# Bosch/Siemens Optimizes Inbound Logistics

Benjamin K. Peterson* Willem-Jan van Hoeve* Laurens G. Debo†
Sham Kekre*

Tepper School of Business Working Paper, January 2010

**Abstract**

Many companies ship parts and materials from a multitude of suppliers scattered over a large area. Because the number of supplier relationships may be large and complex, the procurement problem is often divided among parties that scarcely coordinate. As a result, opportunities exist to reduce total logistical costs by consolidating orders and sharing trucks. For the North-American division of the Bosch/Siemens Home Appliances Corporation, we study how to design and implement such coordination while maintaining flexibility with respect to market forces and changes in demand volumes. Our proposed solution combines the incoming freight from different suppliers in what we call 'flex-runs', to better utilize truck space. The result is an expected transportation cost reduction of up to 25% with an overall increase in robustness with respect to volume fluctuations.

**Keywords:** Logistics; Freight Consolidation; Integer Programming; Column Generation

## 1 Introduction

The Bosch/Siemens Home Appliances Corporation (abbreviated as B/S/H/) is the world's third largest leading manufacturer of high-end home appliances. In 1991, B/S/H/ launched its dishwasher business in North America, which initially imported dishwashers made in Germany. But in 1997, it started production of its European-designed dishwashers in New Bern, North Carolina. In the years thereafter, the site at New Bern was expanded to include the production of cooktops and laundry washers and dryers. To date, the New Bern site hosts three factories: dishwashing, cooking, and laundry that operate almost independently.

Currently, each manufacturing plant at B/S/H/ orders its supplies separately, even though a number of suppliers are shared among the plants. Most suppliers ship weekly or bi-weekly. Third-party logistics providers are hired externally, and most often 'full-truckloads' can be negotiated. For smaller shipments, 'less-than-truckloads' are usually applied. A full-truckload, or FTL, is a full, dedicated trailer with a typical weight capacity of 45,000 lb. In contrast, freight shipped by less-than-truckload, or LTL, is grouped with other shipments by the carrier and is subject to lower weight limits. The exact pricing of FTL and LTL shipments depends on many factors, the most important of which are the total distance traveled and total weight transported. In addition, the recently introduced fuel surcharges play an increasingly important role. However, FTL shipments

*Tepper School of Business, Carnegie Mellon University
†Booth School of Business, University of Chicago

| Type | Amount | Cost | Utilization |
|---|---|---|---|
| LTL | 25% | 37% | n/a |
| FTL | 72% | 47% | 38% |
| Other | 3% | 16% | n/a |

Table 1: Current shipping situation. For each shipment type the relative amount shipped and relative cost are given.

are usually more efficient (in terms of the unit cost per pound per mile) when the amount shipped is at least 10,000 lb (see, e.g., (5)). For smaller amounts, LTL shipments typically provide a lower unit cost.

Even though B/S/H/ uses FTL routes whenever possible, LTL routes are applied in many cases, especially when demand is greater than anticipated. Namely, when the amount to be shipped would cause an overload of the FTL truck, the remaining freight must be transported by an LTL shipment. The current shipping situation at B/S/H/ is depicted in Table 1. It shows the relative amount shipped and the relative cost spent for each of the shipment types LTL, FTL, and 'Other', where 'Other' refers to ad-hoc shipping methods such as parcel shipment. Note that even though 25% of all freight is shipped using LTL, it constitutes 37% of the total transportation costs. FTL, on the other hand, is used to ship 72% of all freight, and only constitutes 47% of the total cost.

Market fluctuations have a huge impact on sourcing volumes. The current operating plan is able to react quickly to changes in volume: each supplier is ordered from independently, so a change in volume (either positive or negative) does not affect the shipments of other suppliers. Moreover, the addition of LTL trucks allows rapid adaptation to volume increases. This flexibility comes with a cost, however. LTL shipments are much more expensive relatively than FTL shipments. Furthermore, the weekly or bi-weekly FTL routes are often under-utilized: the average FTL utilization is 38% (see Table 1). This means that, on average, FTL trucks are using only 38% of their capacity.

The observations above have led us to identify the following opportunities for cost improvement: *(i)* Since we have found that the three plants share a number of suppliers, is it possible to cut costs by consolidating shipments from one supplier to different plants? *(ii)* Is it possible to cut costs by consolidating shipments from different suppliers?

At first glance, it might seem impossible to put this potential into practice without giving in on flexibility with respect to volume fluctuations. That is, shipment consolidation might appear to be orthogonal to flexibility. Nevertheless, we show in this paper that consolidating shipments can significantly reduce transportation costs, while the resulting solution is actually *more robust* with respect to volume changes.

Our proposed solution involves the careful selection of suppliers into consolidated routes that we call *flex-runs*. Flex-runs are essentially flexible milk-runs (the common term for consolidated routes in logistics) designed to handle volume swings along the route. The basic idea of route consolidation is to reduce the relative contribution of the LTL shipments by converting them into (consolidated) FTL shipments. This is a well-known concept that is typically used by logistics providers to efficiently combine shipments of different customers. An illustrative example of such consolidation is given in Figure 1. In this small example, we need to ship freight from four

