Imaging, but our customers as well. These supply-chain analysis techniques have now been adopted by our flatbed scanner and photosmart inkjet printer product lines following the success in our digital camera business.”

Boyd Lyon, Enterprise Supply-Chain Manager, Inkjet Supplies Business, Hewlett-Packard, writes: “Within the Inkjet Supplies Business, it is well recognized that our success continues to depend on our ability to compete as an integrated network. We are challenged with continued pressures to reduce inventories, supply-chain complexity, and product costs while simultaneously breaking new ground in creating new, disruptive customer advantage. These tools and methodologies are enhancing our understanding of the broader systems we manage, and providing new tools for shaping action across a distributed operations network. We are able to ask the right questions, focus on the right opportunities, and speed the process of recognizing benefits.”