Implementing Supply Routing Optimization in a Make-To-Order Manufacturing Network

John Foreman[1], Jérémie Gallien[2], Julie Alspaugh[3], Fernando Lopez[4], Rohit Bhatnagar[5], Chee Chong Teo[6], Charles Dubois[7]

October 7, 2008

Abstract

Dell's supply chain for desktops involves Asian vendors shipping components by sea to several U.S. plants. While suppliers are responsible for shipping enough inventory, Dell can re-route and expedite their shipments while in transit and also transfer on-hand inventory in order to balance supply across sites. This paper describes the development, implementation and impact of the process and optimization-based control system now used by Dell to address this supply routing challenge for its US-bound monitors. This new methodology is estimated to have reduced Dell's inventory re-positioning costs for monitors by about 60%.

1 Introduction

Like many other manufacturers, Dell Computers has seen growth come with a significant increase in the complexity of its operations. Over the last ten years, this evolution has specifically taken the following forms for Dell's North American desktop division: (i) increase of the number of assembly plants and warehouse facilities; (ii) replacement of most US-based suppliers with Asian suppliers; and (iii) extension of the product line offered to customers. Although these changes have directly impacted most of Dell's operational functions, they have in particular drastically complicated the mission of its procurement group. Indeed, that group has thus become responsible for maintaining the availability of more components in more locations, working with suppliers having longer transportation lead-times.

In order to address this supply availability challenge, Dell has long relied on Vendor-Managed Inventory (VMI) relationships, whereby its suppliers are responsible for maintaining a sufficient inventory of components in each one of Dell's relevant locations, relative to a demand forecast periodically communicated by Dell, e.g. 14 Days of Supply in Inventory (DSI). As part of that relationship, component inventory continues to be owned by suppliers

---

[1]   MIT Operations Research Center, Cambridge, MA 02142
[2]   MIT Sloan School of Management, Cambridge, MA 02142 (corresponding author, e-mail: jgallien@mit.edu)
[3]   Dell Computers, Inc., Round Rock, TX 78682
[4]   Dell Computers, Inc., Round Rock, TX 78682
[5]   Nanyang Technological University, Singapore
[6]   Nanyang Technological University, Singapore
[7]   Ecole des Mines de Paris, France

until only a couple of hours before that inventory is pulled onto Dell's production line (or warehouse pick process) for assembly, and suppliers are mostly free to follow any schedule of shipments as long as it meets the service level requirements just mentioned. In order to benefit from transportation volume discounts however, these shipments are typically done through ocean and air carriers directly contracted by Dell. Also, Dell centralizes inventory and shipment information, in part because it often uses several suppliers for the same component. As a result, Dell has retained the function of managing both the routing of its pipeline inventory and the transhipments of its on-hand inventory between various facilities (*supply routing*), even though that inventory may still be on suppliers' books at the time when it is moved or re-routed (see Reyner 2006 and Kapuscinski et al. 2004 for more background and references on Dell's business model, supply chain and history).

The supply routing function just defined is particularly important for components such as desktop chassis and monitors, which account for a substantial proportion of total supply transportation costs. These components are typically shipped by ocean from Asia to the US in full containers of a single part type because of their large volume and weight, and therefore also have long delivery lead-times. As a result, gaps between actual realized demand in each assembly or warehouse facility and the corresponding forecasts driving these shipments can become quite large over this transportation delay, causing potentially large imbalances in the inventory positions of Dell's various sites. In turn, these imbalances may cause customer delivery delays due to component shortages as well as additional inventory holding costs, which Dell is keen to reduce through various means. Specifically, in North America Dell can change at some cost the final destination of any container still on the Pacific ocean (*diversion*) as well as its planned ground transportation mode (*expedition*) up until a couple of days before it is disembarked in Long Beach, CA; available ground transportation modes include, with increasing cost and decreasing lead-time, the default rail and truck mode; a single driver truck-only mode; and a two driver (team) truck-only mode. In addition, Dell can also perform transhipments (*transfers*) of on-hand inventory between its facilities. The available transportation modes for these include a set schedule of pre-contracted truck "milk runs" between Dell US facilities, which have low relative cost but limited capacity, as well as specially affreighted single or team trucks. Figure 1 illustrates the supply-chain structure and the associated supply routing decisions just defined, and also shows the four main locations in the US where chassis and monitors are shipped for assembly and/or inclusion into customer

2

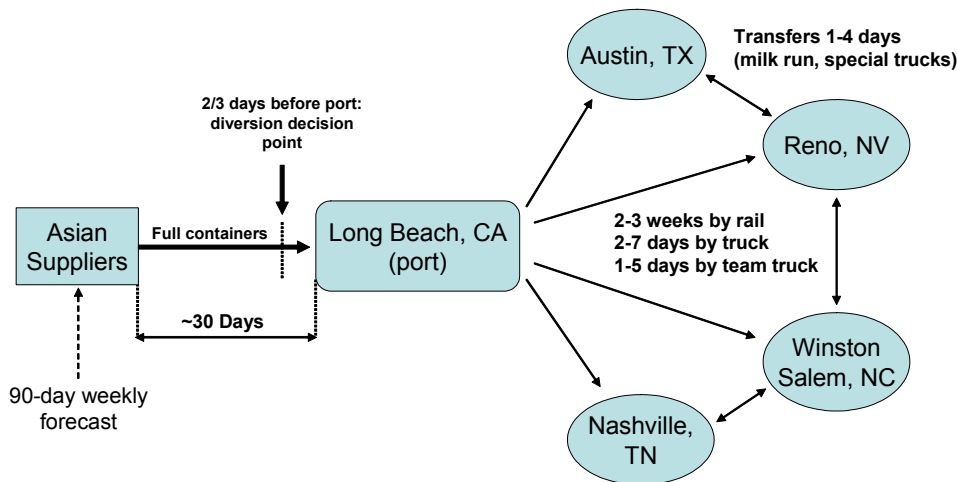orders as well as typical transportation lead-times.

Figure 1: Dell's Supply-Chain Structure and Supply Routing Decisions for US-Bound Desktop Chassis and Monitors.

The volume of material continuously going through the supply chain just described is very significant: a rough estimate from Dell's public 10-K filing for fiscal year 2008 reveals that tens of thousands of units of each component type must have been shipped every week on average to Dell US facilities over that period. The challenge of making all the associated diversion, expedition and transfer decisions in a timely and cost-effective manner thus constitutes a supply chain control problem that is both important and difficult. The present paper summarizes the collaboration conducted by Dell Computers and university researchers over several years to develop the optimization-based control system now used by Dell to address this supply routing challenge for its US-bound monitors. It contributes to the operations management literature by providing a detailed description of a real-world control challenge which has not been discussed extensively so far even though it is critical to the operation of an important supply chain. It also describes a model for addressing this challenge along with a process for implementing it which have been tested and validated by practice.

The remainder is organized as follows. After a discussion of the related literature in §2, we present the two main successive phases of that collaboration in §3 and §4. When describing each phase, we first define the processes and tools developed (§3.1, §4.1), then discuss implementation issues (§3.2, §4.2) as well as observed impact, both quantitative and

3

qualitative (§3.3, §4.3). Finally, §5 contains concluding remarks pertaining to the limitations of our work, ongoing related developments, possible future research and key learnings from this collaboration. An important notational convention used throughout this paper consists of using symbols in bold for random variables, and the same symbols with no highlight for their mathematical expectations, e.g. $d \triangleq \mathbb{E}[\mathbf{d}]$. Also, notations with an upper bar denote cumulative quantities, e.g. $\bar{d}_t = \sum_{k=1}^{t} d_k$. In order to protect the confidentiality of Dell's sensitive information, some of the numerical data included in this paper is disguised.

2 Literature Review

The reader may have noted from §1 that the high-level structure of Dell's supply chain for large desktop components in North America closely resembles that considered in the inventory distribution model of Eppen and Schrage (1981) except that, using the terminology introduced in that paper, Dell only performs the *allocation* function (splitting incoming quantities among final destinations) and has delegated the *ordering* function (determining incoming quantities) to its suppliers. Among common features, Dell's supply routing problem also involves the centralized allocation of incoming inventory among several facilities where it is stored and consumed, and its cross-docking disembarkment operation in Long Beach, CA exactly matches the definition of a "stockless depot" considered in Eppen and Schrage (1981). Consequently, many of the results and insights described in the body of literature on multi-echelon inventory distribution theory that started with that seminal paper (see Axsäter, Marklund and Silver 2002 for a recent survey) are conceptually relevant to the problem considered. In addition, several of these papers analyze extensions of Eppen and Schrage's assumptions which capture key operational features of Dell's supply chain: Among others, Federgruen and Zipkin (1984) consider facilities with non-identical cost structure facing non-stationary demands, Jönsson and Silver (1987) consider transhipments between facilities and Jackson (1988) considers non-identical transportation lead-times to the facilities. Another relevant extension to the theory of multi-echelon inventory control, albeit one only carried out so far for serial systems, is the inclusion of expedited transportation modes, as described in Lawson and Porteus (2000).

In spite of all these papers' contextual relevance, both our goal and methodology differ substantially from theirs. Specifically, our objective is to develop and implement an operational system for a large real life supply chain, as opposed to deriving theoretical insights

from a stylized model. Consequently, our approach sacrifices tractability for realism and operational applicability, and the model we formulate is a mixed-integer program solved over a rolling horizon using numerical (branch and bound) algorithms, as opposed to say a dynamic program – see Chand, Hsu and Sethi (2002) for a more general review of rolling horizon methods.

Some insights on the very supply chain motivating our work may be gained from Kapuscinski et al. (2004), which describes the development and implementation by Dell of replenishment models for its component inventory. That paper is thus an important complement to ours, in that it focuses on how the inventory ordering decisions, which we assume exogenous in the problem we consider, should be generated by Dell's suppliers as part of their VMI relationship (see §1).

The paper most related to ours however is Caggiano, Muckstadt and Rappold (2006), which considers operational models for inventory and capacity allocation decisions in a multi-item reparable service part system with a central warehouse and field stocking locations. In particular, their Extended Stock Allocation Model (ESAM), which leaves the repair decisions aside, is similar in many respects to the one we describe in §4: it is a mathematical program meant to be solved on a rolling horizon basis; its decision variables comprise inventory allocation and expedited shipment decisions; its objective function includes a transportation component and a newsboy-like backorder component; it assumes deterministic lead-times and an exogenous supply pipeline. However, the ESAM is still simpler than our model, in that it does not capture transhipments, considers a single expedited transportation mode, assumes a linear transportation cost structure and ignores transportation scheduling and capacity constraints. These differences are material, as the solution approach ultimately followed in Caggiano, Muckstadt and Rappold (2006) consists of developing heuristic solutions exploiting the structural properties which can be established in their setting, while we compute instead solutions to an approximate (linearized) version of our model (see §4.1). Most importantly however, the present paper describes an actual implementation by an industrial firm of the optimization model presented along with an assessment of its impact, and thus offers a grounded perspective on the many important practical issues involved. This practical focus is reflected by the structure of this paper, whereby we now describe in turn the two successive phases followed by Dell as part of that implementation.

3 First Phase: Process Design

3.1 Development   The first phase of this project, which is described more extensively in Reyner (2006), started in the Spring of 2005 with the determination by Dell executives that the increasing trend in expediting and transfer costs which had been observed over the previous months should be corrected. A team that was tasked then to investigate this problem and issue recommendations identified two root causes. The first was that supply routing decisions (i.e. diversions, transfers and expeditions) were being made by members of several internal groups within Dell, and that the coordination between these groups only occurred on an ad-hoc basis. The second root cause identified was that the information relevant to supply routing decisions (at a minimum, the demand forecast, current inventory and supply pipeline) was scattered, difficult to obtain in a timely fashion, and occasionally unreliable.

The associated solution designed and implemented later that year comprised two key components. The first was organizational, and involved the creation (and staffing) of a specific job definition named *supply chain analyst*, with the explicit responsibility of gathering and analyzing all relevant information in order to make and implement all inventory routing decisions for a specified set of components. The second component was the development of a spreadsheet-based information acquisition and visualization tool in support of that role, which became known as the *Balance Tool*. As Figure 2 illustrates, this tool simultaneously displays all available demand forecasts and scheduled supply deliveries for each component in each of the relevant factories and warehouses over a rolling horizon of several weeks, with a planning period of one day. Combining that information with the current inventory on hand and backlog in the various sites allows Dell to compute a projected net inventory and equivalent DSI (*days of supply in inventory*) levels in all the relevant locations over this horizon, and highlight any anticipated shortages. Specifically, the Balance Tool uses a color code and associated categories of DSI levels to show each day of the horizon in each facility as red (critical situation), yellow (should be monitored) or green (sufficient inventory). Based on daily automatic updates of the information displayed for each component and using special entry cells, the supply chain analysts can then manually explore the implications of all possible routing decisions. For example, a container scheduled to arrive in Austin in the later part of the horizon could be diverted to Nashville and expedited by team truck, which