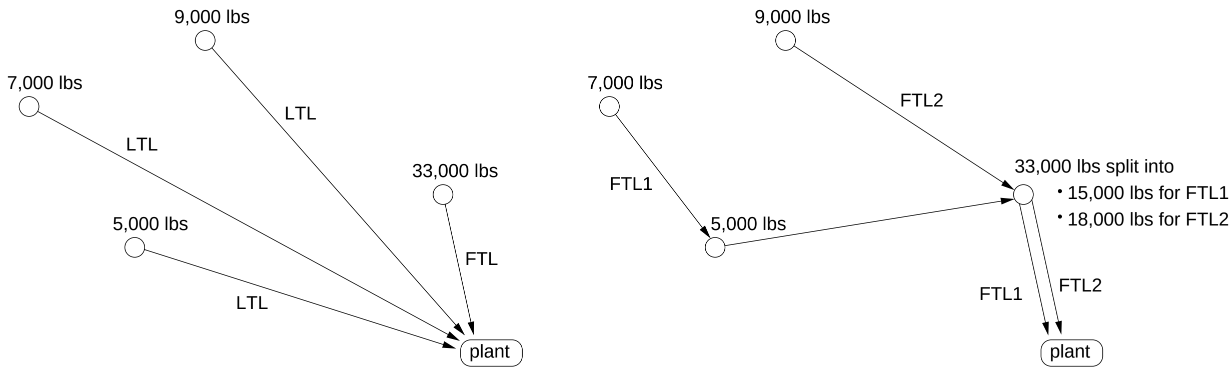


Figure 1: Shipment consolidation. The relatively expensive three LTL and one FTL routes are combined into two relatively cheaper FTL routes.

suppliers (the nodes) to one plant. The amount to be shipped is indicated above each supplier node. In the picture on the left, we apply the current strategy: if the amount to be shipped is less than 10,000 lb, we use an LTL, otherwise we use an FTL. In the picture on the right, we combine the three routes into two FTL flex-runs, FTL1 and FTL2. Given that the per-unit cost for FTL is lower than for LTL, we expect these two flex-runs to be less expensive than the scenario on the left. In fact, as our experimental results will show, we are able to design flex-runs that are expected to reduce the total inbound transportation cost for B/S/H/ by up to 25%.

An important aspect of shipment consolidation is to maintain flexibility with respect to variation in shipment volumes. To account for this, we design our flex-runs in such a way that the expected volume change for one supplier offsets the expected volume change for the other suppliers. In other words, by exploiting negative correlations between the volume changes of suppliers, we are able to guarantee a lower bound (typically 90-99%) on the probability that the FTL trucks on our flex-runs will remain within capacity. Therefore, our flex-runs turn out to be more robust than the current situation and there is less need for ad-hoc LTL shipments.

Designing and finding the *optimal* set of flex-runs is a computationally challenging problem (NP-hard to be precise). We propose to model and solve the problem as an integer linear program using column generation. With this technology, we are able to solve the B/S/H/ problem within 1% of optimality when we allow up to three suppliers in each flex-run. An initial set of our proposed flex-runs is currently being implemented at B/S/H/.

The remainder of the paper is structured as follows. We first provide a detailed problem description in Section 2. In Section 3, we present the outline of our solution, based on an integer linear programming model using column generation. In Section 4, we describe how we generate our columns representing the flex-runs. We provide detailed computational results and an analysis of our proposed solution in Section 5. In Section 6, we discuss some practical implications of our solution. Finally, we present the main conclusions in Section 7.

## 2 Detailed Problem Description

B/S/H/ has about 600 active suppliers for its three assembly plants in New Bern, NC. Among these, only 75 suppliers make up 80% of the total shipping volume. We chose to focus our study on these highly active suppliers. Each makes at least 40 shipments per year to one or more of the B/S/H/ plants, making them practical targets for one or more weekly flex-runs.

So far, logistics planning at B/S/H/ has focused more on the production schedule of the plants than on freight costs. Each supplier is assigned to a material planner, each of whom has several suppliers to manage. The material planner is responsible for monitoring the production schedule and ensuring that all parts and materials from their suppliers are present and ready when needed by the assembly line. Material planners take quantity discounts into account when placing orders but not shipping costs. Each order is assigned a truck based on its weight: LTL for orders under 10,000 lb, and FTL otherwise. For this reason, many orders between 10,000 and the maximum 45,000 lb are carried on partially-empty or even mostly-empty FTL trucks.

Trucks ordered from LTL carriers are priced according to published, zip code-based rate tables. A fixed discount percentage is also applied that is negotiated annually between B/S/H/ and each carrier. The amount of this discount depends on the quantity and frequency of LTL usage by B/S/H/ and on how the movements of these trucks fit into the carrier's overall network strategy. Sometimes the discount percentage varies by the shipment's state of origin.

Trucks ordered from FTL carriers are priced according to a fixed, negotiated rate per load. Each year, B/S/H/ distributes a list of locations to each of its carriers from which it plans to order many FTL shipments. A multiple-round bidding process takes place that identifies the lowest price a carrier will accept for each shipment made from each location. An exclusive contract is granted to each carrier for those locations on which it is the lowest bidder. B/S/H/ can then order FTL shipments at these fixed prices when needed. Because FTL shipments are not priced by weight, they are somewhat flexible to changes in quantities if they are under capacity.

Data was collected from records kept by a third-party logistics provider. The data identified the date, source, type, weight, and cost of each shipment. Unfortunately, the records did not indicate to which of the three plants each shipment was made. Thus, while we could analyze the benefits of consolidating any set of supplier shipments, we could not analyze the intermediate, and possibly more manageable, benefit of consolidating only within each plant.

## 3 Proposed Solution

Our proposed solution is to group suppliers into carefully-designed flex-runs. As stated before, these flex-runs are consolidated routes (or milk-runs) that are flexible with respect to fluctuations in shipping volumes. Our experimental results demonstrate that transportation costs can be decreased up to 25%, while changes in demand are better handled and trucks are better utilized.

