Interfaces
Vol. 35, No. 3, May–June 2005, pp. 238–247
ISSN 0092-2102 | EISSN 1526-551X | 05 | 3503 | 0238

DOI 10.1287/inte.1050.0143

Hewlett-Packard Combined OR and Expert Knowledge to Design Its Supply Chains

Claude Laval, Marc Feyhl, Steve Kakouros
Strategic Planning and Modeling, Hewlett-Packard, MS 59, 5 Avenue Raymond Chanas, 38053 Grenoble-Cedex 9, France
{claude.laval@hp.com, marc.feyhl@hp.com, steve.kakouros@hp.com}

Although one can eliminate human bias from scenario definitions when modeling supply chain networks by using a pure operations research approach, the resources, data, and time needed for this approach may prohibit its use in fast-paced business environments. Hewlett-Packard's (HP's) strategic planning and modeling (SPaM) team refined an approach using true optimization coupled with scenario analysis, selectively focused on the most critical parameters of the design equation, and used this approach to solve a complex supply chain design problem for HP's imaging and printing group. The clear and easily demonstrable results helped management to understand their derivation and act confidently. The combined approach gave equal emphasis to optimization and to scenario analysis.

Key words: facilities—equipment planning: location; industries: computer, electronic.
History: This paper was refereed.

Since 1994, with assistance from the in-house strategic planning and modeling (SPaM) consulting team, supply chain groups at Hewlett-Packard (HP) Corporation have been applying quantitative scenario-based analysis to supply chain design problems (Lee and Billington 1995). SPaM is a small team of operations research (OR) specialists within HP who provide internal support to HP product divisions to improve their efficiency, cost-effectiveness, and profitability. SPaM has developed a variety of OR modeling techniques and adapted them to specific supply chains, improving their performance (Cargille et al. 1999). The authors of this paper are members of the SPaM team and represent SPaM's activities in Europe, the Middle East, and Africa (EMEA).

HP has replaced intuitive approaches with analyses based on hard data and modeling alternative scenarios to support its supply chain decisions. The project sponsors and team still base the preliminary definitions and assumptions in the scenarios on their intuition and expertise. For problems of a well-known type, intuition is a good guide in defining initial scenarios. For green-field situations in which existing operational constraints are not considered and in situations in which analysts must think outside of the box, however, intuition based on past work may bias the overall approach and reduce the probability of reaching a truly optimal solution, especially when analysts believe that the current situation or base scenario is already close to optimal.

By using a mathematical approach, close to classical OR techniques, one could employ a mathematical formulation together with sophisticated optimization techniques to better define scenarios. However, the arcane formulations, human resources, data, and time required for a pure OR approach can prohibit its use or restrict its benefits (Murphy 1998). Furthermore, many of the options within the ranges to be considered may turn out to be peripheral, thus wasting cycles or iterations on the evaluation of noncritical factors.

We developed a systematic way of combining expert knowledge and intuition for scenario-based modeling with pure optimization techniques to solve complex problems in supply chain network design. In developing this approach, we drew on SPaM's expertise in supply chain modeling and used simple optimization tools that can run on a PC.

Defining the Problem

HP's imaging and printing group (IPG) manufactures and markets HP's printing and digital imaging products worldwide. Within IPG, the regional EMEA supply chain organization is responsible for distribution and manufacturing activities in this region—a multibillion dollar operation.

To improve its operations, IPG wanted to reduce the number of its contract manufacturing (CM) partners in the EMEA region. These partners postpone and distribute products built in Asia or Eastern Europe. To eliminate partners, IPG would have to decide how to reassign the affected volume of postponement activities (Figure 1). Could the remaining product completion centers (PCCs) handle the existing volumes of postponement? Were these PCCs located correctly? If EMEA needed new sites, where should it locate them?

In its initial discussions with the EMEA managers sponsoring the project, SPaM proposed a green-field approach, in which it would consider an unrestricted selection of existing and new sites without considering the cost of moving, opening, closing, or changing facilities. By using this approach, it would avoid the biases that would otherwise exist for the current network design. It would then consider costs of moving to the new network at a later stage in the project. The objective was to minimize supply chain costs using two alternative specifications for response times and make recommendations that would help EMEA to

Figure 1: Some of HP's imaging and printing group potential postponement sites in Europe.

determine the optimal number of postponement facilities and their location for each product, to identify potential leveraging opportunities across product categories through pooling common postponement locations, and to identify which unique facility should serve each regional market for products processed at multiple locations.

