European Journal of Operational Research 207 (2010) 456–464

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Integrated optimization of customer and supplier logistics at Robert Bosch LLC

Hakan Yildiz [a,*], R. Ravi [b], Wayne Fairey [c]

[a] *Eli Broad College of Business, Michigan State University, East Lansing, MI 48824, USA*
[b] *Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA 15213, USA*
[c] *Customer Service, Logistics and Planning (ChP/CLP), Robert Bosch LLC, Charleston, SC 29418, USA*

ARTICLE INFO

Article history:
Received 30 April 2009
Accepted 31 March 2010
Available online 3 April 2010

Keywords:
Supply chain management
Logistics
Routing
Network flows
Application

ABSTRACT

Large automotive supply chains typically involve manufacturers pulling materials from their suppliers along the chain, usually by using round-trip truckload routes. The return trips on these routes are used to return empty containers back to the suppliers. The mismatch between the amount of materials and empty containers results in underutilization of the return trips. A supplier can utilize this unused capacity by identifying a subset of promising customer routes that can be combined with its existing supplier routes to save overall costs of the system. Such an integration also leads to other supply chain coordination benefits such as the potential of using crossdocks, more frequent milkruns and ensuing reductions in inventories.

We undertake such an integrated study of the inbound logistics from suppliers and the outbound logistics to customers at Robert Bosch LLC, a leading automotive parts manufacturer. We identify the opportunity for significant cost savings by using a mixed-integer programming model that matches opposite flows from and to the customers and suppliers. We consider the problem from a supply chain coordination perspective, where Bosch makes all the transportation arrangements for its customers and suppliers based on the centralized optimum solution, and outline its additional benefits.

1. Introduction

In the automotive industry, traditionally each company in the supply chain plans and pays for the shipments from its suppliers. For instance, automotive manufacturers pick up and pay for the cost of shipments from the first tier suppliers; first tier suppliers do the same for the supplies they get from second tier suppliers and so on and so forth. The nature of these supplies is interesting in the sense that while there is material flow from higher tier to lower tier suppliers, there is also significant flow of empty containers (return dunnage) in the other direction as shown in Fig. 1.

There are two main types of transportation offered by carriers: truck-load (TL) and less than truck-load (LTL). TL carriers are dedicated trucks carrying the load of a shipper directly from an origin to a destination potentially stopping along the way to fill the truck. The final destination can be different than the origin, which makes the route a one-way route; or it can be the same, which then makes the route a round-trip route (also referred to as a "*milkrun*"). In either case, there may be intermediate stops between the origin and destination. LTL carriers, on the other hand, provide a shared service on the same truck and use an airline-type hub-and-spoke system with shipments. An LTL carrier usually assumes the responsibility for routing each shipment from the origin to the destination.

The cost structures of these two types of carriers are very different and it may be more advantageous to use one or the other depending on the destination(s), the amount, and the density of the shipment. Since the unit cost (i.e. cost per pound per mile) of round-trip TL carriers is significantly less than the unit cost of one-way TL carriers or LTL carriers, companies try to utilize round-trip TL trucks that carry supplies from their suppliers and carry empty containers in the rest of the round trip.

Notably, the amount of space occupied by empty containers is significantly less than the space they occupy when they are full. This is due to the fact that some of the containers are designed to collapse when they are empty and thus occupy less volume. Moreover, not all shipments are made with returnable containers but rather with materials such as cardboard, which are disposed after they are used. Because of this, the ratio of the volume of returned empty containers to the volume of products shipped ranges from 0 to 1. Furthermore, the weight of the empty containers is also significantly less than their weight when they are full. Interestingly, the relatively low weight of the containers also allows higher stacking of containers, which may not be possible when the containers are full due to total weight restrictions on trucks. Thus, the amount of reverse flow of empty containers does not match the amount of product flow. This leaves plenty of unused

* Corresponding author. Tel.: +1 517 4326439; fax: +1 517 4321112.
E-mail address: yildiz@msu.edu (H. Yildiz).