The computational heart of our approach is an integer programming model based on column generation, where the columns of the model correspond to truck routes. This idea is originally due to (1), and has been successfully applied to several (large-scale) vehicle routing applications, see e.g., (3, 2, 6). In our approach, we specify for each route the suppliers visited as well as the amount picked up at each supplier and the expected cost to execute the route. In a pre-processing phase, we generate routes that include all possible combinations of up to three suppliers. Naturally, we could include more than three suppliers per route and use delayed column generation to manage the increased problem size. However, routes of such length are not desirable for B/S/H/. As will be discussed in the next section, we generated several million routes from which we selected about 600,000 useful candidates. To find a solution, the integer programming model selects those routes that satisfy all demand constraints, while minimizing the total cost.

More formally, let $R$ denote the set of all generated routes. Let $S$ denote the set of all suppliers. For each route $r \in R$, let $c_r$ denote the cost. For each route $r \in R$ and each supplier $s \in S$, let $p_{rs}$ denote the proportion of the freight picked up at supplier $s$ by route $r$. We introduce variables $x_r$ for all routes $r \in R$ representing the weekly frequency of route $r$, where $x_r$ is a non-negative integer. The problem of finding the set of routes that pick up all freight at minimum total cost can then be formulated as the following set-covering problem:

$$
\begin{aligned}
\min \quad & \sum_{r \in R} c_r x_r \\
\text{s.t.} \quad & \sum_{r \in R} p_{rs} x_r \geq 1 \qquad \forall s \in S \qquad (1) \\
& x_r \geq 0, \text{ integer} \qquad \forall r \in R.
\end{aligned}
$$

Note that constraints (1) are of the form 'greater than or equal', whereas commonly equality is used in set covering problems. The reason for this is that we allow some of our generated routes to be combined to (artificially) pick up slightly more than the required amount, provided that this would dominate other solutions in terms of cost and overload risk; see the *Flex-Run Design* section for more details.

This integer programming model contains one variable for each candidate route and one constraint for each supplier. It has one nonzero coefficient for each supplier visited by each route. In our experiments, this model found optimal or near-optimal solutions for problems of 75 suppliers and about 600,000 routes.

# 4 Flex-Run Design

To find an optimal trucking plan, we generate a set of flex-runs (routes) from which our integer program chooses an optimal (or near-optimal) set. The design of our flex-runs consists of three components: route generation, route pricing, and modeling flexibility.

## 4.1 Route Generation

The routing aspect of a flex-run consists of the supplier stops, and for each stop a pickup proportion of the weekly shipment volume. Therefore, we first divide each supplier's weekly shipment into a set of equally-weighted packets, each to be assigned to one route. This allows shipments to be split up in a variety of ways, but with finitely many combinations. We chose to divide shipments into packets of around 2,000 lb, though a higher or lower precision may be used in general. For suppliers shipping more than 10,000 lb per week, we divided the delivery into $\lfloor weight/2000 \rfloor$ equal packets. For suppliers shipping less than 10,000 lb per week, we decided not to allow delivery splitting as it would impose an excessive inconvenience on the supplier. In other words, for each supplier $s \in S$, the number of packets is given by:

$$\text{\# packets}(s) = \begin{cases} 1 & \text{if weight}(s) \ < \ 10{,}000 \text{ lb/week} \\ \lfloor \text{weight}(s)/2000 \rfloor & \text{if weight}(s) \ \geq 10{,}000 \text{ lb/week} \end{cases}$$

where weight($s$) denotes the total weight of the weekly freight to be picked up at supplier $s$.

We first generated one-stop LTL routes for all suppliers. If the shipping requirement for a route is less than 10,000 lb, we added a single LTL route for the entire shipment. Otherwise, we added an LTL route for each integer number of packets not exceeding 10,000 lb, which is a reasonable LTL upper bound. An example is shown in the table below, for two suppliers A and B with weekly shipments of 8,000 lb and 11,000 lb, respectively.

| Supplier | Weekly Shipment | Capacity of LTL Routes Added |
|---|---|---|
| A | 8,000 lb | 8,000 lb |
| B | 11,000 lb | (1/5)×11,000 lb = 2,200 lb<br>(2/5)×11,000 lb = 4,400 lb<br>(3/5)×11,000 lb = 6,700 lb<br>(4/5)×11,000 lb = 8,800 lb |

We next generated FTL routes for all combinations of one, two and three suppliers; it is undesirable for B/S/H/ to have more than three stops on a route. The number of all such combinations is $\sum_{k=1}^{3} \binom{75}{k} = 70{,}375$. Note that each supplier combination can be used with various shipment configurations. For each combination, we checked all possible orders in which the suppliers could be visited before driving to B/S/H/, and selected the sequence with the lowest total mileage. We used *Microsoft MapPoint NA 2006* with the *MileCharter* add-in to compute a distance matrix for this calculation.

Finally, we need several copies of each route to cover all meaningful combinations of pickup proportions. We first generated all possible combinations for each route, which totaled $\prod_s$ # packets($s$) combinations, where # packets($s$) again denotes the number of packets into which the shipment of supplier $s$ is divided.

Note that not all generated routes will be included in the model. We describe below how we select the flex-runs that are flexible enough to handle volume changes without overloading. In addition to excluding potential route combinations that exceed the allowed overload risk, we also exclude combinations that are unnecessarily small. For example, an FTL route shipping one tenth of a 20,000 lb shipment would be a waste of capacity. However, since lower weight assignments are less likely to overload, we included them as long as their expected excess weight, computed as $E[max(0, weight - 45000)]$, was at least 1 lb less than that of another included route.