the Balance Tool would reflect by removing that container from Austin's supply line on its original scheduled arrival date and adding it to that of Nashville on a closer date (determined by the difference between the transportation times from Long Beach to Austin by rail and to Nashville by team truck, respectively). The resulting new inventory and DSI levels in both sites resulting from such a move would then be instantly displayed, showing for example the extent to which this action would help correct a projected shortage situation in Nashville in the short term when Austin is projected to have excess inventory later in the horizon. Finally, the length of the planning horizon was chosen so that it would always include any containers located before the diversion cut-off point of 2-3 days before port, assuming the longest possible ground transportation lead-time and then adding an additional time buffer.
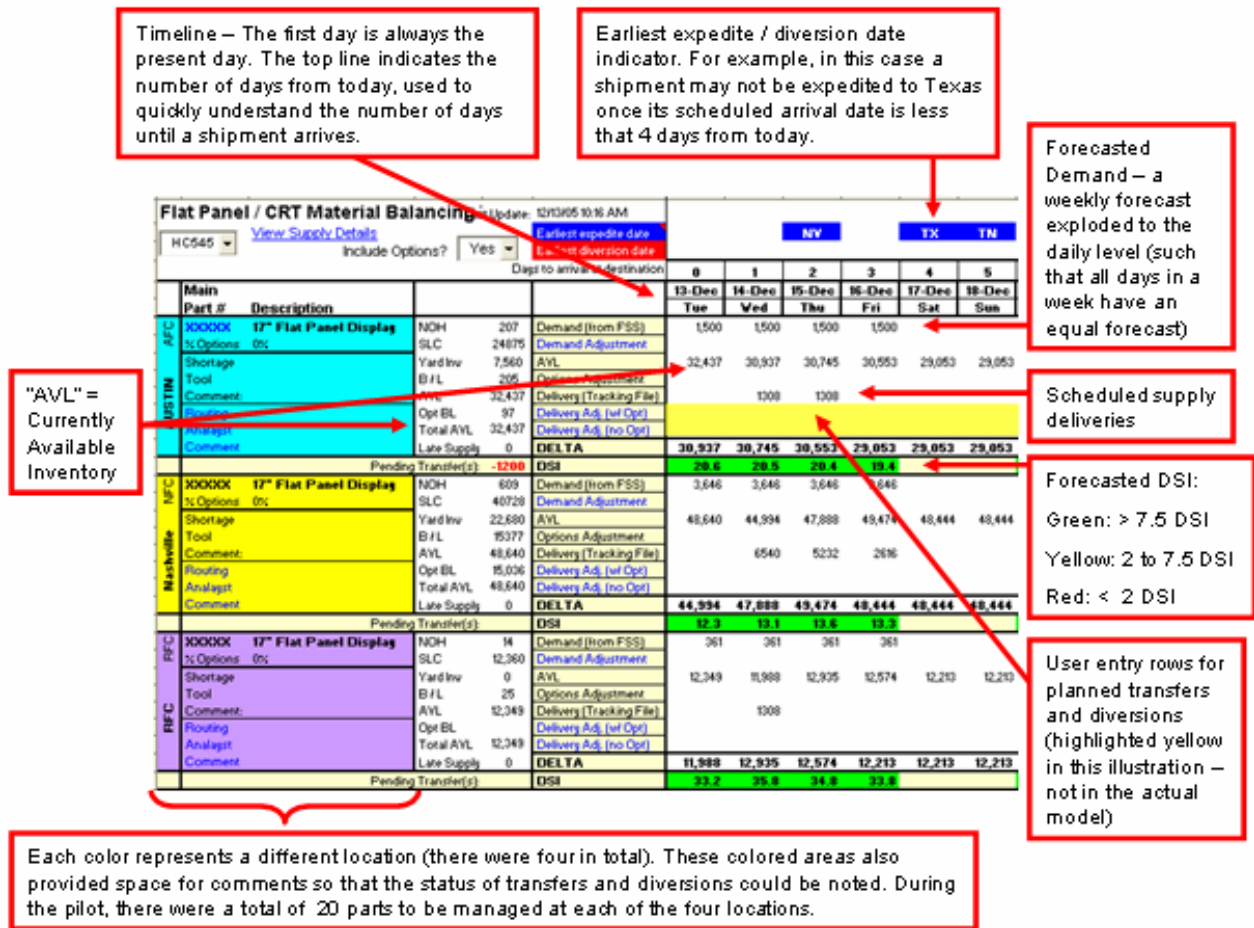

Figure 2: Visualization of Dynamic Routing Decisions with the Balance Tool (from Reyner 2006)

3.2 Implementation   The creation of the supply routing analyst position was welcomed

7

as a relief by the various groups previously involved in making these decisions, in part because many involved regarded supply routing as a non-explicit yet time-consuming part of their mission. Also, there was a fairly large consensus that the previous ad-hoc process could be improved upon, and a sense of urgency created by the executive directions mentioned in §3.1. The Balance Tool was implemented with the spreadsheet program Microsoft Excel, augmented with Visual Basic macros which automated certain functions such as retrieval of external data as well as creation and deletion of parts. A first technical challenge consisted of identifying the diverse heterogeneous databases containing the relevant data (demand forecasts, available inventory and projected deliveries) and developing ways to query them. Another challenge consisted of working with the groups primarily responsible for entering any new information to these databases in order to ensure that these entries would be systematically submitted in a timely manner and consistent format. This proved particularly challenging with the external suppliers and carriers, who were responsible for submitting expected arrival dates of all incoming shipments and updating them whenever appropriate. Difficulties stemmed in part from incentive issues as well as the variety of information systems used at the time by Dell to communicate with suppliers.

A key implementation and development strategy in this phase was to organize a live pilot of the newly designed decision process for a selected subset of components (monitors) relatively early on (September 2005). The associated pressure to generate actual decisions allowed the analyst to quickly identify and address many improvement opportunities for the Balance Tool and the communication formats used with suppliers and carriers. It also allowed the analyst to quickly identify the information sources which were most problematic (see above), and thus dedicate time to checking input information only selectively. Finally, the pilot forced a grounded reflection on how supply routing decisions should be made in specific situations and helped quantify the impact of the new process, as we discuss next.

3.3 Impact   The financial impact of this first phase was estimated using a fairly coarse methodology. Specifically, managers reviewed all the decisions made over a limited period of time during the live pilot described above, along with the associated input data. In each case, they determined which alternative decisions would likely have been made under the previous process, along with their associated transportation costs. Because the part shortages were generally perceived to have improved during the pilot, the transportation cost savings

associated with the new process that were calculated in this way were considered meaningful. Although this methodology involves many subjective and arguably biased inputs, its results were still deemed valid by Dell's managers and led to their conclusion that the new process reduced the transportation costs associated with supply routing by about 40% (Reyner 2006).

This quantitative impact estimation seems to have been easily accepted because it had a clear qualitative explanation. Specifically, the new process generated comparatively more rail diversions, which only involve a small bill of lading splitting fee, and fewer transfers and expeditions, which are considerably costlier decisions. This can be traced back to the organizational structure previously underlying supply routing decisions, whereby responsibilities for continuity of supply were allocated based in part on when in the future that supply would be needed for production. For example, the Production Control group is concerned with on hand inventory, the Purchasing group monitors any supplier capacity problems potentially impacting production one or two months into the future, etc. In this setting, rail diversions had been a neglected lever because they required more information than transfers and needed to be executed about a couple of weeks before delivery, a domain whose organizational ownership was unclear. In contrast, the new process allowed for all supply routing decisions to be considered in a coordinated fashion.

Dell started using the new process described above continuously for all its monitors (about two dozen different part types) in January 2006, shortly after the live pilot and its assessment were completed. It was also deployed later that year for chassis. In spite of their many benefits, these deployments also revealed additional improvement opportunities, linked to the absence of formal decision rules for how routing decisions should be made as a function of the input data available. Relying exclusively on the analysts' judgement proved problematic from a time efficiency standpoint, because of the high number of parts to manage, the very high number of potential decisions involved for each part, the large amount of relevant information and the high decision-making frequency: while forecasts can change daily, for example whenever a large customer order "drops", the analysts were thus only able to review and affect the status of any particular part once a week on average. From a resiliency standpoint, it also seemed problematic for Dell to depend entirely on a handful of individuals for such frequent and critical control decisions. Finally, the Balance Tool only characterized expected shortages very coarsely through the net inventory levels displayed and the color

code described above, and did not provide an estimate of the transportation costs associated with the decisions considered. It was thus suspected that even an expert analyst could easily make sub-optimal decisions in this setting. These considerations all motivated the second phase of this collaboration, which is described next.

4 Second Phase: Optimization

The second phase of this project started in September 2006. Its objective was to develop and implement a formal optimization model that could automatically generate Dell's supply routing decisions as part of the process described in §3, or at least serve as a decision support system to assist the supply routing analysts.

4.1 Development It was quickly realized that any optimization model serving the purpose just stated would need to quantify the main trade-off involved in Dell's supply routing decisions, namely the tension between transportation costs on one hand and shortage costs on the other hand. While expressing the transportation costs incurred as a function of the routing decisions considered is relatively straightforward as will be seen shortly, the critical modeling challenge was to quantify the benefits associated with these decisions, that is the overall change of expected shortage costs in all of the sites where the projected inventory levels were affected. We describe our work on this problem in §4.1.1, then present the resulting complete optimization model in §4.1.2.

4.1.1 Expected Shortage Costs We adopted a standard linear structure $B \sum_{t \in \mathcal{T}, \ell \in \mathcal{L}} v_{t\ell}$ for the total expected shortage costs predicted in all facilities $\ell \in \mathcal{L} \triangleq \{$Austin, Nashville, Reno, Winston-Salem$\}$ for a specific part over the rolling horizon $t \in \mathcal{T} \triangleq \{1, .., T\}$ considered, where B is a unit daily shortage cost rate and $v_{t\ell}$ is the expected average shortage level for future day t in location ℓ. In practice, shortage costs stem from a variety of factors including primarily: order cancellations by impatient customers; expedited shipping to customers with late orders; substitutions of more expensive components for the same price; lost profit from customers turned away by long posted lead-times; price concessions on future orders... We refer the reader to Kapuscinski et al. (2004) and Dhalla (2008) for more exhaustive and detailed descriptions. While B is assumed available input data for now, we return to this issue in §4.2, and develop next an expression for $v_{t\ell}$ as a function of the supply routing decisions considered and the inventory and demand data available.

In the absence of forecasting errors, maintaining barely positive net inventory levels would suffice to ensure supply continuity; the quantity of component inventory that should be maintained at or shipped to any location must therefore depend directly on these errors. Our first step thus consisted of characterizing the distribution of actual demand relative to the forecast available for that quantity at the time when routing decisions need to be made, as this information was not available to us at the outset. This empirical study of the cumulative forecast error (see §A.1 in the Online Appendix) both suggested the structure and provided the standard deviation input data $\sigma_{t\ell}$ for the stochastic model

$$\sum_{k=1}^{t} \mathbf{d}_{k\ell} \sim N(\sum_{k=1}^{t} f_{k\ell}, \sigma_{t\ell}), \tag{1}$$

where: $\mathbf{d}_{t\ell}$ is the random variable representing demand on day t for a given part in a given location $\ell \in \{\text{Austin, Nashville, Reno, Winston-Salem}\}$, as estimated at the beginning of the current day (always indexed by 1 in our rolling horizon model), so that $\sum_{k=1}^{t} \mathbf{d}_{k\ell} \triangleq \bar{\mathbf{d}}_{t\ell}$ is the cumulative demand for the next t days; $N(f, \sigma)$ refers to a normal distribution with mean f and standard deviation σ; $f_{t\ell}$ is the (deterministic) forecast of the same quantity generated by Dell and provided to the supply-chain analyst on day 1, so that $\sum_{k=1}^{t} f_{k\ell} \triangleq \bar{f}_{t\ell}$ is the corresponding cumulative forecast of demand up to day t; and finally $\sigma_{t\ell}$ is the standard deviation of the forecasting error $\bar{\mathbf{d}}_{t\ell} - \bar{f}_{t\ell}$. Note that the forecasting error study mentioned above did identify some systematic biases. However, because they were relatively small and convincingly explained by the forecasting team, we decided to ignore them as part of our model. The notations $f_{t\ell}$ and $d_{t\ell} \triangleq \mathbb{E}[\mathbf{d}_{t\ell}]$ will thus be used interchangeably from now on.

The inventory dynamics over the rolling horizon considered are described in our model by the following balance equation, which assumes that any unmet demand is backlogged:

$$\mathbf{I}_{(t+1)\ell} = I_{1\ell} + \sum_{k=1}^{t} s_{k\ell} - \bar{\mathbf{d}}_{t\ell} \text{ for } t \geq 1. \tag{2}$$

In (2), $\mathbf{I}_{t\ell}$ is the (random) net inventory level available at the beginning of day t in location ℓ, as predicted at the beginning of day 1 (so that $I_{1\ell} = \mathbf{I}_{1\ell}$ is deterministic input data), and $s_{t\ell}$ is the net result of deliveries into and transfers out of location ℓ on day t (which is directly affected by the supply routing decisions we seek to determine). Note that $s_{t\ell}$ is assumed to be deterministic in our model, which ignores supply uncertainty. This is justified by the fact that in Dell's setting supply uncertainty is small relative to demand uncertainty given the

11

(daily) time granularity considered. As a result, the ranges of lead-times appearing in Figure 1 are essentially driven by the differences across destinations, as opposed to any potential unpredictable variability affecting the lead-time associated with a given transportation mode on a specific leg. Also, that assumption does not affect the operational applicability of the model output, as will be seen further.

Next, we characterize the average shortage level $\mathbf{v}_{t\ell}$ for day t in location ℓ predicted at the beginning of day 1 as

$$\mathbf{v}_{t\ell} \triangleq (\mathbf{I}_{t\ell} - \mathbf{d}_{t\ell})^- . \tag{3}$$

The expression $\mathbf{I}_{t\ell} - \mathbf{d}_{t\ell}$ for the average net inventory level during day t which appears in (3) corresponds to the most pessimistic assumption for the daily schedule of supply and demand. That is, demand is assumed to occur entirely at the very beginning, and supply deliveries at the very end, of day t. This approach was followed because the expressions derived from other assumptions (say continuous supply and demand processes) are less tractable, detailed hourly demand and supply data was not easily accessible, and because of Dell's expressed desire to err on the conservative side when predicting shortages[8]. Substituting (2) and (1) in (3) yields

$$\mathbf{v}_{t\ell} \sim [N(I_{t\ell} - f_{t\ell}, \sigma_{t\ell})]^- \quad \text{for } t \geq 1, \tag{4}$$

which characterizes the distribution of average shortages in terms of decision variables and input data[9]. The expectation of the random variable in (4) is thus given by the standard expression

$$v_{t\ell} = \sigma_{t\ell}\phi\left(\frac{f_{t\ell} - I_{t\ell}}{\sigma_{t\ell}}\right) + (f_{t\ell} - I_{t\ell})\Phi\left(\frac{f_{t\ell} - I_{t\ell}}{\sigma_{t\ell}}\right), \tag{5}$$

where ϕ and Φ are the standard normal p.d.f. and c.d.f., respectively.