Defining the Data for the EMEA Imaging and Printing Group Network

The SPaM team working on this project had to take great care in defining the data parameters and their granularity.

The potential data combinations were huge. IPG had thousands of ship-to points (demand areas) all over EMEA. The product portfolio included several thousand items. Furthermore, all cities over a certain size in EMEA were potential locations for postponement activities. We needed the following data:

—Transportation costs and transit times for all combinations of postponement location and demand area,

—Demand by product for each demand area, and

—Handling and transformation costs for each postponement location by product.

For facilities that already handled a particular product, data existed and was usually easy to collect. Nevertheless, the granularity of the existing financial and logistics reports that served as the sources of the data necessitated intensive review and, in some cases, further analysis. We could usually extrapolate data for products not handled at an existing facility from the existing data for that facility (Figure 2).

Obtaining data for potential new locations presented the most difficulty. We needed complete consistency in our assumptions about processes among locations and products. To overcome these problems, we used various data sources, financial databases, contracts with existing partners, databases from past projects, quotations from partners, and publicly available comparisons of cost factors (International Institute for Management Development 2003).

Data Granularity and Sensitivity

In modeling projects, one must reduce the complexity of the real world to provide meaningful and timely results. For our project, we had to reduce

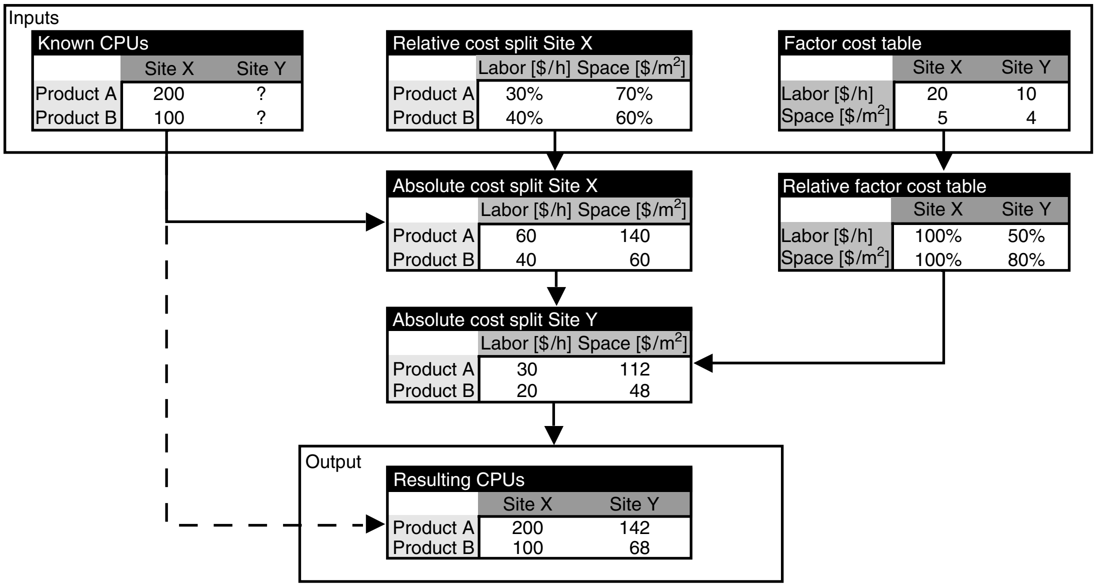

Figure 2: In this simplified example of our extrapolation technique applied to manufacturing overhead (MOH), we have two variable cost types, two facilities (sites), and two products. In our actual project, we used seven variable-cost types, seven fixed-cost types, 10 locations, and 20 products. We used this extrapolation technique to determine fixed and variable costs for manufacturing and completing products at all locations in a consistent manner. We measured variable costs as cost per unit (CPU). We assumed that we know the variable costs—CPUs—for products A and B at site X. We want to know the costs at site Y (Known CPUs). Furthermore, we assumed that all costs for products originate in the factors' labor and floor space. We know the split of the CPU of a product into these two factors for Site X (relative cost split Site X). We also know the factor costs for both cost types and sites (factor cost table). Based on multiplying the known CPUs at site X and the cost split at site X, we can calculate the cost split per product and factor at site X in absolute terms (absolute cost split site X). By taking the factor costs at site X as 100 percent—the reference—we calculate the factor costs at site Y relative to site X (relative factor cost table). By multiplying the absolute costs per factor and product at site X with the relative factor costs at site Y, we can calculate the absolute costs per factor and product at site Y (absolute factor cost table site Y). The final result, we can calculate the CPUs per product at site Y by adding the factor costs per product (resulting CPUs).

all four dimensions—the postponement location, the demand area, the distribution method, and the product category—to manageable ranges:

First, we chose a reasonable number of potential postponement locations to limit our data-collection efforts, which would have been huge for locations with no current HP operations. Therefore, the upper bound on the postponement dimension was driven by the feasibility of obtaining data. On the other hand, choosing too few potential postponement locations would prevent us from reaching true optimum because we might drop an optimal location from the set of scenario alternatives. We grouped similar locations and chose a representative one. To do this selection with ultimate rigor, we would have collected data for all locations. Instead, we used existing data exclusively and relied on internal manufacturing and logistics experts to choose the best representative locations. We also eliminated several locations by testing all the locations against qualitative must criteria (mostly concerned with HP outsourcing policies). In the final model, we had 10 alternative postponement locations, five of which had existing facilities. We included all existing locations regardless of their attractiveness to facilitate calculation of a base case. We also wanted to see if the existing facilities were by chance ideally located and if not, to compare them with the ideal locations.

The granularity of demand areas would determine the accuracy of transportation costs and therefore the accuracy of our calculations of the implicit center of gravity (Bowerman et al. 1997). It also affected the calculation of average transit times, an important constraint in our model. For each country, we ran a side analysis to determine whether it made sense to group it with neighboring countries, to consider it one demand area by itself, or to split it into several demand areas. We found that 20 demand areas, each represented by a major city within the region, gave the best trade-off between accuracy and effort.

We introduced the following alternatives for the distribution methods:

(1) Indirect: Products shipped to the warehouses of channel partners.

(2) Replenishment: Products shipped to an intermediary stocking point serving one or more countries.

(3) Direct: Products shipped directly to end users.

Because each of these distribution methods applies to all 20 geographical demand areas, we had to consider 60 demand areas. For modeling purposes, we reduced the properties of the three distribution methods to different freight rates that depended on individual transit-time constraints and shipment profiles.

For the analysis, we differentiated products using the following parameters:

—Product weight, which drove transportation costs, and

—The processes required to transform or complete the product, which determined labor costs, space costs, equipment costs, and overhead costs.

Among HP's printing and imaging products, these parameters could vary by a factor of 1,000. Besides these quantitative differences, we considered qualitative factors, such as the possibility of sharing manufacturing equipment or capabilities. We concluded that 20 product categories offered the best trade-off for the problem.

The Model

We developed an approach that combines optimization and scenario analysis to produce the various alternatives. The HP managers wanted to understand how we arrived at our proposed alternatives so that they could make a decision with confidence. We had to bridge the gap between their knowledge and the technical sophistication of the model.

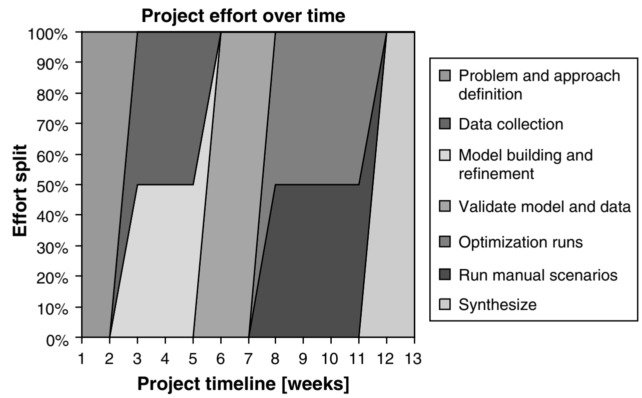

Figure 3: The project timeline shows that the combined approach accelerated the optimization and scenario runs in the analysis phase between weeks 7 and 11.