### 4.2 Route Pricing

For each of the generated LTL routes, we looked up a rate quote on the web site of the LTL carrier used by B/S/H/ in the region of the supplier. For the FTL routes, we created a model that captures the cost structure as follows.

For a list of previously negotiated FTL routes used by B/S/H/, we plot the cost against the length of those routes in Figure 2. From this figure, it is apparent that a close relationship exists between

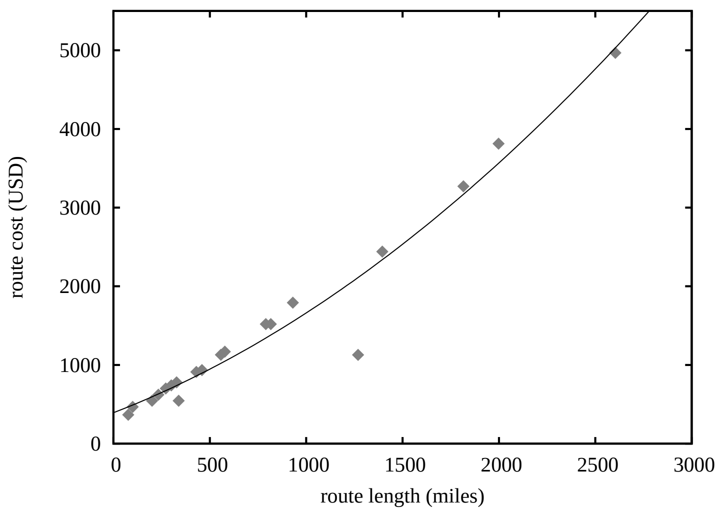


Figure 2: Cost (in USD) of previously negotiated FTL routes with respect to their length (in miles). The cost numbers are rescaled for confidentiality.

the mileage of a route and the minimum price an FTL carrier would accept to haul it. We fit a quadratic curve to this data and used it to estimate prices for our proposed FTL routes. We also added a stopover charge found on the carrier's website for those routes visiting more than one supplier.

### 4.3 Modeling Flexibility

As mentioned above, we only wish to include those routes that are flexible enough to handle volume swings. To model the probability that a generated route would overload, we define a random variable $X_s$ for each supplier $s \in S$ representing its weekly shipment. We define $\mu_s$ as the average of $X_s$ and $\sigma_{st}$ as the covariance between the weekly shipment two suppliers $s, t \in S$. We denote the covariance matrix by $\Sigma$. For a multiple-stop route $r$, we compute the mean and variance of the total weekly shipment as follows:

$$\mu_r = \sum_{s \in S} p_{rs} \mu_s, \quad \sigma_r^2 = \sum_{s \in S} \sum_{t \in S} p_{rs}^T \Sigma p_{rt}$$

From this we generate for each route a Gamma distribution $Gamma(k, \theta)$ with the same mean $\mu_r$ and variance $\sigma_r^2$ to model the probability distribution of the load of the route. We compute the parameters $k$ and $\theta$ as follows:

$$k = \mu_r^2 / \sigma_r^2, \quad \theta = \sigma_r^2 / \mu_r$$

We then select a service level $\alpha$, the minimum probability that the route will be within capacity. If the $\alpha$%-quantile of the $Gamma(k, \theta)$ is over the truck capacity, the requirement is not met and the route is excluded. That is, if more that $1 - \alpha$% of the distribution is above the 45,000 lb, that weight combination is considered a high overload risk and excluded from consideration.

While a normal distribution may seem more applicable for the load distribution, we find it significantly underestimates the frequency of capacity overload, since much of the outlying area of the normal distribution lies in negative values when variances are as high as those we encountered.

**Example 1** *Consider the following illustrative weekly pick-up amounts for two suppliers A and B:*

| *week* | *1* | *2* | *3* | *4* | *5* | *mean* |
|---|---|---|---|---|---|---|
| *supplier A* | *20,000* | *10,000* | *10,000* | *30,000* | *30,000* | *20,000* |
| *supplier B* | *15,000* | *10,000* | *5,000* | *20,000* | *25,000* | *15,000* |

*Suppose we design a weekly route r that picks up 50% of the weight of supplier A, and 100% of the weight of supplier B. That is,* $p_{rA} = 0.5$ *and* $p_{rB} = 1.0$*. The total expected mean weekly load of route r, based on the data, is 25,000 lb, with a standard deviation of approximately 11,401 lb. The Gamma distribution that models the load of this route has parameters* $k = 4.81$ *and* $\theta = 5200$*, and is depicted in the figure below.*

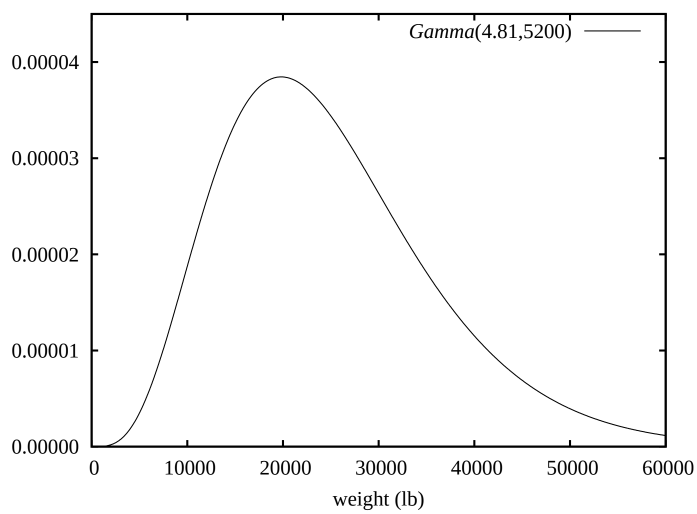


*The cumulative distribution of* $\text{Gamma}(4.81, 5200)$ *at 45,000 lb is 0.9421, which means that this route would be accepted if we impose an non-overload guarantee of 95%.*

The resulting set of flex-runs, containing all LTLs and FTLs of one to three suppliers using all sensible weight combinations, number around 600,000. Consequently, our IP model consists of 600,000 variables and 75 constraints (for 75 suppliers).