Because the r.h.s. of (5) as a function of $f_{t\ell} - I_{t\ell}$ is convex, the upper envelope of a finite number of its tangents constitutes an approximation that can be made arbitrarily accurate and is particularly efficient from a computational standpoint. To construct this approximation, we first determine upper and lower bounds $I_{t\ell}^{LB}$ and $I_{t\ell}^{UB}$ for $I_{t\ell}$ that are independent of decision variables and only depend on input data. A tight upper bound $I_{t\ell}^{UB}$ is easily computed by observing that in any given situation, the maximum expected net inventory

---

[8] The alternative definition $\mathbf{v}_{t\ell} \triangleq (\mathbf{I}_{t\ell} + s_{t\ell})^-$ is equally tractable and is the most optimistic in the sense just defined. It thus allowed us to verify that the modeling assumption reflected by (3) was relatively immaterial.

[9] Random variables and their distributions are used interchangeably in (4), as no related ambiguity arises here.

level in location ℓ at the beginning of time t would be obtained by instantly transferring to ℓ all inventory from other facilities, and re-routing towards ℓ with the fastest ground transportation mode (team truck) all divertable containers which can arrive to ℓ by time t. Likewise, a tight lower bound $I_{t\ell}^{LB}$ corresponds to the situation where all inventory available in ℓ is transferred immediately to other facilities, and all containers initially bound to ℓ are diverted away while demand continues to deplete this facility[10]. Secondly, we calculate iteratively a discrete set of sampling points $\mathcal{P}_{t\ell} \subset [I_{t\ell}^{LB}, I_{t\ell}^{UB}]$ indexed by p, and the slopes $a_{t\ell p}$ and intercepts $b_{t\ell p}$ of the corresponding tangents to the r.h.s. of (5), using numerical implementations of ϕ and Φ along with the maximum error rule method described in Rote (1992)[11]. This approximation is amenable to implementation as part of a linear optimization model, which we now proceed to describe.

4.1.2 Optimization Model Formulation

The basic optimization model we developed to generate supply routing recommendations for each individual part at a specific point in time is the following mixed integer program:

Input Data: Besides the rolling horizon $\mathcal{T} \triangleq \{1, .., T\}$ and the set of relevant locations $\mathcal{L} \triangleq \{\text{Austin, Nashville, Reno, Winston-Salem}\}$ previously introduced, the **part** considered is characterized by a maximum number of parts per truck Q and a number of parts per pallet J; **Incoming supply** consists of a set $\mathcal{C}$ of containers indexed by i, each containing a quantity of parts q_i with a current destination $\ell_i \in \mathcal{L}$ and an expected arrival date $A_i \in \mathcal{T}$. Containers that are still divertable (typically all containers still on the ocean and at least 2 or 3 days away from port) form a subset $\mathcal{C}^{RT} \subset \mathcal{C}$, while the containers in the complement set $\mathcal{C}^{NRT} \triangleq \mathcal{C} \backslash \mathcal{C}^{RT}$ may no longer be re-routed before they arrive to destination. The expected arrival date at the port (Long Beach, CA) of containers i in $\mathcal{C}^{RT}$ is denoted $A_i^{LB} \in \mathcal{T}$. Often containers travel as a group of multiple containers all sharing the same bill of lading, and therefore the same destination, transportation mode and expected arrival time. Containers with the same bill of lading may be split however, provided they belong to $\mathcal{C}^{RT}$. In this case, the carrier creates as many new bills of lading as the new resulting number of container groups traveling together, incurring an administrative fee of c^{BL} times the number of new

---

[10]   The equations for $I_{t\ell}^{LB}$ and $I_{t\ell}^{UB}$ in terms of input data are straightforward and omitted here.
[11]   This algorithm initiates with $\mathcal{P}_{t\ell} = \{I_{t\ell}^{LB}, I_{t\ell}^{UB}\}$. In each iteration, tangents are constructed for each new point in $\mathcal{P}_{t\ell}$, and the x-axis values of the intersection of tangents corresponding to adjacent points in $\mathcal{P}_{t\ell}$ are added as new points. The algorithm terminates when the maximum difference between the y-axis values of these intersections and the corresponding function values reaches a specified upper bound.

bills of lading created. Bills of lading are indexed by $j \in \mathcal{J}$, and the subset of containers sharing each bill of lading j is denoted $\mathcal{C}_j$, so that $\mathcal{C} = \cup_{j \in \mathcal{J}} \mathcal{C}_j$; **Current net inventory** (on-hand inventory minus backorders) currently available in each location $\ell \in \mathcal{L}$ (that is at the beginning of day 1 of the rolling horizon) is denoted $I_{1\ell}$; **Forecast of demand** in location $\ell \in \mathcal{L}$ during day t is denoted $f_{t\ell}$, while the cumulative forecast of demand from day 1 to day t (included) is denoted $\bar{f}_{t\ell}$; **Container ground transportation modes** between the port and Dell's facilities are indexed by $m \in \mathcal{M}^{RT} \triangleq \{$rail,single truck,team truck$\}$ and characterized for each destination $\ell \in \mathcal{L}$ by a cost per container $c_{\ell m}^{RT}$ and an average lead-time $L_{\ell m}^{RT}$ (expressed in days). The potential expected delivery date in location ℓ of any divertable container $i \in \mathcal{C}^{RT}$ is thus $A_i^{LB} + L_{\ell m}^{RT}$; **Specially affreighted transfers** of inventory between two facilities ℓ and ℓ' in $\mathcal{L}$ are characterized by their expedition mode $m \in \mathcal{M}^{SP} \triangleq \{$single truck,team truck$\}$, their cost per truck $c_{\ell\ell' m}^{SP}$ and their lead-time $L_{\ell\ell' m}^{SP}$; **Milk run transfers** of inventory from facility ℓ to facility ℓ' are characterized by their schedule of departures $S_{t\ell\ell'}^{MR}$ (equal to 1 if a run from ℓ to ℓ' is scheduled on day t and 0 otherwise), their cost per pallet $c_{\ell\ell'}^{MR}$, the maximum number of pallets of a given part allowed in each run R, and their lead-time $L_{\ell\ell'}^{MR}$; **Shortage cost parameters** include the unit daily shortage cost rate B as well as each slope $a_{t\ell p}$ and intercept $b_{t\ell p}$ of the approximating tangents to the expected shortage cost function indexed by $p \in \mathcal{P}_{t\ell}$ for each day t and location ℓ. An associated upper bound for the absolute value of the expected net inventory level at the beginning of day t in location ℓ is $M_{t\ell} = \max(|I_{t\ell}^{UB}|, |I_{t\ell}^{LB}|)$ (see §4.1.1).

Decision Variables: **Container routing** decisions are captured by binary variables $y_{i\ell m}$ equal to 1 if container $i \in \mathcal{C}^{RT}$ is routed from the port to facility ℓ using transportation mode $m \in \mathcal{M}^{RT}$, and 0 otherwise. In addition, binary variables $z_{j\ell m}$ take the value 1 if at least one container $i \in \mathcal{C}_j$ from bill of lading j is routed to facility ℓ using transportation mode $m \in \mathcal{M}^{RT}$, and 0 otherwise; **Special transfer** decisions are captured by integer variables $X_{t\ell\ell' m}$ representing the number of full trucks sent from facility ℓ to facility ℓ' on day t using expedition mode $m \in \mathcal{M}^{SP}$, binary variables $x_{t\ell\ell' m}$ equal to 1 if a less-than-full truck is used between ℓ and ℓ' on day t with mode m and 0 otherwise, and continuous variables $w_{t\ell\ell' m} \leq Q$ representing the number of parts carried in that truck[12]; **Milk run transfer** decisions are captured by integer variables $r_{t\ell\ell'}$ representing the number of pallets included in the run from

---

[12]   The integrality of $x_{t\ell\ell' m}$ is immaterial in light of the quantities at stake here.

facility ℓ to facility ℓ' on day t; **Inventory variables** include the expected net inventory level $I_{t\ell}$ at the beginning of day $t > 1$ in location ℓ, its positive part $I_{t\ell}^+$ and negative part $I_{t\ell}^-$, and a binary indicator variable $I_{t\ell}^1 = 1_{\{I_{t\ell}\geq 0\}}$; **Expected average shortages** during each day t in each location ℓ are captured by continuous variables $v_{t\ell}$ (approximately, see §4.1.1).

Formulation:

$$
\text{Min} \sum_{i\in\mathcal{C}^{RT},\ell,m\in\mathcal{M}^{RT}} c_{\ell m}^{RT} y_{i\ell m} + \sum_{j\in\mathcal{J}} c^{BL}\left(\sum_{\ell,m\in\mathcal{M}^{RT}} z_{j\ell m} - 1\right)
$$

$$
+ \sum_{\{t,\ell,\ell',m\in\mathcal{M}^{SP}:\ell\neq\ell'\}} c_{\ell\ell'm}^{SP}\left(X_{t\ell\ell'm} + x_{t\ell\ell'm}\right) + \sum_{\{t,\ell,\ell':\ell\neq\ell'\}} c_{\ell\ell'}^{MR} r_{t\ell\ell'} + B\sum_{t,\ell} v_{t\ell} \tag{6}
$$

subject to:

$$
I_{t\ell} = I_{1\ell} - \bar{f}_{(t-1)\ell} + \sum_{\{i\in\mathcal{C}^{NRT}:\ell_i=\ell,:A_i\leq t-1\}} q_i + \sum_{\{(i,m)\in\mathcal{C}^{RT}\times\mathcal{M}^{RT}:A_i^{LB}+L_{\ell m}^{RT}\leq t-1\}} q_i y_{i\ell m}
$$

$$
+ \sum_{\{(\tau,\ell',m)\in\mathcal{T}\times\mathcal{L}\times\mathcal{M}^{SP}:\ell'\neq\ell,\tau+L_{\ell'\ell m}^{SP}\leq t-1\}} \left(QX_{\tau\ell'\ell m} + w_{\tau\ell'\ell m}\right) + \sum_{\{(\tau,\ell')\in\mathcal{T}\times\mathcal{L}:\ell'\neq\ell,\tau+L_{\ell'\ell}^{MR}\leq t-1\}} Jr_{\tau\ell'\ell}
$$

$$
- \sum_{\{(\tau,\ell',m)\in\mathcal{T}\times\mathcal{L}\times\mathcal{M}^{SP}:\ell'\neq\ell,\tau\leq t-1\}} \left(QX_{\tau\ell\ell'm} + w_{\tau\ell\ell'm}\right) - \sum_{\{(\tau,\ell')\in\mathcal{T}\times\mathcal{L}:\ell'\neq\ell,\tau\leq t-1\}} Jr_{\tau\ell\ell'} \tag{7}
$$

$$
\sum_{m\in\mathcal{M}^{RT},\ell} y_{i\ell m} = 1 \qquad \forall i \in \mathcal{C}^{RT} \tag{8}
$$

$$
z_{j\ell m} \geq y_{i\ell m} \qquad \forall j \in \mathcal{J}, \ell \in \mathcal{L}, m \in \mathcal{M}^{RT}, i \in \mathcal{C}_j \tag{9}
$$

$$
I_{t\ell} = I_{t\ell}^+ - I_{t\ell}^- \qquad \forall t \in \mathcal{T}, \ell \in \mathcal{L} \tag{10}
$$

$$
I_{t\ell}^+ \leq M_{t\ell}.I_{t\ell}^1 \qquad \forall t \in \mathcal{T}, \ell \in \mathcal{L} \tag{11}
$$

$$
I_{t\ell}^- \leq M_{t\ell}.(1 - I_{t\ell}^1) \qquad \forall t \in \mathcal{T}, \ell \in \mathcal{L} \tag{12}
$$

$$
\sum_{\{(\ell',m)\in\mathcal{T}\times\mathcal{L}\times\mathcal{M}^{SP}:\ell'\neq\ell\}} \left(QX_{t\ell\ell'm} + w_{t\ell\ell'm}\right) + \sum_{\{\ell'\in\mathcal{L}:\ell'\neq\ell\}} Jr_{t\ell\ell'} \leq I_{t\ell}^+ \qquad \forall t \in \mathcal{T}, \ell \in \mathcal{L} \tag{13}
$$

$$
w_{t\ell\ell'm} \leq Q.x_{t\ell\ell'm} \qquad \forall t \in \mathcal{T}, (\ell,\ell') \in \mathcal{L}^2, m \in \mathcal{M}^{SP} \tag{14}
$$

$$
r_{t\ell\ell'} \leq R.S_{t\ell\ell'}^{MR} \qquad \forall t \in \mathcal{T}, (\ell,\ell') \in \mathcal{L}^2 \tag{15}
$$

$$
v_{t\ell} \geq a_{t\ell p}(f_{tl} - I_{t\ell}) + b_{t\ell p} \qquad \forall t \in \mathcal{T}, \ell \in \mathcal{L}, p \in \mathcal{P}_{t\ell} \tag{16}
$$

$$
y_{i\ell m}, z_{j\ell m}, x_{t\ell\ell'm}, I_{t\ell}^1 \in \{0,1\} \tag{17}
$$

$$
X_{t\ell\ell'm}, r_{t\ell\ell'} \in \mathbb{N} \tag{18}
$$

$$
w_{t\ell\ell'm}, v_{t\ell}, I_{t\ell}, I_{t\ell}^+, I_{t\ell}^- \geq 0 \tag{19}
$$

The objective (6) is the sum of all transportation costs associated with the decisions con-

sidered, including container re-routing (first term), bill of lading fees (second term), special trucks (third term) and milk runs (fourth term), along with the corresponding expected shortage costs (last term). Note that our choice of minimizing total costs, as opposed to say minimizing transportation costs subject to a service level constraint on total expected shortages, is dictated by the problem at stake. Specifically, Dell's suppliers are responsible for all initial shipment decisions (see §1), which are thus exogenous to the routing problem considered. As a result, such a service level constraint could lead to infeasibility problems. We also observe that the transportation costs captured in (6) do not include all inbound transportation costs paid by Dell, and in particular exclude the cost of ocean transportation – this is because only the portion corresponding to ground transportation is affected by the decision variables. Finally, (6) does not account for any inventory costs which could arise from excessive inventory in a given location. While it would be straightforward to add a term summing the on-hand inventory variables $I_{t\ell}^+$ multiplied by an inventory holding cost rate, it turns out that the relevant costs associated with excessive inventory mostly stem here from the additional storage required in the warehouses adjacent to its factories when the overall amount of inventory across all parts exceeds a threshold. While the inventory holding costs incurred by Dell's suppliers in those warehouses may in turn affect Dell in important ways, these primarily depend on the overall quantity of inventory shipped (as opposed to the allocation of this inventory across sites), which is an exogenous quantity. In light of these considerations and because the inventory storage costs incurred historically represent only a very small fraction of the transportation costs, it was decided to leave these costs out of the optimization model.