0377-2217/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2010.03.044

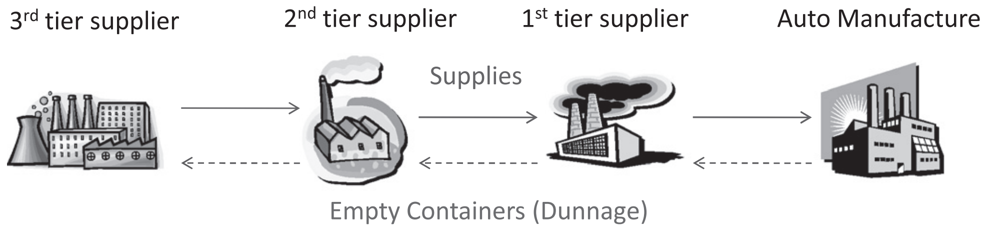

Fig. 1. Logistics in the auto industry involve transfer of materials and dunnage in opposite directions.

capacity on the return segment of a round-trip route. With this observation, we realize the following opportunity for a first tier supplier.

2. Matching opposite flows: A unique opportunity

A company that has customers and suppliers located in the same region, which is far away from the manufacturing plant, can combine the flow of inbound supplies and empty supplier containers with the flow of outbound finished products and empty customer containers. This is illustrated in the left half of Fig. 2, where customers, represented by circles, and suppliers, represented by squares, are located in the Midwest and Canada and they are far away from the manufacturing plant located in Charleston, SC (ChP). The product flows are represented by solid lines whereas the empty container flows are represented with dashed lines, which are drawn thinner than the solid lines to represent the imbalance mentioned earlier.

This opportunity is not possible for automotive manufacturers since trailers that can carry automobiles are very different to the trailers that carry their supplies stored in containers. There is also less potential for a second tier supplier who may not have as much volume in both directions. But a first tier supplier can utilize this opportunity by taking over the delivery of its products to its automotive manufacturer customers.

2.1. Crossdocking

A crossdock is a facility where materials are unloaded from incoming trucks and loaded into outbound trucks with little or no storage in between. Such a facility located close to the customers and suppliers can also be used as a consolidation point to be able to match the two opposite flows better, as depicted in the right half of Fig. 2. Locational decisions, such as locating a crossdock, are usually made on a strategic level as they have long term effects. In the case of using a crossdock facility, one alternative is owning the facility by making an investment. Another alternative is renting the facility, which may require less capital investment than the first one, but still require a commitment for a period of time. A third alternative is using an existing facility operated by a 3rd Party Logistics Firm (3PL) and paying a fee on a per container basis. This last alternative might be desirable if the total volume that will be going through the crossdock is not very large and if the company does not desire to make an investment or a long term rental contract.

3. Motivation

Robert Bosch LLC, or Bosch in short, is a leading international automotive parts manufacturer that has several plants in North America. Bosch has applied lean manufacturing practices in its