## 5 Computational Results

In this section we provide a detailed analysis of the solution we proposed to B/S/H/. In addition to a single solution, B/S/H/ management wanted us to test the robustness of our model parameters. These parameters include the maximum number of stops on each flex-run, the minimum service level $\alpha$ that we require for each flex-run, and the shape of the weight distribution assumed for each flex-run load. We performed several experiments to test different values of the parameters and confirm the accuracy of the model. Throughout this section, cost figures have been rescaled by a fixed multiplier for confidentiality.

For our experiments we have used the integer programming solver CPLEX 11 (4) on a 2.4Ghz machine with 8GB of RAM running Windows Vista 64 Pro. For the 75-supplier instances, we report the solutions after a duality gap of 1% was reached, i.e., the solution found is provably within 1% of the optimal solution. Finally, we note that for the larger problem instances on around 600,000 variables, most of the available RAM was used. All reported solutions were found within 24 hours total.

## 5.1 Number of Stops per Route

Our truck routing model contains 75 supplier locations that must be visited by one or more trucks each week. As described in our route generation procedure, we considered only those routes visiting at most three suppliers. This both limits the size of the model for computational tractability and limits the amount of coordination required for B/S/H/ to implement the solution. To study the effects of this maximum-stops parameter on an optimal routing cost, we tested the inclusion of routes with more stops on a smaller example problem containing 15 suppliers. For a problem of this size, the number of these routes is small enough to enumerate and solve to optimality. We also used a distribution of weight and locations similar to our B/S/H/ data.

In Figure 3 we display the results of this experiment by means of a table and a plot. In this figure, '# Suppliers' indicates the number of suppliers in the problem, 'Max Stops per Route' indicates the maximum number of stops (i.e., suppliers) on each route, 'Min Reliability' indicates the reliability of each route to be within the truck capacity (in this experiment it is set to 95%), and 'Optimal Cost' indicates the optimal weekly shipping cost for this problem.

As the results show, cost savings improve as more stops are added to the feasible routes, but the marginal value of the extra stops diminishes quickly. To observe this effect on our full 75-supplier problem, we tested our model with a maximum of one, two and three stops per

| # Suppliers | Max Stops per Route | Min Reliability | Optimal Cost |
|---|---|---|---|
| 15 | 1 | 95% | $30,553 |
| 15 | 2 | 95% | $19,292 |
| 15 | 3 | 95% | $14,427 |
| 15 | 4 | 95% | $12,916 |
| 15 | 5 | 95% | $12,714 |
| 15 | 6 | 95% | $12,714 |
| 15 | 7 | 95% | $12,714 |
| 15 | 8 | 95% | $12,714 |

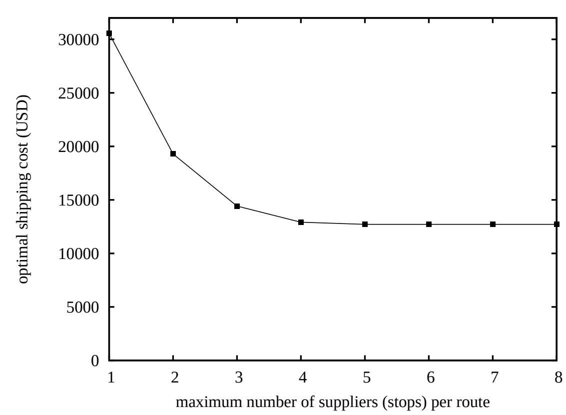


Figure 3: The effect of increasing the number of stops per route on the optimal shipping cost, for a 15-supplier benchmark problem.

| Schedule | # Suppliers | Max Stops per Route | Min Reliability | # Feasible Routes | Est. Weekly Fixed Cost | Savings |
|---|---|---|---|---|---|---|
| Original | 75 | 1 | - | 75 | $134,063 | - |
| Optimization 1 | 75 | 1 | 95% | 293 | $123,331 | 8.0% |
| Optimization 2 | 75 | 2 | 95% | 13,052 | $105,337 | 21.4% |
| Optimization 3 | 75 | 3 | 95% | 592,975 | **$100,000** | 25.4% |

Table 2: The effect of increasing the number of stops per route on the shipping cost and expected savings, for the 75-supplier problem of B/S/H/.

| Schedule | # Suppliers | Max Stops per Route | Min Reliability | # Feasible Routes | Est. Weekly Fixed Cost | Savings |
|---|---|---|---|---|---|---|
| Original | 75 | 1 | - | 75 | $134,063 | - |
| Optimization 3a | 75 | 3 | 99% | 291,110 | $121,899 | 9.1% |
| Optimization 3b | 75 | 3 | 95% | 592,975 | **$100,000** | 25.4% |
| Optimization 3c | 75 | 3 | 90% | 851,446 | $87,518 | 34.7% |

Table 3: The effect of increasing the reliability for each route to remain within capacity on the shipping cost and expected savings, for the 75-supplier problem of B/S/H/.

route. The results are depicted in Table 2. In this table, we compare four different transportation schedules: the original schedule of B/S/H/, and our three optimized schedules based on routes including a maximum of one, two, and three stops per route, respectively. Here '# Feasible Routes' indicates the number of feasible routes that were generated, 'Est. Weekly Fixed Cost' indicates the estimated weekly fixed transportation cost of the schedule, and 'Savings' indicates the savings of the optimized schedules with respect to the original schedule.