Constraints (7) are inventory balance equations defining the relationship between the expected net inventory variables $I_{t\ell}$ and the inventory currently available ($I_{t\ell}$), the demand forecasts ($\bar{f}_{(t-1)\ell}$), the pipeline of non-routable containers ($\sum_{\{i \in \mathcal{C}^{NRT}:\ell_i=\ell,:A_i \leq t-1\}} q_i$), and all the supply routing decisions considered (all subsequent terms in the r.h.s.). Constraints (8) ensure that every container is routed to exactly one destination through one transportation mode. Constraints (9) ensure that the term $\sum_{\ell,m} z_{j\ell m} - 1$ appearing in the objective corresponds indeed to the number of new bills of lading created for the containers initially included in bill of lading j as a result of the routing decisions. Constraints (10)-(12) ensure that variables $I_{t\ell}^+$, $I_{t\ell}^-$ and $I_{t\ell}^1$ correspond indeed to the positive part, negative part and

non-negativity indicator of variable $I_{t\ell}$, respectively. Constraints (13) state that the total inventory transferred out of any facility ℓ during a given day t, either through special trucks or a milk run, may not exceed the inventory on hand expected to be available in that facility at the beginning of that day. Constraints (14) ensure both that the quantity of parts recommended for transportation aboard a less-than-full specially affreighted truck does not exceed its capacity, and that the binary variables signalling the existence of such trucks take values consistent with their definition. Similarly, (15) enforces both the capacity and the scheduling restrictions of milk runs between facilities. Note that the variables $S_{t\ell\ell'}^{MR}$ are only introduced here to simplify exposition, as for implementation purposes it is more computationally efficient to only define variables $r_{t\ell\ell'}$ over the set of indices (t, ℓ, ℓ') such that there exists a run from ℓ to ℓ' on day t. Finally, constraints (16) together with the minimization objective ensure that in any optimal solution (and any solution computed through a branch and bound MIP algorithm) the variables $v_{t\ell}$ implement indeed the approximate expected average shortage level during day t in location ℓ which is described in §4.1.1.

An important observation is that the formulation (6)-(19) only considers one part at a time, and thus implicitly assumes that parts are independent from each other. This is partly justified for monitors and chassis, which are shipped in full containers of a single part type. It does represent a limitation of our approach however, and we return to this issue in §5.

4.2 Implementation   The technical implementation of the model described in §4.1 was performed using the development environment OPL Studio linked with the optimization engine CPLEX 9.1, using Microsoft Excel as a repository for the static input data (costs, lead-times, forecast accuracy parameters, shortage costs) and also to visualize the output data (see Figure 3 for an illustration[13]). In addition, links were created with some of Dell's existing databases in order to automatically import the dynamic input data (current inventory levels, forecasts, pipeline inventory) whenever required. An unexpected data acquisition challenge was created by the difficulty to obtain accurate rail transportation costs. This stemmed from the division of Dell's total inbound transportation costs for accounting purposes in two categories: *embedded* costs which correspond to a default transportation mode (for monitors, transit by ocean and rail) and are included in the price per part charged by suppliers; and

---

[13]   Figure 3 note: The entries "Red Ball" appearing in the 7th column of the table under the heading "Mode" refer to the name of a carrier contracted by Dell to perform the milk runs between sites described in §1, which has become synonymous with that transportation mode within Dell.

re-positioning costs which include any additional ground transportation costs incurred (i.e. expeditions by truck and transfers) and are paid separately. Because it was known however that the costs of transportation by rail between Long Beach and Dell's US facilities only differ by very small amounts relative to all the other ground transportation costs and fees, it was decided that (6) would only capture re-positioning costs. After checking through sensitivity analysis that this would have little impact if any on the decisions recommended, we thus set $c_{\ell m}^{RT} = 0$ for $m =$"rail". Finally, the pre-processing necessary to compute the piecewise linear approximations to the expected shortage cost functions (see §4.1.1) was implemented using Microsoft Visual Basic. The creation of this software tool from complete specifications required approximately 6 months of full-time work by an experienced developer familiar with optimization theory at an introductory graduate course level. We refer the reader to §A.2 in the Online Appendix for a more detailed description of this software (including additional screen copies of its interface), and to Foreman (2008) for the source code.

Enact Decisions

Part #	Decision	Cont #	BOL #	Orig	Dest	Mode	QTY	QTY parts	Time Sensitivity
GY946	DIV	MERH7845625	UYD845972	Nashville	Austin	Rail	Container	1560	3
GY946	DIV	JDRB8432575	UYD845972	Nashville	Austin	Rail	Container	1560	3
GY946	DIV	JDFB6485462	UYD845972	Nashville	Reno	Rail	Container	1560	3
DF923	DIV	WDKJ8762485	JRB174585	Austin	Austin	Team Truck	Container	2184	4
DF923	DIV	FKEI8432675	JRB174585	Austin	Austin	Team Truck	Container	2184	4
DF923	DIV	CEGS7918425	JRB174585	Austin	Austin	Team Truck	Container	2184	4
DF923	RB	N/A	N/A	Nashville	Austin	Red Ball	5 pallets	420	3
DF923	RB	N/A	N/A	Nashville	Austin	Red Ball	8 pallets	672	7
DF923	RB	N/A	N/A	Nashville	Austin	Red Ball	8 pallets	672	10
DF923	TRK TRNS	N/A	N/A	Winston Salem	Austin	Truck	1 ftl 1 ltl	4368	1
DF923	TRK TRNS	N/A	N/A	Nashville	Austin	Truck	1 ltl	2173	6
DF923	TRK TRNS	N/A	N/A	Winston Salem	Austin	Team Truck	3 ftl 1 ltl	8620	6
DF923	TRK TRNS	N/A	N/A	Winston Salem	Austin	Team Truck	1 ftl	2184	11
FG645	DIV	MERH8141846	GUYD269447	Austin	Reno	Rail	Container	616	1
FG645	DIV	FWGR2694487	GUYD269447	Austin	Reno	Rail	Container	616	1
FG645	DIV	JDRB8174585	GUYD269447	Austin	Reno	Rail	Container	1392	1

Enact Selected   Cancel

Figure 3: Output Interface of the Model Software Implementation

It was decided at the outset that the model implementation would be initially limited to monitors, which are very large contributors to Dell's total inbound transportation costs. A first preliminary step consisted of gathering a large and representative collection of input data sets in order to evaluate the computational time associated with executing the branch and bound algorithm on realistic problem instances, and possibly identify the run parameters for this algorithm yielding the fastest computations. This study demonstrated that while achieving optimality occasionally required more than an hour for these problems, a subopti-

mality gap equivalent to a bill of lading creation fee (the smallest individual transportation cost component appearing in the objective function (6)) was almost always achieved in a matter of seconds using standard search strategies. As a result, the achievement of a sub-optimality gap equal to that amount was set as our algorithm termination criteria, and we did not further investigate the computational solution time for this problem, nor attempt to reduce it through additional efforts.

An important implementation issue was to decide how the model output data should affect operational decisions. While each optimization run produces a set of recommended supply routing decisions for each part, some of these are decisions that can possibly be only enacted on some distant day in the future, since the optimization model considers a rolling horizon of several weeks. As seen in the last column of the model output table in Figure 3, the software was thus made to calculate and display the *time sensitivity* of each decision, defined as the number of days before the opportunity to enact that decision will disappear (e.g. a generated diversion decision affecting a container on a vessel five days away from Long Beach would have a time sensitivity of three days if the diversion cutoff point for this part was two days before port). This enabled the analysts to only enact the decisions with a time criticality lower than a set threshold (for example the number of days before the next anticipated run on that part), with the overall goal of waiting for as long as possible for the most recent data before committing to any decision.

The ultimate execution plan envisioned consisted of performing an optimization run for every part on a daily basis. However, a key aspect of the implementation strategy was to first go through a pilot period of several weeks during which the model output would be systematically compared to the supply routing decisions generated manually with the Balance Tool beforehand, with the following objectives: (i) improve the software interface and functionalities with observations grounded in practice; and (ii) build an archive of input and output data in order to evaluate the qualitative and quantitative impact of the model (see §4.3). An important implementation hurdle faced at that point was to determine which value(s) should be used in practice for the unit shortage cost rate B introduced in §4.1. As the model's main input parameter for resolving the trade-off between shortage and transportation costs, it had a significant impact on the output: with a low value of B most decisions generated could be container diversions with no expeditions and some milk-run transfers,

while in the same situation a high value of B could generate many container expeditions and transfers through team trucks. At the same time however, no study previously performed within Dell was available to guide the implementation team towards an objective value for that parameter. The strategy decided then was two-fold: For the long term, an in-depth study of Dell's shortage costs was initiated, following a methodology similar to that described by Oral et al. (1972) (see Dhalla 2008 for more details); In the short term, B was to be treated as a control lever which the supply chain analyst could initially adjust, with the goal of achieving through experimentation the same trade-off between transportation costs and service level as the one that was implicitly associated with the decisions made to date with the Balance Tool. That short term strategy clearly would not be optimal with respect to the key trade-off just mentioned. However, it would still hopefully generate in a more efficient manner recommended supply routing decisions that would be consistent in terms of that trade-off, at least after the initial experimentation phase would converge to a relatively stable set of values for B. In addition, these decisions could still possibly produce substantial savings in transportation costs.

The determination by the supply analyst of a value for the unit cost rate per shortage B achieving a *status quo* in the sense just stated proved more difficult than anticipated. In other words, the experimentation process described above did not seem to converge very fast, at least initially. Through extensive interviews and analysis of specific cases, we became convinced that this was due to intrinsic differences in the metrics respectively used by the model and by the analyst when assessing expected shortages. Specifically, the analyst would primarily evaluate the criticality of a given supply situation by inspecting on the Balance Tool the DSI levels projected in all of Dell's facilities over the planning horizon, as described in §3.1. Given only limited time and input data, the analyst thus relied on an empirical notion of the relationship between these DSI levels and the corresponding expected shortages. While that approach usually provided a good qualitative appreciation of the criticality of the situation in any given site, it sometimes seemed too coarse for correct comparisons of the situation across sites and time periods, which are needed to inform inventory balancing and expedition decisions. In particular, the implicit empirical relationship just mentioned seemed to occasionally ignore the increase in variability of the projected net inventory levels with time, which is driven by the error of the cumulative demand forecast. Also, it sometimes

failed to capture the differences of that variability across sites, which could be significant because sites could face fairly distinct demand patterns. For example, one of the facilities in Dell's network is used to fulfill a larger proportion of the "option" orders for monitors only (i.e. without a computer system), which are harder to forecast accurately than regular orders. In contrast, the model would base its assessment of expected shortages on a stochastic model informed by historical forecast accuracy data, which for example captured the two effects just described through the dependence of its input parameter $\sigma_{t\ell}$ on both time t and location ℓ, respectively (see §4.1.1). Incidentally, this model feature was perhaps the most difficult for the implementation team to convey to relevant stakeholders within Dell. This stemmed occasionally from a lack of familiarity with probability concepts, but more generally because the use of quantitative forecast accuracy data represented for Dell a cultural departure from a pervasive reliance on DSI levels as the sole relevant explicit information when estimating future shortages.

In the end, the implementation team resolved the question of which values for B should be chosen in order to implement the current implicit shortage costs through a study of historical data. Specifically, we constructed a database where each entry corresponds to a set of routing decisions made by the analyst for a given part on a given day, and includes both the associated transportation costs as well as the corresponding reduction in total expected shortages over the planning horizon, as estimated by the model using all relevant input data available at the time. As Figure 4 illustrates, we then performed a linear regression with forced zero intercept of the reduction in expected shortages achieved (dependent variable) as a function of the re-positioning transportation costs incurred (independent variable) for each part over that dataset, which spanned several weeks of decisions.An interesting aspect of these regressions is that their fit provided a measure for the consistency of the analysts' historical decisions with respect to the trade-off between transportation costs and expected shortages, as determined through our stochastic evaluation model. From this standpoint, it was found that these regressions yielded in fact a fit with the data which was higher than expected, as reflected by their relatively high R^2 values (the value of .75 reported in Figure 4 is typical). Consequently, we decided to use their slopes as an (inverse) estimate of the unit shortage cost rate B corresponding to the current implicit trade-off. These new starting points greatly facilitated the determination of which unit shortage cost rate values should
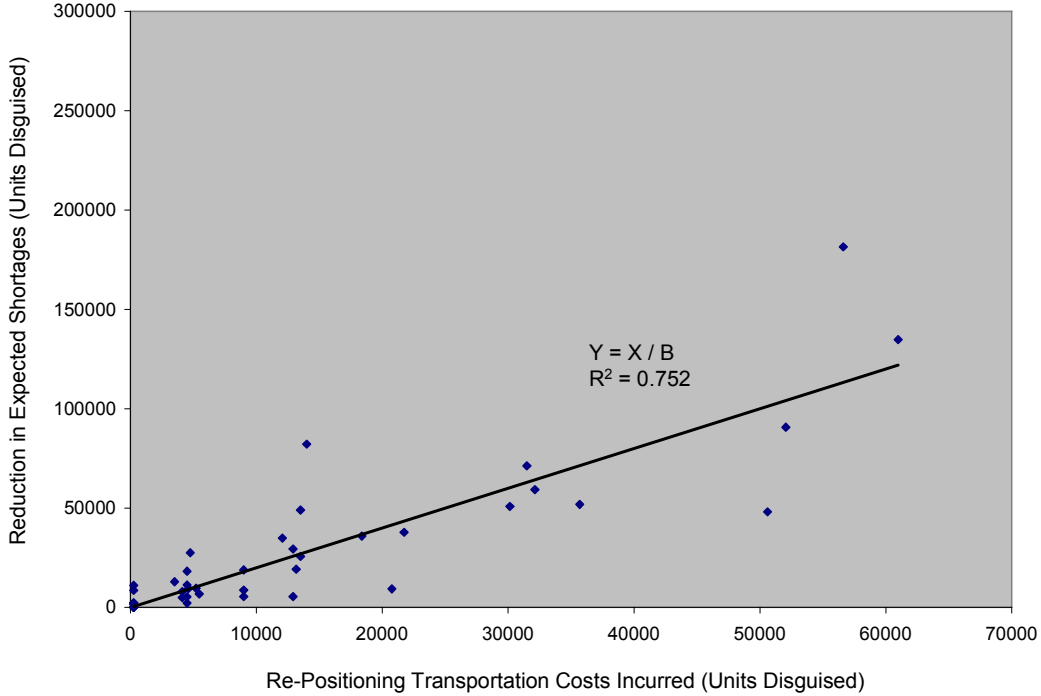
Figure 4: Linear Regression of Expected Shortage Reductions Achieved Against Re-Routing Transportation Costs Incurred

be used initially.