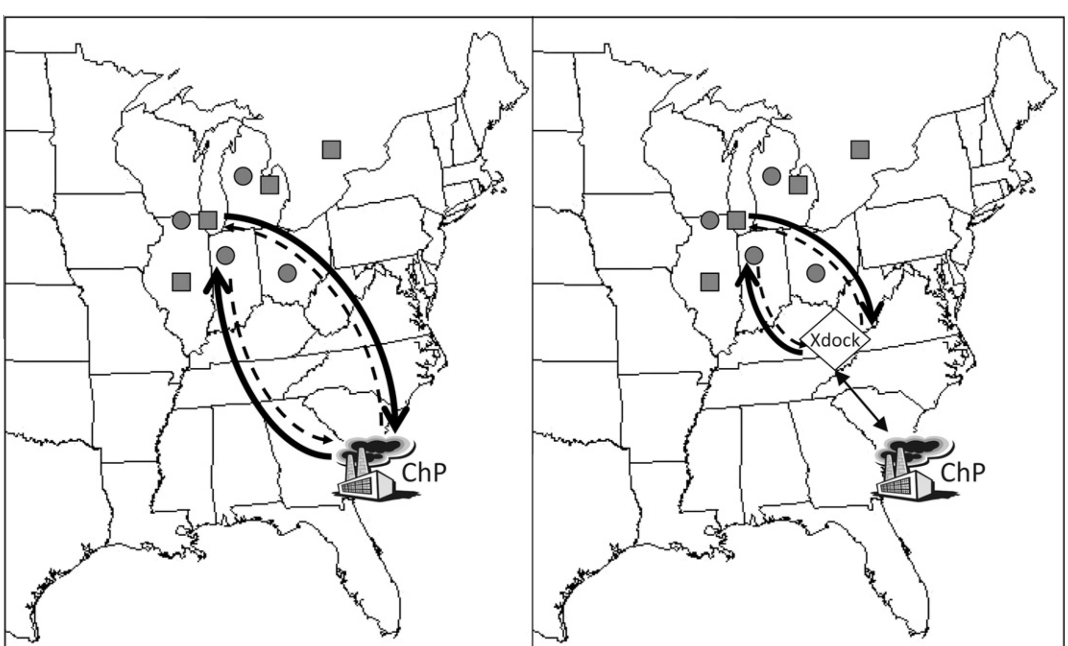

Fig. 2. Uneven flows from and to the suppliers and customers, represented by squares and circles, respectively, in both directions going over long distances can be matched for better coordinated utilization, sometimes along a crossdock.

production system for the last few years. Lean manufacturing is the production of goods using less of everything compared to traditional mass production. Lean manufacturing is a generic process management philosophy inspired by the Toyota Production System (Womack et al., 1991) and implemented in Bosch as the Bosch Production System or BPS (Made in Europe Magazine, 2004). Our study was motivated by Bosch's desire to improve its logistics operations and it was done at Bosch's Charleston, SC plant (Bosch/ChP). In addition to direct transportation cost savings, attaining additional benefits, such as reducing inventories, is also desired since one of the principles of lean is evening out the production flow by reducing batch sizes and increasing delivery frequencies internally and, if possible, externally.

4. Problem description

In this project, we look into selecting a subset of customers and suppliers of Bosch/ChP and combining their shipments on the same routes. We also consider the use of a crossdock located close to the customer and supplier base to use as a consolidation center. We investigate whether such changes bring net savings to the system, where the system consists of Bosch/ChP, customers (automotive manufacturers) and suppliers (second tier suppliers). Every node in the system may have both pick ups and deliveries: Each customer has a demand from Bosch/ChP and it may also have a supply if the containers used are returnable. Similarly each supplier has a supply to Bosch/ChP and it may also have a demand if the containers used are returnable. The supply and demand at each node has to be served simultaneously by the same vehicle.

In addition to the vehicle routing aspect of the problem, we also have to decide on whether to use a crossdock as a transshipment node. The crossdock cost is a variable cost that is charged on a per container basis, so there is no fixed rental or investment cost. The location of the potential crossdock is predetermined.

We look at the problem from a supply chain coordination perspective. In the new system, Bosch/ChP will make all the transportation arrangements including shipments for its customers. To implement this change, materials supplied to customers will have to be re-priced to factor in shipping costs to Bosch/ChP. In the current system, Bosch/ChP takes care of the transportation arrangements for its suppliers, and all of its customers handle the transportation by themselves.

5. Literature review

At the core of this problem is vehicle routing, which is defined as the process of selecting paths in a transportation network along which to send physical traffic. The goal is to provide services to spatially dispersed customers via delivery routes that have one or more customers on them.

There are variants of the Vehicle Routing Problem (VRP) based on different aspects, such as the number of depots (one/multiple), cost structure (fixed/variable/both), time to service a particular node (specified/time windows/unspecified), size of fleet (one/multiple), type of fleet (identical/non-identical), nature of demands (deterministic/stochastic), existence of pickups along with demands at each node and vehicle capacity (imposed/non-imposed).