Note that the optimized schedule with 1 stop per route has a lower cost than the original schedule. This is the result of multiple LTL routes being consolidated into single FTL routes for suppliers shipping to multiple plants at B/S/H/.

Because the number of feasible routes increases exponentially with respect to the maximum number of stops, we found it impossible to add routes with more than three stops using out model. The problem files became to large to load into memory. However, we felt that the above data was sufficient to conclude that most of the available savings is captured by our one- to three-stop routes.

### 5.2 Reliability of Remaining within Truck Capacity

The significant cost reduction of our solution is attributable to an increased use of FTL routes, since FTL costs less per pound of capacity than LTL. But to exploit this cost reduction, one must consistently have enough freight to utilize the large, fixed capacity of the FTL. This can be difficult to achieve under conditions of high demand variability: while assigning several suppliers to an FTL route may reduce per-unit costs, it may also result in frequent overloading of the truck capacity. In this case, an additional LTL route may be necessary, negating the savings of the FTL route and increasing managerial costs.

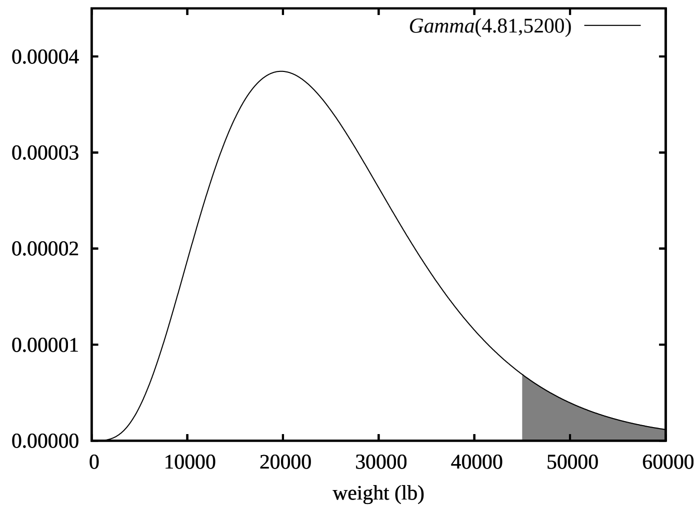


Figure 4: The probability of overloading the route of Example 1, indicated by the shaded region.

To balance the two goals of minimizing cost and maintaining reliability, we consider only those routes in our model which achieve some minimum threshold of reliability. That is, the proportion of supplier freight assigned to the route must be within its capacity during some fraction of the times it operates. In Table 3 we report the results of an experiment comparing the savings at various thresholds of reliability, namely 99%, 95% and 90%, referred to as Optimization 3a, 3b, and 3c, respectively. In these experiments, the maximum number of stops is always three: The experiment 'Optimization 3' in Table 2 is equal to 'Optimization 3b' in Table 3, both costs are indicated in bold. Based on these results, we decided the solution with 95% reliability was the best trade-off of route costs and managerial costs.

### 5.3 Estimating the Capacity Overload

In order to further quantify the cost of unreliability in our solution, we wished to measure the expected cost incurred during the 5% or less of cases in which FTL route capacity is overloaded. For this we used the gamma distribution fit to each route's weekly load in our route-generation procedure. To find the expected overload weight of each FTL route, we integrated the excess function $max(0, weight - 45000)$ over the distribution. We then multiplied this weight by the approximate cost per pound of an additional LTL route to find the expected overload cost of the FTL route. Recalling Example 1, a route with mean load 25,000 lb and standard deviation 12,000 lb has an expected overload weight of 455 lb; this corresponds to the shaded area in Figure 4.

Table 4 shows the expected overload cost associated with our proposed solution. We estimate this cost for all feasible routes that were included in the model ('All Feasible Routes'), and for the subset of routes included in the solution to two optimizations. The first ('Min Fixed Cost Routes') corresponds to the optimal solution that we have considered thus far, using a fixed cost for each route. The second ('Min Total Cost Routes') adds to the fixed cost the overload cost for each route. Therefore, the objective is slightly higher, and the solution changes slightly.

For each of these three route sets, we report the approximate weekly fixed cost of the routes,

| | All Feasible Routes | | Min Fixed Cost Routes | | Min Total Cost Routes |
|---|---|---|---|---|---|
| Approximate Weekly Fixed Cost | $2,874,994,855 | | **$100,000** | | $100,144 |
| Expected Weekly Overload Cost | + $19,177,002 | + | $4,091 | + | $3,489 |
| Total Weekly Cost | = $2,894,171,856 | = | $104,090 | = | $103,634 |
| Actual Overload Cost | $11,063,102 | | $1,821 | | $1,561 |

Table 4: Estimating the overload cost for the complete set of generated routes, and for the routes that are included in the optimal solution. These results are based on the 2007 data (12 months) from B/S/H/.

| | All Feasible Routes | | Min Fixed Cost Routes | | Min Total Cost Routes |
|---|---|---|---|---|---|
| Approximate Weekly Fixed Cost | $2,850,174,001 | | $98,059 | | $98,431 |
| Expected Weekly Overload Cost | + $17,916,572 | + | $3,873 | + | $3,391 |
| Total Weekly Cost | = $2,868,090,573 | = | $101,932 | = | $101,822 |
| 8-Month Calibrated Overload Cost | $21,776,023 | | $1,701 | | $1,583 |
| 4-Month Uncalibrated Overload Cost | $8,233,524 | | $5,696 | | $5,174 |

Table 5: Estimating the accuracy of the overload cost model based on the Gamma distribution, by calibrating it on eight months of the 2007 B/S/H/ data, and testing it against the remaining four months.

the expected weekly overload cost, and the total weekly cost of the routes. For all these numbers, the load of the routes are modeled by a Gamma distribution. We compare this (in the last row of Table 4) to the actual overload costs, based on the B/S/H/ routes of 2007. Note that the latter is computed from the data itself, not the distribution fit to the data.