The systematic comparison between the model output and the decisions generated manually by the analyst beforehand also led to several improvements of the implemented software. A first improvement consisted of eliminating the "flipped expeditions" initially observed as part of the model output. This would arise when two sites in short supply were scheduled to receive at some point in the future some containers loaded by a common supplier in the same ship, and therefore with the same expected port arrival date. As illustrated in Figure 5, the model could then recommend to use expedited ground transportation (e.g. team truck) for all containers, but also switch the containers' destinations. We found out however that, for reasons not captured by the initial model (an expedition decision entails a bill of lading creation expense independently of the chosen destination), both the carriers and the supply chain analysts prefer the simpler communications associated with a small number of destination changes, provided this does not impact transportation costs. To capture this preference we introduced the additional objective function term $\sum_{i \in \mathcal{C}^{RT}, \ell \in \mathcal{L} \setminus \{\ell_i\}} y_{i\ell m}$, which essentially adds a dollar penalty for such destination changes. This modification indeed eliminated all such "flipped expeditions" without affecting the transportation costs of computed solutions,

and thus improved the simplicity of the model output.

Part #	Decision	Cont #	BOL #	Orig	Dest	Mode		QTY	QTY parts	Time Sensitivity
657FT	DIV	YRLF8463748	RYT549786	Nashville	Austin	Team Truck	Container		2184	0
657FT	DIV	HCEF5798165	RYT549722	Austin	Nashville	Team Truck	Container		2184	0

Figure 5: Example of Model Output with Swapped Container Destinations

Another feature addition was prompted by the realization that the routing analysts occasionally found issues with the demand forecasts imported by the model from Dell databases, in particular those covering the next 7 days of demand. In the analysts' opinion, these were not updated by the forecasting team as frequently as desirable for daily decision-making, at least at the time when the pilot was run. As a result, the routing analysts often took it upon themselves to correct these forecasts so they would better reflect the actual parts consumption patterns they had observed over the previous days, which could of course lead to drastically different routing (in particular transfer) decisions relative to those recommended by the model with the original forecasts. In addition, our review of the historical demand and forecast data suggested that these ad-hoc forecast corrections often seemed justified. From an organizational standpoint, we believe that such forecast corrections are better done centrally by the forecasting team, perhaps by making better use of relevant decentralized input data such as these recent actual parts consumption patterns. However, we also recognized that some hurdles pertaining to the implementation of such coordination between the forecasting and the procurement teams would likely take some time to overcome. This created a need for the software to support the ad-hoc forecast correction practice just described. Specifically, we added a feature whereby the historical consumption of each part in every site is stored in a database covering the past ten days of actual demand, and any major discrepancy between some time series-based forecasts constructed from that database and those provided by the forecasting team is automatically highlighted. The analyst can then decide to automatically modify the model input data $\bar{f}_{t\ell}$ by replacing the original forecast for the next 7 days of the horizon with the alternative one based on time series calculations.

Finally, a particularly topical question at the beginning of the pilot was to determine how large orders from retailers distributing Dell's computers should be captured by the model. That question arose in a context of strategic change for Dell, which in 2007 started a far-reaching transformation involving some distribution partnerships with large retailers,

whereas it had almost exclusively relied on direct channels to customers until then. As a result, large customer orders for a single type of computer with some advance lead-time became more frequent. In particular, the supply routing analysts were starting to receive notes informing them of committed schedules of large retailer deliveries for specific parts, which they were asked to plan for in addition to the existing forecasts for direct channels. The approach followed to account for these special orders with the model was initially the same as that employed as part of the manual decision process supported by the Balance Tool, and consisted of simply adding these large customer orders to the existing forecasts for the day of their shipment deadline. That method however resulted in transfers and diversions to sites with large retail orders which were sometimes thought to be excessive. We determined that this resulted from a substantial overestimation of demand variability (and therefore expected shortages) in those sites, as the original demand model resulting from our forecast accuracy study evaluated the standard deviation of (cumulative) demand $\sigma_{t\ell}$ as a coefficient of variation coefficient times the corresponding forecast value $\bar{f}_{t\ell}$ (see §A.1 in the Online Appendix). This did not reflect the fact that these special retail orders have a substantially lower associated uncertainty than the direct channel orders. In order to address this issue, we created a feature to capture these special orders by modifying the means of demand forecasts $\bar{f}_{t\ell}$ correspondingly, but without affecting the forecast standard deviations $\sigma_{t\ell}$ (see §4.1.1 and §A.1 in the Online Appendix). This drastically reduced the seemingly unnecessary diversions and transfers.

4.3 Impact   The quantitative impact evaluation of the model implementation described in §4.2 had an important methodological requirement, which was to account for any effects on both transportation costs and part shortages – a reduction in transportation costs alone is easily obtained by eliminating all ground expedition modes for example, and may thus not represent an improvement if it is associated with an increase of shortages. Conversely, only using team trucks for all ground transportation would likely reduce part shortages, but also drastically increase transportation costs. In order to construct an unambiguous measure of overall impact, one method considered was to use the current implicit shortage cost rate B (see §4.2) in order to estimate shortage costs, and then measure any changes in the sum of transportation and shortage costs. Out of concern that that shortage cost rate was affected by subjective factors however, Dell executives expressed that it would be desirable to not

rely on its inferred value for impact evaluation purposes.

For this reason, we followed an alternative methodology consisting of computing a posteriori the reduction of re-positioning transportation costs achieved by the optimization model relative to the legacy process, under the additional constraint that its output may not result in higher shortages than that achieved historically. More specifically, for a representative group of monitors $\mathcal{K}$ representing approximately half of total monitor sales and over a period of 14 weeks in 2007 preceding the implementation of the optimization-based process for those monitors, we recorded every individual routing decision made by the analysts using the existing manual process and the Balance Tool described in §3, along with all the corresponding input data (inventory, forecasts, supply line) available at the time when these decisions were made. From that database, we were thus able to construct an instance of the optimization problem (6)-(19) for every week that the analysts made a set of routing decisions for each monitor k within that group. Note that the set of historical routing decisions recorded that week along with their corresponding expected shortage variables $\hat{v}_{t\ell}^k$ (and associated secondary variables) constitute a feasible solution to that problem instance, with re-positioning transportation cost $\hat{C}_k$ and total objective value $\hat{C}_k + B \sum_{t,\ell} \hat{v}_{t\ell}^k$. Our impact assessment was then based on the solution to the modified optimization problem obtained by minimizing only the transportation cost components of (6) subject to the previous constraints (7)-(19) along with the additional constraint that $\sum_{t,\ell} v_{t\ell}^k \leq \sum_{t,\ell} \hat{v}_{t\ell}^k$. Denoting by C_k the optimal value of that modified objective (i.e. the lowest re-positioning transportation costs achievable when allowing no more expected shortages than achieved historically), Figure 6 contains a plot of the weekly re-positioning transportation costs $\sum_{k \in \mathcal{K}} \hat{C}_k$ incurred historically for all these parts as well as data labels indicating the corresponding relative total reduction $\frac{\sum_{k \in \mathcal{K}} \hat{C}_k - C_k}{\sum_{k \in \mathcal{K}} \hat{C}_k}$ achieved by the optimization model[14].

When summed over all 14 weeks of the data collection period defined above, the cumulated transportation cost savings associated with these optimization model runs represent approximately 46% of the total incurred historically, which provides an aggregate measure for the impact of this implementation. However, these relative savings seem to depend on the overall scarcity of supply, which is driven by the total quantities of components shipped by suppliers

---

[14] While the qualifyer re-positioning may occasionally be omitted for brevity in this section, we emphasize that, as stated in §4.2, the only transportation costs considered here are re-positioning costs which do not include ocean and rail transportation.
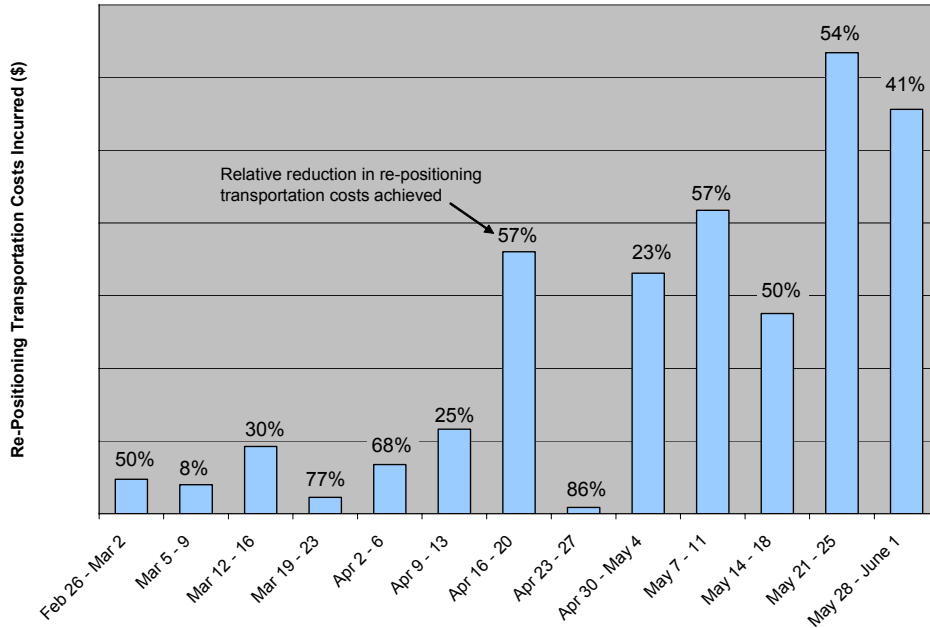
Figure 6: Weekly Re-Positioning Transportation Costs Incurred and Relative Reduction Achieved by the Optimization Model for Six Monitors from February 26 to June 1, 2007

relative to demand and is thus exogenous to the routing model considered here. This can be seen from Figure 6, where the average weekly transportation costs plotted increase several folds in the second half of the data collection period (April 16 – June 1) compared to its first half (February 26 – April 13). This increase corresponds to an industry-wide shortage of glass substrates and color filters which began to impact the deliveries of flat panel monitors by Dell's suppliers in the middle of April that year (Uno 2008), and in turn resulted in additional transportation costs (in particular expeditions). This affected the corresponding relative transportation cost savings, which can be evaluated independently for the first and second halves of the data collection period at 38% and 48% respectively. These observations suggest that the lower of these last two numbers constitutes a better estimation for the relative transportation cost savings attributable to the optimization model during normal periods characterized by appropriate overall supply quantities. It is noteworthy however that the relative benefits derived from the optimization model seem to increase during drastic shortage situations; our explanation of this observation is that under the legacy process, the analysts are typically required then to execute a higher number of routing decisions every day, leaving them with less time for performing extensive analyses of these decisions.

More generally, we wanted to identify the main qualitative reasons explaining the cost

26

savings attributed to the optimization model. This led us to inspect the output of many of the optimization runs we conducted a posteriori as described above, and compare them with the historical transportation decisions made by the analysts with the same input data. Although we cannot provide an exhaustive description of these qualitative comparisons due to space constraints, the two representative examples illustrated by Figures 7 to 10 convey the main insights we obtained. Figure 7 shows a disguised but qualitatively representative version of the Balance Tool interface for a specific 15 inch monitor and a portion of the planning horizon as it appeared to the analyst on March 13, 2007. It shows a situation with an apparent excess of inventory relative to predicted demand in Nashville and Winston-Salem (NCO), and a shortage of inventory appearing in Austin and Reno (RFC) at some point over the horizon considered. The situation in Austin would be particularly preoccupying at that point, as the shortages there are predicted to be higher and occur sooner than in Reno, which is only attributed a small demand forecast. Indeed, the (disguised) total number of expected shortages across all sites and days in the (complete) horizon predicted by our shortage model in the case where no action would be taken then is 50,000 unit-days of shortages (i.e. a measurement corresponding for example to predicted shortages of 2,500 units across all Dell sites in each day of a 20 day horizon). Note also that no upcoming deliveries of containers by suppliers for that component are visible within the planning horizon, leaving transfers as the only supply routing decisions available.