The basic version of the problem is defined as follows: there is only one depot where the vehicles begin their routes and we are given its location. We also know the locations and demands of customers. We have a number of available vehicles that have a finite capacity. This can be a physical capacity, such as the weight or volume, or it can be a restriction imposed due to the nature of the system, such as the total distance or total time of the route. The goal is to find a set of routes that minimize the total distance the vehicles travel.

VRP is a very difficult problem to solve. In fact, even with the very recent developments, the size of the problems that can be solved by exact solution approaches are in the order of a hundred customers. Among the most effective exact methods are the branch and cut method in Lysgaard et al. (2004), the branch and cut and price method in Fukasawa et al. (2006) and the set partitioning approach with cuts in Baldacci et al. (2008).

Due to the difficulty of VRP, most of the techniques developed are heuristics, including local search methods, population based algorithms and learning mechanisms. These methods generate quality solutions very fast but do not guarantee optimality. The best performers, in terms of accuracy and computing time seem to be the hybrid methods presented in Mester and Bräysy (2005), Tarantilis and Kiranoudis (2002) and Prins (2004) that combine population search and local search. A broad summary of recent work on VRP is given in the book by Golden et al. (2008) including a categorized review on metaheuristics by Gendreau et al. (2007).

In the version of the VRP that we are considering in this paper, every node may have both pick ups and deliveries. The general class of such problems is referred to as Pickup and Delivery Problems (PDPs) in the literature. To the best of our knowledge and based on the examination of a very recent survey on PDPs (Berbeglia et al., 2007), the version of the VRP considered in this paper has not been studied before. The closest versions are the *VRP with Simultaneous Pickups and Deliveries*(VRPSPD), and *VRP with Pickups, Deliveries and Transshipments*(VRPPDT), as defined in Berbeglia et al. (2007). Recent works on VRPSPD include Bianchessi and Righini (2007), Hoff et al. (2009), and Gajpal and Abad (2009). These papers present and compare various heuristics including tabu search and ant colony system. There are also a few studies on VRPPDT including the papers by Cortés et al. (2010), Mitrovic-Minic and Laporte (2006), and Kerivin et al. (2008).

6. Modeling assumptions

Cost linearity

In order to compare the cost of the new system with the existing one, we need to know the current cost of customers' pickups from Bosch/ChP. To estimate those costs, we assume that the cost of a truck is linear in the percentage fill rate (i.e. if a full truck costs C, half a truck costs $\frac{C}{2}$). This allows us to assign a cost figure to the pickups from Bosch/ChP by the customers. For instance, assuming that the length of a route does not change after removing Bosch/ChP from it, the cost associated for the pickup from Bosch/ChP is directly proportional to the percentage of material pickup from Bosch/ChP with respect to the whole load in that truck. If 30% of the load of a truck is picked up from Bosch/ChP, the cost of this pickup is equal to the 30% of the whole route. This linearity assumption implies that once the shipment from Bosch/ChP to a customer is removed from customer's current route in the new system, the customer can and will fill in the free space with its other suppliers' material, which are already on the existing route of the customer, possibly by decreasing the frequency of that route.

This assumption can be justified by noticing the fact that customers will re-optimize their new routes after removing Bosch/ChP from their list of stops. We can not know the effect of not having Bosch/ChP on any of their routes as this requires the knowledge of all of their suppliers and their pickup schedules from them. This effect can be an increase in total cost if the customers can not fill in the empty space efficiently. Or it can be a decrease in total cost if removing Bosch/ChP from their route significantly reduces the length of the route or increases the utilization of their existing route capacities. Because of this lack of knowledge, we assume that customers' networks are large enough that removing Bosch/ChP from their system will have a neutral effect. This assumption is approved by several experts in the Logistics Department of Bosch/ChP and by the experts of the 3PL company currently used by Bosch/ChP.