We spent two months collecting data and building the model and another month or two running the different scenarios on the model and analyzing the results (Figure 3). Running each iteration of the model took approximately five minutes.

Before we developed the combined approach, we would have modeled the problem using either optimization or scenario-based modeling.

Optimization Modeling

Using optimization, the analyst feeds data and constraints into a modeling program and lets the program search for the best combination out of a range of possibilities (Cohen 1989). The enthusiastic OR practitioner might be tempted to consider a full-scale optimization with rich dimensions and a highly detailed objective function. However, each added dimension (even each added detail) drastically increases the data required and therefore the project team's work.

Scenario-Based Modeling

Using a scenario-based approach, one analyzes a limited set of scenarios. In traditional scenario analysis, the analyst uses scenarios that describe alternative instances of a future world under some hypotheses of improvement (McBurney and Parsons 2002, 2003). Each scenario represents a complete picture of the world, with its own set of input data. For our problem, each scenario would consist of a set of postponement locations, the product groups handled at each postponement location, and the demand areas for each product group served by each postponement location.

After the scenarios have been identified, a model is developed, typically spreadsheet-based. For every scenario, the model calculates the overall performance in terms of cost and transit time. A good scenario-based model can provide all sorts of breakdowns by manufacturing location and demand area. The level of detail in the cost analysis depends on the granularity of the input data.

Scenario-based models can even use a stochastic or random logic to reflect uncertainties in demand or transit time that affect inventory levels and delivery performance. This approach is rarely limited by the model but rather by the predefined scenarios. The expertise of the project team and sponsors determines the scenarios' definition. Scenario-based models work well in a stable world but become brittle when the objective is to explore uncharted territory. The temptation is to consider only scenarios that fit the stakeholders' mental models. This bias can be inadvertent. Even worse, intentional biases can result from political influence, for example, partner-specific constraints can be introduced in the scenario definition.

Optimization and Scenario-Based Modeling

We compared scenario-based and optimization-based approaches:

—Scenario-based approaches identify the best scenario out of a predefined set of scenarios. Optimization produces an optimal solution within defined constraints based on the objective function.

—Scenario-based approaches evaluate only the scenarios the project team manually defines. Optimization automatically outputs the optimal scenario but gives limited information about other scenarios that may be almost as good.

—Scenario-based analysis is good when the analyst knows the approximate solution and needs to fit it in real-world requirements (for example, after optimization). However, scenario-based analysis does not necessarily provide an approximately optimal solution. Optimization is good for creating new solutions or thinking-out-of-the-box solutions and for narrowing down the scope for further analysis using scenario-based methods.

Combined Modeling

The idea of a combined modeling approach is not new. Camm et al. (1997, p. 139) describe a "...hybrid approach...that closely links expert human judgment and mathematical optimization." Their scenario analysis was a manual validation and verification of the results obtained by distribution and product-sourcing optimization models. They then displayed the final results within a geographical information system (GIS).

The effort of Camm et al. for Procter and Gamble (P&G) involved a team of 500 to develop and set up the model. A core team of 10, including two members of the SPaM team developed the IPG model; the extended team for the model included a further 20 people. P&G's larger scale model focused on the North American supply operation and featured a more fractured approach. The P&G team optimized product sourcing and distribution as separate models, and once they had manually verified these models using scenario and risk analysis, they assumed an overall optimal state. The P&G team used proprietary software for its optimization model and mapped the results in a GIS for human validation. For the SPaM optimization model, we used a public-domain solver. We formulated the scenario model in Excel with a graphical display. The key personnel in the supply chain team could easily understand the results. The SPaM model had a great advantage: we could break it down into different transition steps, which allowed us to clearly identify the optimum transition-path cost scenario. Also, it was a simple task to run the model again to check the different alternatives produced through optimization.

We gave equal weight to optimization and scenario modeling in the SPaM model. Crucially, HP managers did not see the scenario analysis as part of the backup or validation process but as an integral part of delivering the final result.