On that day, the analyst decided in fact to order a transfer of 5000 parts from Winston-Salem (NCO) to Austin with three full specially affreighted team trucks, for a (disguised) cost of $30,000. NCO was chosen as the location providing inventory because it had the largest amount of inventory available, both in absolute terms and when evaluated through DSI levels. Also, note that NCO has a forecasted demand about 30% lower than that of Nashville over the horizon considered, so that a transfer of a given quantity out of that facility results in a larger decrease of its DSI level, a metric which is closely monitored. Finally, observe that no inventory was transferred to Reno, presumably because the potential corresponding transportation costs were not justified by the minor and distant predicted shortages at stake in that location. These decisions therefore suggest a qualitatively correct appreciation by the analyst of the overall directions, criticality and time-sensitivity of inventory imbalances across sites, and indeed decreased by 59% the total expected shortages

		14-Mar Wed	15-Mar Thu	16-Mar Fri	19-Mar Mon	20-Mar Tue	21-Mar Wed	22-Mar Thu	23-Mar Fri	26-Mar Mon	27-Mar Tue	28-Mar Wed	29-Mar Thu	30-Mar Fri	2-Apr Mon	3-Apr Tue
AUSTIN / AFC	Demand (from FSS)	566	566	566	656	656	656	656	656	656	656	656	656	656	656	656
	Demand Adjustment															
	AVL	6,000	5,434	4,868	4,302	3,646	2,990	2,335	1,679	1,023	367	-289	-945	-1,60˙	-2,257	-2,913
	Options Adjustment															
	Delivery (Tracking File)															
	Delivery Adj. (w/ Opt)															
	Delivery Adj. (no Opt)															
	DELTA	5,434	4,868	4,302	3,646	2,990	2,335	1,679	1,023	367	-289	-945	-1,601	-2,257	-2,913	-3,569
	DSI	9.6	8.6	7.6	5.6	4.6	3.6	2.6	1.6	0.6	-0.4	-1.4	-2.4	-3.4	-4.4	-5.4
Nashville / NFC	Demand (from FSS)	348	348	348	393	393	393	393	393	393	393	393	393	393	393	393
	Demand Adjustment															
	AVL	19,320	18,972	18,624	18,276	17,883	17,490	17,097	16,704	16,311	15,918	15,526	15,133	14,740	14,347	13,954
	Options Adjustment															
	Delivery (Tracking File)															
	Delivery Adj. (w/ Opt)															
	Delivery Adj. (no Opt)															
	DELTA	18,972	18,624	18,276	17,883	17,490	17,097	16,704	16,311	15,918	15,526	15,133	14,740	14,347	13,954	13,561
	DSI	54.5	53.5	52.5	45.5	44.5	43.5	42.5	41.5	40.5	39.5	38.5	37.5	36.5	35.5	34.5
RFC	Demand (from FSS)	33	33	33	41	41	41	41	41	41	41	41	41	4˙	41	41
	Demand Adjustment															
	AVL	639	606	573	540	499	458	417	375	334	293	252	211	169	128	87
	Options Adjustment															
	Delivery (Tracking File)															
	Delivery Adj. (w/ Opt)															
	Delivery Adj. (no Opt)															
	DELTA	606	573	540	499	458	417	375	334	293	252	211	169	128	87	46
	DSI	18.5	17.5	16.5	12.1	11.1	10.1	9.1	8.1	7.1	6.1	5.1	4.1	3.1	2.1	1.1
NCO	Demand (from FSS)	226	226	226	240	240	240	240	240	240	240	240	240	240	240	240
	Demand Adjustment															
	AVL	21,209	20,983	20,757	20,530	20,290	20,050	19,810	19,570	19,330	19,090	18,850	18,609	18,369	18,129	17,889
	Options Adjustment															
	Delivery (Tracking File)															
	Delivery Adj. (w/ Opt)															
	Delivery Adj. (no Opt)															
	DELTA	20,983	20,757	20,530	20,290	20,050	19,810	19,570	19,330	19,090	18,850	18,609	18,369	18,129	17,889	17,649
	DSI	92.8	91.8	90.8	84.5	83.5	82.5	81.5	80.5	79.5	78.5	77.5	76.5	75.5	74.5	73.5

Figure 7: Disguised Copy of the Balance Tool Interface for a 15 inch Monitor on March 14, 2007

predicted by our stochastic model, down to about 20,450 unit-days of shortages (note that because the overall supply quantity is exogenous, in many situations such as this one routing decisions may not reduce expected shortages below a certain level). In the same situation however, the optimization model recommended two regular truck transfers of 1,665 parts each (this quantity corresponds to a full truckload for that part) from Nashville and NCO respectively, along with a schedule of subsequent milk run transfers from Nashville to Austin containing each the maximum number of parts allowed – this solution is illustrated by Figure 8, which also shows the impact of these decisions on the predicted inventory and DSI levels. By construction, that solution achieved the same total expected shortages as the analyst's, however its total transportation cost amounts to $20,010, which represents a 33% reduction relative to the cost incurred historically. Remarkably, the total quantity of inventory transferred to Austin according to that solution (5,175) is very similar to that decided by the analyst, which is a by-product of the additional constraint on expected shortages. However, it exploits the lower transfer cost to Texas from Nashville (Tennessee) than from NCO (North Carolina), and is immune to considerations about the potential perceptions of high DSI levels in NCO – the reason here why the model does not recommend in fact all inventory to be transferred from Nashville is that this would generate more expected shortages for that facility in the later part of the horizon, which is not shown in Figure 8. Another source of

		14-Mar Wed	15-Mar Thu	16-Mar Fri	19-Mar Mon	20-Mar Tue	21-Mar Wed	22-Mar Thu	23-Mar Fri	26-Mar Mon	27-Mar Tue	28-Mar Wed	29-Mar Thu	30-Mar Fri	2-Apr Mon	3-Apr Tue
AUSTIN (AFC)	Demand (from FSS)	566	566	566	656	656	656	656	656	656	656	656	656	656	656	656
	Demand Adjustment															
	AVL	6,000	5,434	4,868	7,992	7,336	7,040	6,385	5,729	5,433	4,777	4,481	3,825	3,169	2,873	2,217
	Options Adjustment															
	Delivery (Tracking File)	Truck														
	Delivery Adj. (w/ Opt)			3330												
	Delivery Adj. (no Opt)			360		360			360		360			360		45
	DELTA	5,434	4,868	7,99	7,336	7,00	6,385	5,729	5,43	4,777	4,41	3,825	3,169	2,83	2,217	1,6
	DSI	9.6	8.6	14.1	11.2	10.7	9.7	8.7	8.3	7.3	6.8	5.8	4.8	4.4	3.4	2.4
NASHVILLE (NFC)	Demand (from FSS)	348	348	348	393	393	393	393	393	393	393	393	393	393	393	393
	Demand Adjustment															
	AVL	19,320	17,307	16,599	16,251	15,498	15,105	14,712	13,959	13,566	12,813	12,421	12,028	11,275	10,882	10,444
	Options Adjustment			RB		RB			RB		RB			RB		RB
	Delivery (Tracking File)															
	Delivery Adj. (w/ Opt)	-1665	-360		-360				-360		-360			-360		-45
	Delivery Adj. (no Opt)															
	DELTA	17,307	16,599	16,251	15,498	15,105	14,712	13,959	13,566	12,813	12,421	12,028	11,275	10,882	10,444	10,051
	DSI	49.7	47.7	46.7	39.4	38.4	37.4	35.5	34.5	32.6	31.6	30.6	28.7	27.7	26.6	25.6
RFC	Demand (from FSS)	33	33	33	41	41	41	41	41	41	41	41	41	41	41	41
	Demand Adjustment															
	AVL	639	606	573	540	499	458	417	375	334	293	252	211	169	128	87
	Options Adjustment															
	Delivery (Tracking File)															
	Delivery Adj. (w/ Opt)															
	Delivery Adj. (no Opt)															
	DELTA	606	573	540	499	458	417	375	334	293	252	211	169	128	87	46
	DSI	18.5	17.5	16.5	12.1	11.1	10.1	9.1	8.1	7.1	6.1	5.1	4.1	3.1	2.1	1.1
NCO	Demand (from FSS)	226	226	226	240	240	240	240	240	240	240	240	240	240	240	240
	Demand Adjustment															
	AVL	21,209	19,318	19,092	18,865	18,625	18,385	18,145	17,905	17,665	17,425	17,185	16,944	16,704	16,464	16,224
	Options Adjustment															
	Delivery (Tracking File)															
	Delivery Adj. (w/ Opt)	-1665														
	Delivery Adj. (no Opt)															
	DELTA	19,318	19,092	18,865	18,625	18,385	18,145	17,905	17,665	17,425	17,185	16,944	16,704	16,464	16,224	15,984
	DSI	85.4	84.4	83.4	77.6	76.6	75.6	74.6	73.6	72.6	71.6	70.6	69.6	68.6	67.6	66.6

Figure 8: Routing Decisions Recommended by the Optimization Model for the Example Illustrated by Figure 7

cost difference is the use of regular trucks as opposed to team trucks, which results from the model's determination that the corresponding lead-time difference of one day (delivery on March 15 instead of March 16) does not justify this additional cost in light of the predicted inventory situation in Austin over these couple of days – as seen in Figure 7 Austin is still predicted to have 5.6 DSI on March 19 absent any transfer decisions, also this time period (March 15-19) is situated very early in the rolling horizon. As discussed in §4.2, the analysts tend to infer the criticality of shortages based on DSI levels alone, whereas the model also takes into account whether that level is predicted early or late in the planning horizon, which affects the variability of the corresponding cumulative demand forecast, and therefore the estimation of expected shortages. As a result, for a given DSI level the analysts tend to overestimate expected shortages relative to the model in the early part of the horizon, and underestimate them in the more distant part. Finally, the model solution also exploits the lower transportation costs associated with milk run transfers (RB) than with specially affreighted trucks, even though the capacity restrictions of milk run transfers result in a higher number of individual transfer decisions. In addition, milk run transfers for a given leg are only available on specific days, and therefore require the additional step of checking their current weekly schedule. These last observations explains why the analysts, who are subject to time pressure and human cognitive limitations, are unlikely to devise this type of

transportation plan, which is more cost effective but also more complex.

More generally, when working with the Balance Tool alone the analysts tended to order few routing decisions affecting large quantities of parts. As a result, they missed important cost reduction opportunities associated with staggering deliveries and using multiple transportation modes to the same destination. Figure 9 illustrates a representative example, and shows a disguised portion of the Balance Tool interface for a 20 inch monitor in the morning of April 17, 2007. That initial situation is characterized by insufficient inventory in Nashville, with the other facilities showing sufficient inventory levels that are initially comparable in terms of DSI. Also, there are planned container arrivals in Reno on May 7 (960 parts), and in Nashville on May 10 (3564 parts, not visible in Figure 9). Absent any routing decisions in that initial situation, our stochastic model predicts a (disguised) total of 80,000 expected unit-days of shortages over the complete rolling horizon.

	17-Apr Tue	18-Apr Wed	19-Apr Thu	20-Apr Fri	23-Apr Mon	24-Apr Tue	25-Apr Wed	26-Apr Thu	27-Apr Fri	30-Apr Mon	1-May Tue	2-May Wed	3-May Thu	4-May Fri	7-May Mon	8-May Tue
AUSTIN (AFC)																
Demand (from FSS)	612	612	612	612	758	758	758	758	758	775	775	775	775	775	803	803
Demand Adjustment																
AVL	19,569	18,957	18,345	17,733	17,121	16,363	15,605	14,846	14,088	13,330	12,555	11,781	11,006	10,231	9,457	8,654
Options Adjustment																
Delivery (Tracking File)																
Delivery Adj. (w/ Opt)																
Delivery Adj. (no Opt)																
DELTA	18,957	18,345	17,733	17,121	16,363	15,605	14,846	14,088	13,330	12,555	11,781	11,006	10,231	9,457	8,654	7,851
DSI	31.0	30.0	29.0	28.0	21.6	20.6	19.6	18.6	17.6	16.2	15.2	14.2	13.2	12.2	10.8	9.8
Nashville (NFC)																
Demand (from FSS)	478	478	478	478	706	706	706	706	706	724	724	724	724	724	741	741
Demand Adjustment																
AVL	3,832	3,354	2,876	2,398	1,920	1,214	509	-197	-902	-1,608	-2,332	-3,057	-3,781	-4,506	-5,230	-5,971
Options Adjustment																
Delivery (Tracking File)																
Delivery Adj. (w/ Opt)																
Delivery Adj. (no Opt)																
DELTA	3,354	2,876	2,398	1,920	1,214	509	-197	-902	-1,608	-2,332	-3,057	-3,781	-4,506	-5,230	-5,971	-6,712
DSI	7.0	6.0	5.0	4.0	1.7	0.7	-0.3	-1.3	-2.3	-3.2	-4.2	-5.2	-6.2	-7.2	-8.1	-9.1
RFC																
Demand (from FSS)	25	25	25	25	30	30	30	30	30	30	30	30	30	30	25	25
Demand Adjustment																
AVL	663	638	614	589	564	534	504	474	444	414	384	354	324	294	264	1,199
Options Adjustment																
Delivery (Tracking File)																960	
Delivery Adj. (w/ Opt)																
Delivery Adj. (no Opt)																
DELTA	638	614	589	564	534	504	474	444	414	384	354	324	294	264	1,199	1,174
DSI	25.8	24.8	23.8	22.8	17.8	16.8	15.8	14.8	13.8	12.8	11.8	10.8	9.8	8.8	48.0	47.0
NCO																
Demand (from FSS)	320	320	320	320	467	467	467	467	467	486	486	486	486	486	537	537
Demand Adjustment																
AVL	12,311	11,991	11,670	11,350	11,029	10,562	10,095	9,628	9,160	8,693	8,207	7,721	7,235	6,749	6,263	5,726
Options Adjustment																
Delivery (Tracking File)																
Delivery Adj. (w/ Opt)																
Delivery Adj. (no Opt)																
DELTA	11,991	11,670	11,350	11,029	10,562	10,095	9,628	9,160	8,693	8,207	7,721	7,235	6,749	6,263	5,726	5,190
DSI	37.4	36.4	35.4	34.4	22.6	21.6	20.6	19.6	18.6	16.9	15.9	14.9	13.9	12.9	10.7	9.7

Figure 9: Disguised Copy of the Balance Tool Interface for a 20 inch Monitor on April 17, 2007

On that day however, the analyst ordered an immediate transfer of 5000 parts from Austin to Nashville using 4 team trucks, and a ground transportation expedition by team truck of all 3564 parts (3 containers) initially scheduled to arrive in Nashville on May 10, which advanced their arrival date to April 30 (and thus mitigated the predicted shortages in Nashville from April 30 to May 10). The (disguised) total transportation cost of these decisions was $71,400.