A closed system for empty containers

We assume that the returnable container flow is a closed loop. Whatever returnable container sent to customers comes back to Bosch/ChP. Similarly, whatever returnable container received from suppliers is sent back to suppliers.

7. Problem formulation

We developed a network flow model combined with vehicle routing from the crossdock and the plant. The formal formulation is given in the Appendix. In this model, we allow two new means of transportation for each customer and supplier in addition to their current way of transportation. One is being on a milkrun starting and ending either at the plant or the crossdock. In the other alternative, a customer can receive products from and send empty containers to the crossdock using LTL carriers. Similarly, a supplier can send products to and receive empty containers from the crossdock using LTL carriers. The consolidated shipments at the crossdock are shipped either with round-trip (or one-way) TL carriers between the crossdock and the plant or the crossdock may be on milkruns that have some other intermediate stops.

In this model, in addition to the transportation cost, we also have a crossdock cost. The goal is to minimize the total cost in the new system.

The constraints that we need to satisfy are similar to the basic VRP: We need to pick up all supplies and empty customer containers. We need to deliver all customer products and supplier empty containers. Each customer and supplier can use only one mode of transportation and the demands and supplies can not be split into more than one mode and even into more than one truck (This constraint, of course, implicitly assumes that none of the demand or supply values exceed the capacity of a truck. In this project, this assumption holds for the customers and suppliers of Bosch/ChP). There is also a limit on the total weight and total volume of the trucks. As for the volume, at most 70% of the total volume should be utilized while planning for the shipments. This is the current practice of Bosch/ChP to ensure that there will be enough space left for unexpected demand fluctuations, which may not be captured during the planning phase. In addition to these constraints, there are additional constraints imposed by Bosch/ChP to ensure the quality of customer service and have easy-to-coordinate routes.

7.1. Customer sensitive constraints

- *There can be at most one customer per new milkrun.* This ensures that we will not have to coordinate shipments for two customers.
- *If a customer is on a new milkrun, then it should be the first stop on the milkrun.* This ensures that the promised delivery times can be met by Bosch/ChP.
- *Customers should not see degradation in mode of shipment* (LTL carriers are considered less reliable than TL carriers).

In addition, Bosch/ChP also wants the following constraints to be satisfied:

- *The maximum number of stops on a route to be at most 4 (excluding the origin).* This is the current practice Bosch/ChP has with its suppliers and it is thought that a higher number of stops on a route may be difficult to coordinate and schedule.
- *If the crossdock is on a new milkrun, then it should be the last stop.* This gives Bosch/ChP more accurate information and also more control on the arrival time of trucks to Bosch/ChP, as there will be no intermediate stops between the crossdock and Bosch/ChP.

In the next section, we discuss the data collection and preprocessing phases of this project.

8. Data and preprocessing

We collected customer and supplier data for a quarter of a year. We also collected detailed packaging information, such as volume and weight when full and empty and when collapsed if collapsible. After that, we computed necessary statistics such as average demand, supply, weights, space utilization for all customers and suppliers. Each customer has an average demand that needs to satisfied, but it also has an average supply, which is the average amount of empty returnable containers. This supply value is calculated separately for each customer, based on the type of containers used for the shipments of that customer. Similarly, every supplier has a supply that needs to be delivered to Bosch/ChP, but it also has a demand value, which is the average amount of empty returnable containers that needs to be delivered back to the supplier.

Moreover, we identified current customer and supplier routes. The supplier routes were easy to obtain as they are being formed by the 3PL that Bosch works with. But the customer routes were completely unknown. So, to obtain the route information and the demand load on these routes, we interviewed truck drivers to piece together this information when they come to Bosch/ChP for pickups.