Communicating the Approach to Non-OR Practitioners

Ultimately a management team most of whose members were not OR practitioners had to make a decision based on our analysis. Those outside the profession cannot easily grasp some of the combinatorial aspects

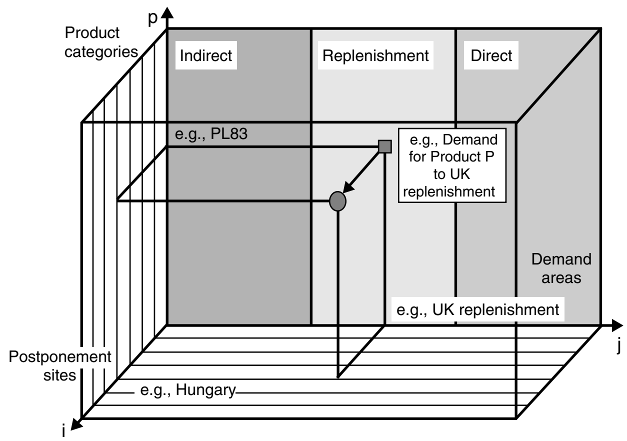

Figure 4: This type of visualization of a decision variable that assigns a product in a demand area in a business model to a postponement facility was helpful in communicating optimization techniques to HP management.

of mathematical programming. We had to explain the model to the project team and the sponsors to ensure that they understood and accepted the proposed solution. OR optimization techniques are not really new at HP, but we had to explain the feasible scenario space and obtain management's support for the mathematical-programming approach (see Figure 4).

Benefits

SPaM's approach was based on optimization to generate scenarios combined with traditional analysis and modeling applied to the scenarios. We were guided by our experience in previous projects in framing the business problem to keep data requirements, time, and cost within reasonable limits. In our example case, we defined three key dimensions: 20 product categories, 10 possible postponement (PCC) locations, and 60 demand areas.

We relied on our supply chain expertise in determining the factors to be considered. In some cases, it made sense to narrow the optimization to a portion of the cost structure. In this project, we focused the optimization on the trade-off between manufacturing and freight costs, and left aside inventory-driven costs, at least in the first phase of the project, and only re-integrated them later as a side analysis.

SPaM developed a medium-scale relaxed mixed-integer program as the first optimization model tailored to the business problem (appendix). We solved the program using a public-domain optimizer. Up to this point, HP had not commonly used optimization in the supply chain architecture environment. The SPaM team had to use a public-domain solver so that the supply chain team could independently check additional variations or permutations. Now that we have demonstrated the value of optimization and the solution the solver provided was proved to be robust, we bought a more sophisticated mixed-integer program (MIP) solver to look at existing areas of the supply chain architecture or to include finer data granularity, more variables, larger amounts of data, and nonlinearities. We were careful in our choice because we did not want to become dependent on a specific piece of software or machine.

We developed the second scenario model with a Microsoft Excel spreadsheet using the same input data and formulations as the optimizer. We used the spreadsheet

—To confirm the optimizer's results and debug it by looking for discrepancies,

—To validate the modeling approach against the business performance of the current supply chain,

—To challenge proposed solutions by exploring alternative scenarios,

—To run sensitivity analysis and offer zooming capabilities to provide insight, and

—To refine the cost analysis, for example, by adding inventory-driven costs in our example.

By iterating between the two models, we converged on the true optimal solution and controlled the size of the problem. The true optimal solution is one that the optimization model produces based on its formulation and data feeds, and also one enriched and validated by the experts and decision makers based on their experience and insight. The analyst uses the spreadsheet model to validate the results of the optimizer's solution, exploring much of the feasible space without using a resource-intensive combinatorial algorithm.

This combined modeling approach has the following advantages:

—Based on experience, practitioners can fine-tune the optimizer's solutions by adding qualitative knowledge not contained in the optimization formulation. This fine-tuning can significantly improve the convergence to the true optimal.

—A detailed understanding of the solution helps analysts to identify data sensitivities and robust solutions.

—By comparing the model's solution with experts' insights, analysts may find additional constraints to add to the initial formulation that reflect new assumptions worth considering before making a final recommendation. This approach increases our objectivity by forcing us to make the assumptions explicit.