A first observation is that the expected overload cost for the routes included in the optimal solution is relatively high, around 4% of the weekly cost, compared to that of all feasible routes, which is around 0.67%. From Table 4 we further derive that the overload costs are somewhat overestimated, but at the same time constitute only a marginal amount of the overall transportation costs. Further, adding the overload costs to the objective did only slightly change the optimal solution, and we therefore omitted the overload costs in our basic model.

Finally, we wish to evaluate how well our model using the Gamma distribution fits the overload cost estimations. To this end, we decided to rebuild our model using only eight months of our available 2007 historical data, and test the resulting solution against the data of the remaining months. That is, estimations of load distribution and reliability used for designing routes were based on only eight of our 12 months of data. We used months 1,2,4,5,7,8,10, and 11 to eliminate seasonality effects. After optimizing this new model, we compared estimated overload costs during the eight calibration months and the remaining four months. The results are reported in Table 5.

From this experiment we see that overload costs will inevitably be minimized during the

| Original Routes | | | | | |
|---|---|---|---|---|---|
| Route Type | #/week | Avg. Mileage | $/week | ton·mi/week | $/ton·mi |
| 1-Stop LTL | 126 | 623 mi | $ 62,853 | 122,706 | $ 0.49 |
| 1-Stop FTL | 39 | 436 mi | $ 71,210 | 258,966 | $ 0.29 |
| **Optimized Routes, 95% Reliability Level** | | | | | |
| Route Type | #/week | Avg. Mileage | $/week | ton·mi/week | $/ton·mi |
| 1-Stop LTL | 19 | 880 mi | $ 11,873 | 30,813 | $ 0.39 |
| 1-Stop FTL | 9 | 128 mi | $ 7,871 | 23,612 | $ 0.33 |
| 2-Stop FTL | 9 | 305 mi | $ 11,807 | 49,309 | $ 0.24 |
| 3-Stop FTL | 23 | 897 mi | $ 68,449 | 341,852 | $ 0.20 |

Table 6: The structure of the original and optimized routes.

period over which the routes are calibrated, as reflected in the row '8-Month Calibrated Overload Cost'. Also, the optimization tends to select routes for which the overload cost is underestimated. However, the overload cost of the four uncalibrated months ('4-Month Uncalibrated Overload Cost') is fairly accurate in the sense that the weekly costs are comparable to those of Table 4, which are based on the model that was calibrated on the entire 2007 data set. Therefore, we consider our estimate to be sufficiently accurate to conclude that overload costs negate only a small fraction of the total savings achieved.

### 5.4 Analyzing the Proposed Solution

As indicated above, we chose to submit to B/S/H/ the optimal solution using the 95% reliability level with respect to the overload risk. We next provide a detailed analysis of this solution and compare it to the original shipping situation at B/S/H/.

We first discuss the structure of the original routes and the optimized routes in terms of the amount of LTL and FTL shipments, presented in Table 6. To allow a meaningful comparison, the reported costs of the original routes are based on the model we presented in the *Flex-Run Design* section. In Table 6, we report for each route type the average mileage, the average cost per week, the average ton-mileage per week, and the average cost per ton-mile. A first observation is that the number of routes dramatically decreased from 165 to 60 routes in total, most of which are 3-stop FTLs. Interestingly, the average mileage for the resulting optimized LTL routes is higher than the original LTL routes. Indeed, an analysis of our solution showed that the best candidates for optimized LTL routes are those that are geographically isolated from the others, while other factors are low shipping weight and high variability.

To identify which routes (and suppliers) attribute most to the total cost savings of about 25%, we present in Table 7 the average unit costs for the different route conversion: from LTL to Optimized LTL, from LTL to Optimized FTL, and from FTL to Optimized FTL. The most important contribution stems from the conversion of LTL routes into Optimized FTL routes (from $0.44 to $0.21 per ton-mile), while converting LTLs and FTLs into optimized LTLs and FTLs contribute relatively less to the savings.

| | Original Routes | Optimized Routes 95% reliability | |
|---|---|---|---|
| Type Conversion | \$/ton·mi | ton·mi/week | \$/ton·mi |
| LTL → LTL Optimized | \$ 0.44 | 30,813 | \$ 0.39 |
| LTL → FTL Optimized | \$ 0.44 | 91,893 | \$ 0.21 |
| FTL → FTL Optimized | \$ 0.29 | 258,966 | \$ 0.21 |

Table 7: Cost savings appear when converting the original routes into optimized routes.

| FTL | Original | 95% reliability | 90% reliability |
|---|---|---|---|
| Average Utilization | 38.1% | 48.5% | 55.5% |
| 10%-Peak Utilization | - | 81.9% | 96.5% |
| 5%-Peak Utilization | - | 96.6% | 115.6% |

Table 8: Comparing the average and peak FTL truck utilization of the optimized solution with the original shipping situation.