8.1. Cost structure

We estimated the TL costs based on the expertise of the traffic engineers and studies conducted by the 3PL that Bosch/ChP works with. We used a single fixed per mile cost for TL carriers irrespective of the routes as we see that on average the costs of routes do not differ significantly based on the endpoints of the route. The LTL costs are estimated using the rate software of USF-Holland, an LTL carrier. We tuned these costs by comparing with actual LTL shipment data. The crossdock cost is also a variable cost that is charged on a per container basis for which there is no fixed rental cost or investment.

8.2. Initial analysis

For each potential customer, we computed the savings for the case where the customer does not pick up from the plant and Bosch/ChP delivers it from its crossdock instead. We considered several cases depending on customer pickup type (TL/Oneway TL/LTL) and identified most promising customers based on volume and total potential savings and included these promising customers in the optimization model.

9. Solution and results

Our results are obtained by using the average demand data and supply data over a quarter of a year for the customers and suppliers, respectively. Using this average data, we solved increasingly complex scenarios and the results are summarized in Table 1. In this and later tables, the savings percentage reflects the reduction

Table 1
Results for different scenarios.

Scenario	No. of customers	No. of suppliers	Constraint set	Savings (%)	No. of variables	No. of constraints
1	5	8	Basic	26	6744	2401
2	5	8	All	14	6744	3451
3	5	18	All	25	30201	9244

in cost as a percent of the current costs, which are calculated as follows.

For each supplier, whether served with LTL or TL carriers, we have the necessary data such as the length of the routes and the costs associated with them. For each customer, this is not exactly known since customers handle their pick ups themselves and Bosch has very limited information about their costs. In the proposed new solutions, Bosch/ChP may be removed from some of the customers' existing routes. To calculate the cost change that customers will realize with the removal of Bosch/ChP from their routes, we estimate the cost that can be associated solely to the pickups from Bosch/ChP as follows: For customers that use LTL carriers, this is estimated using the LTL rates between Bosch/ChP and the customer's location. For customers who use TL carriers, this is more complicated. First, the length of the current route, which includes Bosch/ChP as a stop, is calculated. Let's refer to this *CurrentLength*. Then, the length of the new route after removing Bosch/ChP from the current route is calculated. Let's refer to this *NewLength*. Let's also assume that *P* represents the percentage of Bosch/ChP's load on a customer's truck in the current system. Then using the cost linearity assumption stated before, we can estimate the cost reduction in the customer's route due to the removal of Bosch/ChP from the current route as $CurrentLength - (1 - P) * NewLength$. The sum of all the costs for the suppliers and customers is the base total cost. The new system's cost is compared to this base total cost and the savings are given as a percentage.

In the first scenario, we have five customers and eight suppliers and we do not impose the additional customer sensitive constraints. In the second scenario, we have the same set of five customers and eight suppliers, as we have in the first scenario, but now we also impose the customer sensitive constraints. In the third scenario we have the same set of 5 customers but a larger set of 18 suppliers and we impose the customer sensitive constraints.

When we compare the percentage savings between Scenarios 1 and 2, we see that the extra customer sensitivity constraints reduce these savings, which is not surprising and reflects the general trade-offs seen in business. Moreover, when we compare Scenarios 2 and 3, we see that including more suppliers in the model increased the savings significantly as it allowed more options and better matching of the flows from and to the plant.

To illustrate what a solution looks like, we depict the current and future states for Scenario 2 in Fig. 3, where customers and suppliers are represented by circles and squares, respectively. The future state is significantly different to the current state. In the future state, C1 and C2 continue using their old milkruns. C3, S4, S6 and S8 use LTL service from and to the crossdock. There are two milkruns starting and ending at the crossdock. The remaining two suppliers (S5 and S7) use LTL service from and to ChP as they used to do in the current system. There is also a dedicated round-trip TL truck between the crossdock and ChP.