—Expert insight helps us to incorporate scenarios that do not come from the optimizer to challenge the initial solutions, to think out of the box, or simply to compare new solutions to the current situation.

SPaM's approach incorporated both rigorous quantitative modeling and building acceptance of the recommendation. Our model helped us to decompose building the recommended network into individual steps. We needed to establish the implementation sequence and determine the extent of the ideal network. The business gave us information on the implementation (moving) costs for each transition path, and the scenario model calculated the related operational savings. We would then weight the savings against the corresponding costs of implementation to decide which transition path to choose and where to end it.

Our combined approach is particularly suited to green-field scenarios, in which we considered completely new sets of locations for a range of product groups, especially when time was short. We did not include implementation costs in the objective function of the optimization because of their complexity and dependence on the constraints of HP's service-provider strategy. We initially included inventory costs in the scenario model but not in the optimization model, but we now include them in both models. In the optimization model, we used the difference in the risk pooling between the current network design and an alternative network design as the main piece of the overall inventory-driven cost. This incremental piece of the overall inventory-driven cost made up two percent of the total costs.

HP managers saw the SPaM combined model's ability to break down implementation costs and make them visible for each step of implementation as an important strength. It smoothed the project's transition from analysis to implementation planning. Also, we found the scenario model useful for including or pinpointing unexpected geopolitical opportunities. Such opportunities can change the whole setup but are outside the scenarios initially considered. Taking late changes into account is vitally important in the business world.

The SPaM Combined Approach—IPG Results

Our project resulted in a new recommended network of postponement locations that was different from both the current network and from any scenario the project team and sponsors might reasonably have imagined.

We could model single-site scenarios easily, and the optimizer confirmed the best single-site solution. More than one location, however, made the problem more complicated, and the optimizer showed its real benefits. The optimizer looked at the global picture and considered possible synergies across some product categories by combining them together for the same site, hence reducing the fixed costs. This is not something one can weigh intuitively in seeking the right scenario. We ended up with a twofold recommendation for each of the specified transit times:

(1) We retained the current transit-time performance and recommended three postponement locations.

(2) We relaxed the specified transit-time specifications and again recommended three postponement locations, one of which differed from those in the previous recommendation.

With our approach, it was extremely important that we describe the model effectively and that managers understand and accept it because the model was central to the project. Managers wanted to understand the underlying assumptions and the rationale for its recommendations relative to the dimensions and the granularity of the data. IPG managers needed to be confident in implementing the model's recommendations, which promised to save over $10 million in supply chain costs and maintain the existing service levels. We believe that the traditional scenario-based approach would have yielded only one-quarter of this saving, based on scenarios we identified using our intuition and that of the project sponsors. Critically, we completed the project within the specified time line, which would not have been feasible with a full-scale optimization alone. Clearly, the combined modeling approach had value.

Appendix

MIP Problem Formulation

Given:

$j = 1, \ldots, 3N$ demand areas that combine N geographic areas and three business types ($1, \ldots, N$ for business A, $N+1, \ldots, 2N$ for business B, $2N+1, \ldots, 3N$ for business C).

$i = 1, \ldots, M$ possible nodes (PCC locations).

$p = 1, \ldots, P$ product groups.

$D_{jp} =$ demand in units of product p at demand area j.

$wp =$ weight in kilos per product p.

$f_{ji} =$ cost to ship one kilo from node i to node j (outbound freight).

$t_{ji} =$ transit time from node i to node j.

$\mathrm{TA} =$ transit time specification for business A (from PCC to demand area).

$\mathrm{TB} =$ transit time specification for business B (from PCC to demand area).

$\mathrm{TC} =$ transit time specification for business C (from PCC to demand area).

$\mathrm{TTreq}_j =$ transit time specification for demand area j (from PCC to demand area).

$L_{ip} =$ labor cost (variable manufacturing overhead (MOH) or distribution overhead (DOH)) \$ per unit at node i per product p.

$C_{ip} =$ fixed cost per time unit (capital equipment + overhead) at node i for product p.

$O_i =$ fixed cost per time unit (site + fixed overhead (MOH or DOH)) \$ for node i.