		17-Apr Tue	18-Apr Wed	19-Apr Thu	20-Apr Fri	23-Apr Mon	24-Apr Tue	25-Apr Wed	26-Apr Thu	27-Apr Fri	30-Apr Mon	1-May Tue	2-May Wed	3-May Thu	4-May Fri	7-May Mon	8-May Tue
AUSTIN / AFC	Demand (from FSS)	612	612	612	612	758	758	758	758	758	775	775	775	775	775	803	803
	Demand Adjustment																
	AVL	19,569	16,457	15,845	15,058	14,446	13,688	12,930	12,171	11,213	10,455	9,680	8,906	8,131	7,356	6,582	5,779
	Options Adjustment																
	Delivery (Tracking File)	Truck															
	Delivery Adj. (w/ Opt)	-2500		-175 ••••• RB						-200••••• RB							
	Delivery Adj. (no Opt)																
	DELTA	16,457	15,845	15,058	14,446	13,688	12,930	12,171	11,213	10,455	9,680	8,906	8,131	7,356	6,582	5,779	4,976
	DSI	26.9	25.9	24.6	23.6	18.0	17.0	16.0	14.8	13.8	12.5	11.5	10.5	9.5	8.5	7.2	6.2
Nashville / NFC	Demand (from FSS)	478	478	478	478	706	706	706	706	706	724	724	724	724	724	741	741
	Demand Adjustment																
	AVL	3,832	3,354	5,376	7,398	7,095	6,389	5,684	4,978	4,273	3,767	3,043	2,318	1,594	869	145	-596
	Options Adjustment																
	Delivery (Tracking File)													Rail Diversion, Delivery May 14			
	Delivery Adj. (w/ Opt)		2500	2500	175					200							
	Delivery Adj. (no Opt)																
	DELTA	3,354	5,376	7,398	7,095	6,389	5,684	4,978	4,273	3,767	3,043	2,318	1,594	869	145	-596	-1,337
	DSI	7.0	11.2	15.5	14.8	9.1	8.1	7.1	6.1	5.3	4.2	3.2	2.2	1.2	0.2	-0.8	-1.8
RFC / RFC	Demand (from FSS)	25	25	25	25	30	30	30	30	30	30	30	30	30	30	25	25
	Demand Adjustment																
	AVL	663	638	614	589	564	534	504	474	444	414	384	354	324	294	264	239
	Options Adjustment																
	Delivery (Tracking File)															960	
	Delivery Adj. (w/ Opt)															-960	
	Delivery Adj. (no Opt)																
	DELTA	638	614	589	564	534	504	474	444	414	384	354	324	294	264	239	214
	DSI	25.8	24.8	23.8	22.8	17.8	16.8	15.8	14.8	13.8	12.8	11.8	10.8	9.8	8.8	9.6	8.6
NCO / NCO	Demand (from FSS)	320	320	320	320	467	467	467	467	467	486	486	486	486	486	537	537
	Demand Adjustment																
	AVL	12,311	9,491	9,170	8,850	8,529	8,062	7,595	7,128	6,660	6,193	5,707	5,221	4,735	4,249	3,763	3,226
	Options Adjustment																
	Delivery (Tracking File)	Truck															
	Delivery Adj. (w/ Opt)	-2500															
	Delivery Adj. (no Opt)																
	DELTA	9,491	9,170	8,850	8,529	8,062	7,595	7,128	6,660	6,193	5,707	5,221	4,735	4,249	3,763	3,226	2,690
	DSI	29.6	28.6	27.6	26.6	17.3	16.3	15.3	14.3	13.3	11.7	10.7	9.7	8.7	7.7	6.0	5.0

Figure 10: Routing Decisions Recommended by the Optimization Model for the Example Illustrated by Figure 9

The optimization model solution for the same situation is illustrated by Figure 10, and consist of two immediate regular truck transfers of two full trucks each (2,500 parts) from Austin and Winston Salem (NCO) to Nashville, two milk run transfers from Austin to Nashville and a diversion to Nashville by rail of the 980 parts initially scheduled to arrive in Reno on May 7, which postponed their arrival date to May 14 because of the longer lead-time from California to Nashville (see §1 for background). It achieves by construction the same number of expected unit-days of shortages, but costs 53% less than the manual solution implemented historically (or $33,450).

Observe that both the manual and the model solutions involve initial transfers to Nashville of the same quantity of parts (5000). However, the model does not use costly team trucks for these transfers, for reasons that are similar to those explained in the previous example. Also, it spreads the origins of these transfers across two locations (Austin and NCO), which saves many expected shortages in the later part of the horizon in Austin: note that with only 2,875 parts withdrawn from Austin in the model's solution (against 5,000 for the manual one), the last day of the horizon portion shown in Figure 10 (May 8) shows only 6.2 predicted DSI, with continued demand and no subsequent container arrival in Austin in the time horizon beyond that – the situation in Austin from then on is thus significantly worse with the analyst's solution. More generally, this model behavior is consistent with the convexity of expected shortages as a function of the opposite of inventory level, which is an important

31

feature of both our model and the actual situation at stake (see discussion following (5)). Indeed, from that property total expected shortages are lower when the "pain" (that is low inventory levels) is shared across several locations rather than concentrated in one location only. This is also consistent with using two milk run transfers from Austin and a diversion from Reno instead of the simpler but much more expensive expedition by team truck of the containers bound to Nashville in the manual solution. Note also that, in contrast with the model's decisions, that expedition decision by the analyst does not affect the shortages in Nashville beyond May 10 (the original container arrival date), a distant time period with high cumulative demand forecast variability (see interpretation of previous example). Finally, we observe that the cheaper model solutions with multiple origin facilities for inventory transfers and diversions conflict with a second aspect of the former manual decision process besides the analysts' preference for (and possibly ability to handle) only a small number of decisions. Specifically, analysts often needed to quickly evaluate the situation for many different parts and quickly determine whether any specific one deserved some attention. When doing so, they tended to inspect the total number of cells showing in red or yellow on each part's Balance Tool for any day and location because of a low predicted DSI level (see Figure 2), and use that number as an overall indicator of criticality. By extension, they had come to also use that metric as a proxy for total expected shortages when making routing decisions. Indeed, the first reaction of an analyst with whom we shared the model solution illustrated in Figure 10 was that it was a worse solution than the one determined manually, because it happens to entail a larger area of the Balance Tool showing in red. Because of the convexity property just discussed however, that metric can in fact lead to an increase of total expected shortages in some cases, as is shown by the simple example of two locations facing the same demand on a given day with a total of 3 DSI available for both (allocating 1.5 DSI to each minimizes total expected shortages but results in both location showing in red on the Balance Tool, whereas allocating all 3 DSI to a single location only puts the other one in the red). Indeed, we have found several instances in our dataset where, in contrast with the two examples just discussed, the locations receiving inventory according to the analyst and the model solutions are different, with the analyst moving inventory to a location with slightly lower DSI levels but much smaller demand forecasts, because this increased these DSI levels by a larger amount (although these decisions reduced total expected shortages by a smaller

quantity).

5 Conclusion

Although the process and optimization model for supply routing described earlier have been successfully implemented by Dell, this work has several important limitations, all of which motivate ongoing or future research. A first opportunity is the implementation of unit shortage costs resulting from a rigorous evaluation of the main cost components involved. The related study mentioned in §4.2 (Dhalla 2008) is now completed, and has already been used to generate more objective estimates for the value of the unit shortage cost rate B that should be used in optimization model runs. In particular, that study investigated the average margin of customer orders associated with a given part and shipped from a given location, and thus showed how the parameter B should depend not only on the part, but also the location considered – a key factor is that one of the facilities in Dell's network receives a larger proportion of option orders (e.g. for monitors only), for which the cost consequences of delays are milder than for complete system orders. That study also showed that in some cases our (standard) assumption of a linear structure for shortage costs (see §4.1.1) was fairly coarse, in part because the likelihood of order cancellation by a customer does not seem to increase linearly with the number of days of delay relative to the promised delivery date. This motivates ongoing efforts to develop and test an optimization model reflecting these non-linear shortage costs, however because of the significant associated increases in model complexity and data maintenance requirements, it is not clear yet that this work may ultimately affect Dell's practice. Another opportunity would be to capture the dependencies across different parts when generating supply routing decisions. A first avenue would be to extend the current model structure to components which, unlike monitors and chassis, are shipped in mixed containers of several part types. While we did not focus on these "mixed" parts initially because they account for less transportation costs, that extension may still generate substantial savings over time. A more ambitious goal would be to take into account the inventory situation of several components likely to be required by the same customer orders when determining supply routing (and more generally ordering) decisions for each. Interestingly, while the academic literature discusses the potential benefits of this practice (see Song and Zipkin 2003), it does not seem to have impacted operations at Dell yet, in part because of concerns linked to organizational incentives (e.g. two managers

responsible for the supply of different components both saving on expedition costs because of a simultaneous belief that the other manager's component will be short anyway). Finally, another opportunity is to relax the assumption that demand in individual sites is endogenous, i.e. jointly optimize the allocation of customer orders to manufacturing sites and inventory transfer decisions. The approach followed in the present paper seems correct as a first approximation, because Dell ships directly to most of its customers, so that the differences in (unit) outbound shipping costs for complete systems across different manufacturing sites are often substantially larger than the average (bulk) inbound transportation costs for individual components. In certain situations however, for example when transfering customer orders to a different factory may avoid some overtime, such joint optimization could prove profitable.

Despite all these limitations, the financial impact assessment presented earlier (the relative cost reduction estimates of 40% and 38% discussed in §3.3 and §4.3 amount to a cumulative reduction of inventory re-positioning costs for monitors by about 60% since the beginning of this collaboration) suggest that the model described in the present paper is already quite valuable for operational purposes. This is also supported by several recent developments at Dell. Specifically, Dell has committed some resources to implement that model in its European manufacturing network, where the supply chain structure is somewhat more complex because it involves several disembarkation ports where inventory can be held at and re-routed from. In addition, Dell is funding an effort to develop and test an extension of that model to compute recommended quantities, timing and transportation modes for all component shipments between a global Asian warehouse and all of its manufacturing sites worldwide (see Foreman 2008). Finally, we note that many features of the model defined in §4 do not seem specific to Dell, so that part or all of it may also be useful in the future to other firms facing supply routing and/or transportation mode decisions.

From a methodological standpoint, we observed that the small scale but real pilot implementation experiments we organized in each phase of this collaboration (see §3.2, §4.2) were a particularly effective way to improve the process or model tested, but also to motivate and coordinate the joint work of practitioners and academic researchers, who can otherwise be subject to different timelines and environmental constraints. Also, this work gave us an appreciation for the recent advances in computational power, development environments and computational engines for optimization applications, and database software

tools. It led us to revise downwards our estimates of the resources required to implement large optimization-based control systems, which is consistent with recent reports of other similar successful implementations by relatively small teams combining both practitioners and academic researchers (Caro and Gallien 2008, Durbin and Hoffman 2008). This may signal a larger opportunity for many firms to develop their own customized supply-chain control optimization systems, either internally or as part of a collaboration with academia, and possibly a competitive threat to some vendors of generic enterprise software.

References

Caggiano, K. E., J. A. Muckstadt, and J. A. Rappold (2006). Integrated real-time capacity and inventory allocation for reparable service parts in a two-echelon supply system. *Manufacturing Service Operations Management 8*(3), 292–319.

Caro, F. and J. Gallien (2008). Inventory management of a fast-fashion retail network. Technical report, MIT Sloan School of Management, Cambridge, MA.

Chand, S., V. N. Hsu, and S. Sethi (2002). Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing Service Operations Management 4*(1), 25–43.

Dhalla, N. (2008). Evaluating shortage costs in a dynamic environment. Master's thesis, Massachusetts Institute of Technology.

Durbin, M. and K. Hoffman (2008). The dance of the thirty-ton trucks: Dispatching and

scheduling in a dynamic environment. *Operations Research 56*(1), 3–19.

Eppen, G. and L. Schrage (1981). *Centralized Ordering Policies in a Multi-Warehouse System with Leadtimes and Random Demand*, pp. 51–69. North Holland, Amsterdam, The Netherlands.

Federgruen, A. and P. Zipkin (1984). Approximation of dynamic multilocation production and inventory problems. *Management Science 30*(1), 69–84.

Foreman, J. (2008). Optimized supply routing at dell under non-stationary demand. Master's thesis, Massachusetts Institute of Technology.

Jackson, P. L. (1988). Stock allocation in a two-echelon distribution system or "what to do until your ship comes in". *Management Science 34*(7), 880–895.

Jonsson, H. and E. A. Silver (1987). Analysis of a two-echelon inventory control system with complete redistribution. *Management Science 33*, 215–277.

Kapuscinski, R., R. Q. Zhang, P. Carbonneau, R. Moore, and B. Reeves (2004, May-June). Inventory decisions in dell's supply chain. *Interfaces 34*(3), 191–205.

Lawson, D. G. and E. L. Porteus (2000). Multistage inventory management with expediting. *Operations Research 48*(6), 878–893.

Oral, M., M. S. Salvador, A. Reisman, and B. V. Dean (1972). On the evaluation of shortage costs for inventory control of finished goods. *Management Science 18*(6), B344–B351.

Reyner, A. (2006). Multi-site inventory balancing in a global extended supply chain. Master's thesis, Massachusetts Institute of Technology.

Rote, G. (1992). The convergence rate of the sandwich algorithm for approximating convex functions. *Computing 48*, 337–361.

S. Axsater, J. M. and E. A. Silver (2002). Heuristic methods for centralized control of one-warehouse, n-retailer inventory systems. *Manufacturing and Service Operations Management 4*(1), 75–97.

Song, J.-S. and P. Zipkin (2003). *Supply Chain Operations: Assemble-To-Order Systems*, Volume XXX of *Handbooks in Operations Research and Management Science*, Chapter 11. North-Holland.

Uno, T. (2008, January). *Semiannual TFT LCD Materials and Components Report*. Austin, TX: DisplaySearch, L.L.C.