Finally, we analyze the truck utilization of our optimized solution, as compared to the original shipping situation at B/S/H/. As was reported in Table 1, the current shipping situation uses only 38% of the FTL truck capacity, on average. We report in Table 8 the average truck utilization of our proposed solution, as well as the 'peak utilization'. The latter is defined as the average of the highest loads over all trucks. For example, the 10% peak utilization collects for all trucks the 10% highest load volumes, and the average is reported as the 10%-peak volume (in percentage of the capacity). Similarly for the 5%-peak volume. The results in Table 8 clearly show an improvement in average truck utilization with respect to the original situation (from 38.1% to 48.5%). We note that the average truck utilization is deliberately kept at a relatively low level by our optimization model, to allow a reliability buffer for the flex-runs. That this is indeed needed is shown by the peak utilization of the trucks. For example, for the 95% reliability level, the 10%-peak utilization is 81.9% of the truck capacity. This means that many flex-runs will in fact be almost completely filled during roughly 10% of the schedule execution. We note that for the 90% reliability level, the higher chance of overloading is witnessed by the 5%-peak utilization of 115.6% of the truck capacity (i.e., an overload of 15.6%).

## 6 Implementation in Practice

From the perspective of B/S/H/, the initial goal of this project was to gain insight in the potential cost savings for their inbound freight logistics operations. If these cost savings are high, a company-wide implementation could be worthwhile. B/S/H/ further envisioned that the insight gained by this project (particularly the cost savings) could be useful during the contract negotiations with their logistics providers.

We provided B/S/H/ with the findings reported in this work. In particular, we presented them the optimal solution with a 95% reliability level for the overload risk. Our reported expected cost savings of up to 25% led B/S/H/ to initiate an implementation of our solution, first on a small scale.

Based on the solution we reported, B/S/H/ selected a subset of most profitable felx-runs and locations, as candidates to implement in a first phase. For this subset, B/S/H/ provided more detailed information and restrictions that needed to be taken into account. For example, for some of the routes stacking and volume restrictions were added, while for other routes it was important to respect a certain order to visit the suppliers. Our model can easily handle all these additional restrictions during the route generation process: we only include those routes that are feasible with respect to all restrictions. From the solution that was obtained by re-optimizing this model, B/S/H/ selected five flex-runs for actual implementation in practice.

It should be noted that our proposed solution can have major implications throughout the organization at B/S/H/. Most importantly, it requires communication between different logistics planners, sometimes even between different plants. It could therefore be useful in practice to limit such communication by introducing certain planning protocols for shared trucks. Alternatively, one could try to make the communication more transparent through a web-based plant-wide planning tool. In any case, this is an important issue that must not be ignored during the implementation.

Another important issue is the reward and incentive structure at manufacturing companies such as B/S/H/. Currently at B/S/H/, logistics planners are mainly evaluated based on inventory levels (the lower the better), while transportation costs are mostly ignored. Since our proposed solution focuses on optimizing the transportation costs, there is no immediate incentive for logistics planners to put extra effort in the successful implementation of our solution. It would therefore be advisable if the evaluation of the logistics planners would take transportation costs into account as well.

## 7 Summary and Conclusions

We have described an approach to optimize the inbound freight logistics for Bosch/Siemens in North America, a leading manufacturer of home appliances. The computational model of our approach is based on combining individual supplier's shipments into 'flex-runs' (milk-runs that are flexible with respect to shipping volume fluctuations) such that total logistic costs are reduced while maintaining, if not increasing, the reliability of the shipping schedule. To this end, we exploit correlations between the variations in shipping volume for different supplier locations, to ensure that each potential route does not overload with respect to a given level of reliability. The design of our routes further allows the shipping volumes to be combined in various proportions using any desired level of precision. Our approach applies integer linear programming and column generation, where the columns of the model correspond to feasible (flex-run) routes.

We have demonstrated the benefits of our approach on real-life data from Bosch/Siemens. For this data set, we pool volumes from up to three suppliers onto a single route, which allows us to generate the required feasible routes in a pre-processing phase. When more than three stops are demanded, the number of feasible routes quickly grows too large, and delayed column generation can be employed instead.

Our most important finding is that with our optimized solution, expected cost savings of up to 25% can be achieved with respect to the current shipping situation, while at the same time, the robustness of the schedule with respect to changes in shipping volume increases. An initial implementation of our proposed solution is currently initiated at Bosch/Siemens. Our approach is not restricted to the situation at Bosch/Siemens however, but in fact is broadly applicable.

## References


[1] M.L. Balinski and R.E. Quandt. On an integer program for a delivery problem. *Operations Research*, 12(2):300–304, 1964.

[2] C. Barnhart, E.L. Johnson, G.L Nemhauser, M.W.P. Savelsbergh, and P.H. Vance. Branch-and-Price: Column Generation for Solving Huge Integer Programs. *Operations Research*, 46(3):316–329, 1998.

[3] M. Desrochers, J. Desrosiers, and M. Solomon. A New Optimization Algorithm for the Vehicle Routing Problem with Time Windows. *Operations Research*, 40(2):342–354, 1992.

[4] ILOG S.A. ILOG CPLEX 11.0 Reference Manual, 2007.

[5] P.R. Murphy, Jr. and D.F. Wood. *Contemporary Logistics*. Pearson Prentice Hall, 2008. 9th ed.

[6] A. Pajunas, E.J. Matto, M. Trick, and L.F. Zuluaga. Optimizing Highway Transportation at the United States Postal Service. *Interfaces*, 37(6):515–525, 2007.