$R_p =$ regional manufacturing percent for product p.

$u_i =$ the cost to supply one kilo in node i from Rotterdam (inbound freight *).

$v_i =$ the cost to supply one kilo in node i from Hungary (inbound freight *).

$\mathrm{twa}_{jp} =$ time-weight for business A demand area j for product p, all PCCs.

$\mathrm{twb}_{jp} =$ time-weight for business B demand area j for product p, all PCCs.

$\mathrm{twc}_{jp} =$ time-weight for business C demand area j for product p, all PCCs, i.e., percentage of total demand going to demand area j.

We can create a formulation of a relaxed MIP:

Objective function:

$$\begin{aligned}\text{supply chain cost} &= \text{outbound freight} + \text{inbound freight}\\ &\quad + \text{variable MOH} + \text{fixed MOH}.\end{aligned}$$

One product family can be allocated to more than one PCC; if so, each PCC does a set of options implying it serves a region (a set of countries) at least for each type of businesses (indirect and direct).

The transit-time specification is independent of product family and demand area but specific to business type (indirect and direct).

$$\begin{aligned}\min\ & \sum_{i=1}^{M}\sum_{p=1}^{P}\sum_{j=1}^{3N}[(1-R_p)\cdot u_i + R_p v_i]\cdot d_{ijp}D_{jp}w_p\\ &+\sum_{i=1}^{M}\sum_{j=1}^{3N}\sum_{p=1}^{P} d_{ijp}f_{ij}D_{jp}w_p + \sum_{i=1}^{M}\sum_{j=1}^{3N}\sum_{p=1}^{P} d_{ijp}L_{ip}D_{jp}\\ &+\sum_{i=1}^{M}\sum_{p=1}^{P} b_{ip}C_p + \sum_{i=1}^{M} k_i O_i\end{aligned}$$

$$\begin{aligned}\text{s.t.}\ & \text{if } D_{jp}=0, \quad d_{ijp}=0,\\ & \text{otherwise } \sum_{i=1}^{M} d_{ijp} = 1 \quad \forall\{j,p\},\\ & d_{ijp} - b_{ip} \le 0 \quad \forall\{i,j,p\},\\ & b_{ip} - k_p \le 0 \quad \forall\{i,p\},\\ & \sum_{j=1}^{N}\sum_{p=1}^{P} d_{ijp}\mathrm{twa}_j t_{ij} \le \mathrm{TA}(t1),\\ & \sum_{j=N+1}^{2N}\sum_{p=1}^{P} d_{ijp}\mathrm{twb}_j t_{ij} \le \mathrm{TB}(t2),\\ & \sum_{j=2N+1}^{3N}\sum_{p=1}^{P} d_{ijp}\mathrm{twc}_j t_{ij} \le \mathrm{TCB}(t3),\\ & d_{ijp} \in [0,1],\\ & b_{ip} \in \{0,1\},\\ & k_i \in \{0,1\}.\end{aligned}$$

Constraints ($t1$), ($t2$), and ($t3$) ensure that the weighted average of the transit time meets the specifications. Constraints ($t1$), ($t2$), and ($t3$) can be replaced by

$$\sum_{i=1}^{M} d_{ijp} t_{jp} \le \mathrm{TTreq}_j \quad \forall j, p$$

if we want every single transit-time performance to meet the country-specific target for every product group.

Optional constraint to drive the number of PCCs either per product line (PL) or total:

$$\sum_{i=1}^{M} b_{ip} \leq N b_p \quad \forall p \quad \text{or} \quad \sum_{i=1}^{M} k_i = N \quad \text{total.}$$

$d_{ijp} = 1$ if demand from node i to demand area j for product p (default 0) in theory can be a percentage, but there is no reason to allocate this demand to more than one PCC (the optimizer will select the one that provides the lowest cost).

$b_{ip} = 1$ if node i is open for product p (default 0).

$k_i = 1$ if node i is open for "something" (default 0).

(1) In the final solved problem, $M = 10$, $P = 20$, and $N = 20$, resulting in 12,200 variables and 14,600 constraints. With shareware solvers customized for this problem, we could solve it in seconds.