Online Appendix to:
Implementing Supply Routing Optimization in a Make-To-Order Manufacturing Network

John Foreman[1], Jérémie Gallien[2], Julie Alspaugh[3], Fernando Lopez[4], Rohit Bhatnagar[5], Chee Chong Teo[6], Charles Dubois[7]

July 25, 2008

A.1. Forecast Accuracy Study

Assuming a single location and part for now, this study can be described by defining d_t^* as the demand for that part actually observed in that location on day t, and $f_t^{t+\delta}$ as the corresponding demand forecast available at the beginning of day t for day $t + \delta$, so that $f_t^{t'} = d_t^*$ for $t' < t$. Because of the inventory balance equation stated in Section §4.1 of the paper, we were primarily interested in the cumulative forecast error $\bar{\varepsilon}_t^{t+\delta} \triangleq \bar{d}_{t,t+\delta}^* - \bar{f}_t^{t+\delta}$ with $\bar{d}_{t,t+\delta}^* \triangleq \sum_{k=t}^{t+\delta} d_k^*$ (sum of demand observed from day t to day $t + \delta$) and $\bar{f}_t^{t+\delta} \triangleq \sum_{k=t}^{t+\delta} f_t^k$ (forecast of the same quantity available at the beginning of day t). We first collected a large number of actual observations of the cumulative forecast error $\bar{\varepsilon}_t^{t+\delta}$ for all part types, all locations $\ell \in \{$Austin, Nashville, Reno, Winston-Salem$\}$ and all values of δ lower than the length T of the relevant planning horizon for supply routing decisions (about 4 weeks). Considering observations corresponding to disjoint sets of days (i.e. $\{\bar{\varepsilon}_{k\delta}^{k\delta+\delta}, k \in \mathbb{N}\}$) in order to avoid correlation biases due to time period overlaps, we then constructed and studied the associated empirical distributions of $\bar{\varepsilon}_t^{t+\delta}$. This study led to the following observations:

1. These empirical distributions were well fitted by the normal distribution, which is unsurprising in light of the central limit theorem given the definition of $\bar{\varepsilon}_t^{t+\delta}$. In order to establish this, we performed χ^2 and Kolmogorov-Smirnov goodness of fit tests for the hypotheses that the empirical data for the cumulative forecasting error $\bar{\varepsilon}_t^{t+\delta}$ had been

[1]   MIT Operations Research Center, Cambridge, MA 02142
[2]   MIT Sloan School of Management, Cambridge, MA 02142 (corresponding author, e-mail: jgallien@mit.edu)
[3]   Dell Computers, Inc., Round Rock, TX 78682
[4]   Dell Computers, Inc., Round Rock, TX 78682
[5]   Nanyang Technological University, Singapore
[6]   Nanyang Technological University, Singapore
[7]   Ecole des Mines de Paris, France

generated by normal, uniform and gamma distributions, respectively. The typical results of these tests were that the hypothesis of normality was accepted by both tests for all but the smallest values of the forecast horizon δ, whereas the hypotheses that the data had been generated by either the uniform or the gamma distribution were rejected. To illustrate this fit, Figure A.1 provides a plot of the cumulative distribution functions of the empirical data and the normal distribution for a particular part (a 17 inch flat panel) and location (Nashville), and $\delta = 5$ days;
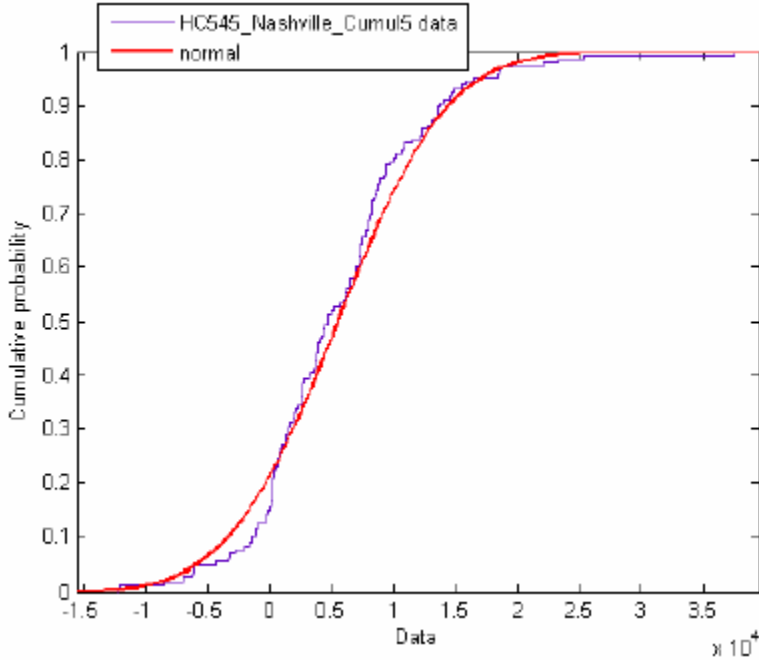
Figure A.1: Empirical Cumulative Distribution of the Cumulative Forecast Error $\bar{\varepsilon}_t^{t+\delta}$ for Part HC545 in Nashville and $\delta = 5$ days, and Normal Cumulative Distribution (Units Disguised).

2. The standard deviations σ_δ of these distributions for $\bar{\varepsilon}_t^{t+\delta}$ were well predicted by a coefficient of variation factor K_δ times the demand forecast $\bar{f}_t^{t+\delta}$, with K_δ only depending on the location and the number of days of demand predicted δ, and exhibiting a decreasing trend with δ.

3. The expected values $\mathbb{E}[\bar{\varepsilon}_t^{t+\delta}]$, representing the systematic forecasting bias, reflected two effects: (i) over the data collection period, the daily forecasts provided to the supply routing analysts were actually obtained by dividing a weekly forecast equally among all days of each week. Because Dell's demand within the week does exhibit a seasonality pattern, that construction method for the daily forecasts induced some bias; and (ii)

Dell's demand exhibited a general downwards trend for some components over a portion of the data collection period, which was slightly underestimated by the forecasting team in each one of its successive forecast revision steps. Because these biases were relatively small overall and accounted for by the two effects just described, we decided however to ignore them as part of our model.

These results suggested the structure and provided the input data for the stochastic model of cumulative demand stated in Section §4.1.1 of the paper.

A.2. Software Implementation

The software implementing the optimization model described in section §4 of the paper was developed using the environment and modeling language Ilog OPL and relied on the integer optimization engine Ilog CPLEX 9.1. The user interface was embedded in several Microsoft Excel spreadsheets, which are illustrated in Figures A.2 to A.4 below, and Figure 3 in the paper. Specifically, Figure A.2 shows the spreadsheet serving as a repository for control commands such as execution of optimization runs, choice of method used to generate the forecasting data, addition and removal of parts, visualization and enactment of decisions. Figure A.3 shows a portion of the spreadsheet developed to enter and modify the static input data, which includes costs, lead-times and forecast coefficients of variation. Figure A.4 shows a portion of the interface developed to visualize jointly some of dynamic input data (inventory, incoming supply, forecast), as well as the output data (transportation decisions, expected inventory and resulting expected shortages, denoted "Expected Backlog" in Figure A.4). Note that the interface shown in Figure A.4 was designed in order to represent the model output and its rationale in a format that would be familiar to the supply chain analysts, hence its similarity with the Balance Tool described in Section §3 of the paper. Finally, we refer the reader to Figure 3 in the paper for a screen copy of the interface developed to represent each individual supply routing decision generated by the optimization model runs.

Part Number	Parts per Pallet	Pallets per Truck:	Time Horizon	Cost per Unit Shortage	Approx. Iterations	Max. Model Runtime (seconds)	Model Transfer & Diversion Cost	Model Shortage Cost	Model Predicted Shortage	Model Total Cost	Do-Nothing Transfer & Diversion Cost	Do-Nothing Shortage Cost	Do-Nothing Predicted Shortage	Do-Nothing Total Cost
RH654	78	26	35	1	3	60	500	13109	32772.42	13608.969	0	18699.13	46747.83	18699.13
DJ034	60	26	35	1	3	60	500	4785.89	11964.72	5285.8878	0	10830.29	27075.73	10830.29
UJ327	84	26	35	1	3	60	37035	1284518	256903.6	1321552.8	0	1828484	365696.8	1828484
XA370	84	26	35	1	3	60	4636	83974	209935	88609.984	0	95607.36	239018.4	95607.36
JB007	60	26	35	1	3	60	61260	444575	222287.3	505834.59	0	620289.5	310144.8	620289.5
XX787	60	26	35	1	3	60	963	60307.4	150768.5	61270.409	0	67610.72	169026.8	67610.72
JG972	60	26	35	1	3	60	750	5082.09	12705.22	5832.0895	0	10040.33	25100.82	10040.33
SC074	45	26	35	1	3	60	73587	523066	104613.2	596652.92	0	1050000	209999.9	1050000
TG005	34	26	35	1	3	60	19136	205719	102859.5	224854.98	0	358062.6	179031.3	358062.6
LG207	34	26	35	1	3	60	500	24090.5	60226.23	24590.49	0	27104.53	67761.32	27104.53
PG208	32	26	35	1	3	60	30887	856528	171305.7	887415.28	0	986415.4	197283.1	986415.4
MA846	60	26	35	1	3	60	500	1824.71	4561.768	2324.7073	0	5545.41	13863.52	5795.41
EZ040	60	26	35	1	3	60	213	4861.73	12154.31	5074.7256	0	5253.594	13133.99	5253.594
HD101	54	26	35	1	3	60	1000	337671	844177.3	338670.91	0	349975.2	874938	349975.2
KK000	54	26	35	1	3	60	4068	8536.52	21341.31	12604.524	0	13365.3	33413.25	13365.3
BO008	54	26	35	1	3	60	0	5452.89	13632.23	5452.891	0	5452.891	13632.23	5452.891
JC011	40	26	35	1	3	60	250	7640.22	19100.55	7890.2215	0	9498.431	23746.08	9498.431
ST989	8	26	35	1	3	60	0	5705.2	14263	5705.2016	0	5705.202	14263	5705.202
GW210	17	26	35	1	3	60	0	10325.3	25813.31	10325.325	0	10325.33	25813.31	10325.33
FK333	6	26	35	1	3	60	0	13126.5	32816.22	13126.486	0	13126.49	32816.22	13126.49
OL008	48	26	35	1	3	60	250	4386.32	10965.79	4636.3166	0	5017.642	12544.1	5017.642
NE911	37	26	35	1	3	60	0	239287	598217.2	239286.9	0	239286.9	598217.2	239286.9
LP100	4	26	35	1	3	60	0	15027.7	37569.19	15027.677	0	15027.68	37569.19	15027.68
WW66	44	26	35	1	3	60	71	2358.96	5897.393	2429.9574	0	2430.084	6075.209	2430.084

Run Parts Macros

Add New Part

Remove Part(s)

Update Model Settings from template: ☑

Compensate for Short Term Forecast Overconsumption ☐

Reset Backlog Counter ☐

Show Part Info ☐

Show Model Info ☐

Show what-if ☐

Hide Decision Info ☑

Visualize Situation

Visualize Decisions B1

Enact Decisions

Deenact Decisions

Bar Decisions

Log Manual Decisions

Figure A.2: Screen Copy of the Control Interface (Disguised Data)

Transfer Settings:

Red Ball

Schedule:

Active Route? ('yes' = 1, 'no' = 0):

From:	To:	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Cost per truck ($):	Leadtime (days):
Austin	Nashville	0	0	0	0	1	0	0	150	1
	Reno	0	0	0	0	0	0	0	0	0
	Winston Salem	0	0	0	0	0	0	0	0	0
Nashville	Austin	0	1	0	0	1	0	0	150	1
	Reno	0	0	0	0	0	0	0	0	0
	Winston Salem	0	0	0	0	0	0	0	0	0
Reno	Austin	0	0	1	0	0	0	0	200	2
	Nashville	0	1	0	0	0	1	0	200	2
	Winston Salem	0	0	0	0	0	0	0	0	0
Winston Salem	Austin	0	0	0	0	0	0	0	0	0
	Nashville	1	0	0	0	0	0	0	100	1
	Reno	0	0	0	0	0	0	0	0	0

Max pallets per part: 8

Truck

To:	From:	Single Driver Truck Transportation Cost per Truck ($):	Single Driver Truck Leadtime (Days):	Team Driver Truck Transportation Cost per Truck ($):	Team Driver Truck Leadtime (Days):
Austin	Nashville	350	2	450	1
	Reno	350	3	450	2
	Winston Salem	350	3	450	1
Nashville	Austin	350	2	450	1
	Reno	350	3	450	2
	Winston Salem	350	1	450	1
Reno	Austin	350	3	450	2
	Nashville	350	3	450	2
	Winston Salem	350	4	450	3
Winston Salem	Austin	350	3	450	1
	Nashville	350	1	450	1
	Reno	350	5	450	3

Figure A.3: Screen Copy of the Interface for Entering and Modifying Cost and Lead-Time Data for Transfer Decisions (Disguised Data)

Part#: ER465

	0	1	2	3	4	5	6	7	8	9
	4/25/2008	4/28/2008	4/29/2008	4/30/2008	5/1/2008	5/2/2008	5/5/2008	5/6/2008	5/7/2008	5/8/2008
Austin:										
Demand Forecast	776	627	627	627	627	627	575	575	575	575
Arrivals and Departur	5278	0	0	0	0	0	0	0	0	0
Inventory	-8214	-2712	-3339	-3966	-4593	-5220	-5847	-6422	-6997	-7572
Expected Backlog	7403.56	3185.86	3874.282	4531.371	5174.813	5811.814	6393.052	6972.714	7551.423	8129.523
Nashville:										
Demand Forecast	623	565	565	565	565	565	538	538	538	538
Arrivals and Departur	0	4368	0	0	0	0	0	0	0	0
Inventory	373	-250	3553	2988	2423	1858	1293	755	217	-321
Expected Backlog	400.8285	883.4898	0	0	12.4433	163.4575	307.379	579.8195	886.9969	1237.53
Reno:										
Demand Forecast	12	12	12	12	12	12	7	7	7	7
Arrivals and Departur	0	910	0	0	0	0	0	0	0	0
Inventory	341	329	1227	1215	1203	1191	1179	1172	1165	1158
Expected Backlog	5.910321	3.085391	0	0	0	0	0	0	0	0
Winston Salem:										
Demand Forecast	563	401	401	401	401	401	443	443	443	443
Arrivals and Departur	0	0	0	0	0	0	0	0	0	0
Inventory	1236	673	272	-129	-530	-931	-1332	-1775	-2218	-2661
Expected Backlog	24.82773	143.4485	370.3103	680.0483	1017.557	1358.52	1784.642	2224.359	2665.29	3106.953

Figure A.4: Screen Copy of the Output Visualization Interface (Disguised Data)

5