(2) The time unit is a month for demand and fixed costs. Transit is time in working days.

(3) Demand is defined per time unit (e.g., monthly), reflecting the countries and subcountries (e.g., south Germany) for three types of business (different delivery frequency, different freight costs, and transit-time requirements):

(a) Business A, i.e., indirect (from PCCs to channel partners),

(b) Business B, i.e., replenishment (through Regional Warehouses),

(c) Business C, e.g., direct.

(4) Outbound freight in \$/kg (average bin range). We ignored the expediting mode.

(5) Inbound: modeled as a unique flow (corresponding to the main units) with the total weight of the finished products:

(a) Either from Rotterdam (port of entry for sea shipments from Asia),

(b) or from one (only) factory in Eastern Europe (Hungary?).

(6) Inbound freight in \$/kg corresponding to full-truckload costs.

(7) Fixed percentage of products produced in the region versus those coming from Asia.

(8) MOH is defined with three elements:

(a) Variable labor (per unit) = f(PCC location, product family),

(b) Fixed cost coming from having a product family in a PCC (capital equipment + product-specific overhead) = f(product family),

(c) Fixed cost coming from running a PCC (building lease, overhead) = f(PCC location).

(9) Remarks on TAT:

(a) With no TAT constraints, there is no chance that more than one PCC will be selected for any single product family. Different TAT requirements may lead to the selection of the center of gravity for quick deliveries and a more remote location for best costs.

The TAT constraint is treated either as a weighted average performance across countries (weight being demand) or more simply as a specification for each and every demand area. In both cases, the constraint is specific to every business type.

Acknowledgments

We are grateful for the helpful comments of two anonymous referees that helped shape this paper. We also acknowledge Astrid van der Staak, Gabriele Woerner, and Reinhard Winkler for their support during this project.

References

Bowerman, Robert L., Paul H. Calamai, G. Brent Hall. 1997. The demand partitioning method for reducing aggregation errors in p-median problems. *Comput. Oper. Res.* **26** 1097–1111.

Camm, Jeffrey D., Thomas E. Chorman, Franz A. Dill, James R. Evans, Dennis J. Sweeney, Glenn W. Wegryn. 1997. Blending OR/MS, judgment, and GIS: Restructuring P&G's supply chain. *Interfaces* **27**(1) 128–142.

Cargille, Brian, Robert Hall, Steve Kakouros. 1999. Part tool, part process—Inventory optimization at Hewlett-Packard Co. *OR/MS Today* **26**(5) 24.

Cohen, Morris A., Hau L. Lee. 1989. Resource deployment analysis of global manufacturing and distribution networks. *J. Manufacturing Oper. Management* **2** 81–104.

International Institute for Management Development (IMD). 2003. *IMD World Competitiveness Yearbook 2003*. IMD Lausanne, Switzerland. http://www02.imd.ch/wcy/.

Lee, Hau L., Corey Billington. 1995. The evolution of supply-chain management models and practice at Hewlett-Packard Company. *Interfaces* **25**(5) 42–46.

McBurney, Peter, Simon Parsons. 2002. Formalizing scenario analysis. A. Darwiche, N. Freeman, eds. *Proc. 18th Conf. on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, San Francisco, CA, 327–334.

McBurney, Peter, Simon Parsons. 2003. Chance discovery and scenario analysis. *New Generation Comput.* **21**(1) 13–22.

Murphy, Frederic H. 1998. The occasional observer: Some simple precepts for project success. *Interfaces* **28**(5) 25–28.

Gabriele Woerner, Director Supply Chain Architecture and Design, IPG Supply Operations EMEA, Hewlett Packard, GmbH, Schickardstrasse 32, Geb. Businesspark, Boeblingen, Germany, writes: "SPaM did an exceptional job in establishing a state-of-the-art optimization model, taking into account all key parameters that drive the location decision. The results identified supply chain cost savings in excess of \$10 million per year. We had an excellent basis to make the right decisions and shape the future IPG supply chain physical infrastructure in EMEA. We implemented the key elements of the recommendation and realized the majority of the potential savings. Thanks a lot for your initiative, commitment and the high quality work you provided in this analysis!"