9.1. Computational issues

The model is formulated using the OPL modeling language and is solved using CPLEX 11.1 on a PC with an Intel Core 2 Duo 3.0 GHz processor and 2 GB of RAM. In terms of computational difficulty, Scenario 3 is significantly more difficult than the other two scenarios. Scenario 1 and Scenario 2 are solved to optimality in 4 min and 5 min, respectively. We allowed the model to run for 24 h for Scenario 3, but it was unable to solve it to optimality. After 24 h, the best feasible solution found by CPLEX corresponded to a cost savings of 25% and the best upper bound found corresponded to a cost savings of 28%. To shorten the running times, we created a new restricted model for Scenario 3: we first classified the suppliers into two groups based on the amount of supplies they supply. In the restricted model, only the largest eight suppliers are allowed on new milkruns. The remaining 10 suppliers are not considered in new milkruns. This restriction made the model significantly easier to solve. The restricted model is solved to optimality in only 11 min, resulting in a solution with 25% cost savings, which is equal to the savings found by running the original model for 24 h. The savings obtained with the restricted model are also very close to the upper bound obtained for the original model (thus, only at most 3% away from the optimal solution of the original model) and the restricted model is solved much faster. Because

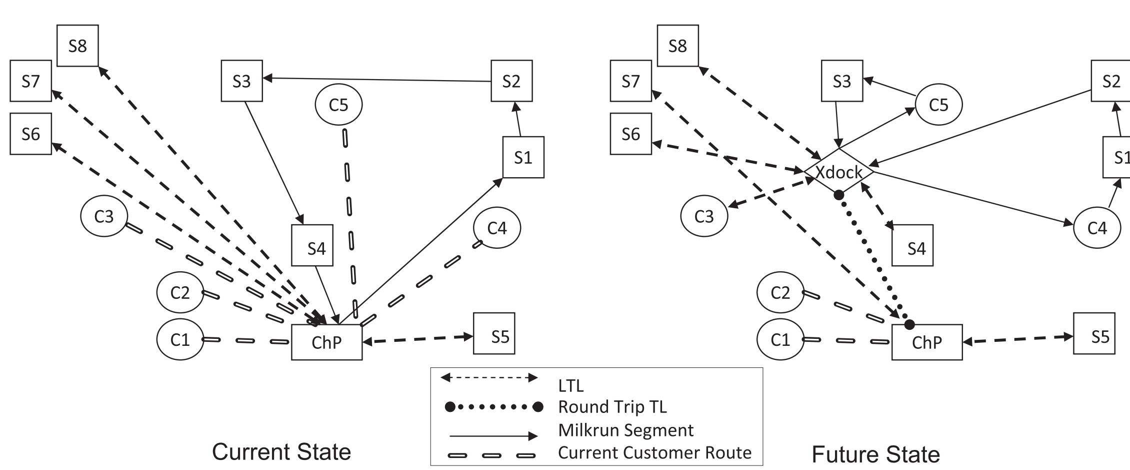

Fig. 3. Illustration of current and future states for Scenario 2.

of this, the restricted model is also used in all the following analysis.

9.2. Robustness analysis

We investigate the robustness of the savings based on the demand variability over the weekdays. For this we compute the average data for each weekday and separately solve the optimization model over these average daily data. To keep the consistency in the transportation modes across the days, we fix the customer transportation modes in the model based on the solution we previously obtained on the overall average data in Scenario 3. Thus, if a customer uses its old milkrun in the solution of Scenario 3, then that customer will be forced to use its old milkrun on each of the week days. Similarly if a customer uses a new milkrun in the solution of Scenario 3, then that customer will be forced to use a new milkrun on each of the days, but in that case we will allow the model to pick different routes on the different days of the week as long as it is a new milkrun. For each week day, the average savings for the week day are shown in Fig. 4, where we see that the savings fluctuate significantly across the days. The savings are highest on Monday and Thursday, because one of the customers, which is placed on a new milkrun based on the result of Scenario 3, has demands on these two days only. The overall average savings is 18%.

We also solve the model simultaneously on the average data of Mondays and Tuesdays, which have significantly different demand patterns. In that case, we force customers and suppliers to have the exact same routes, which we refer to as *stable routes*, on both days. Interestingly forcing the system to have the same routes on both days reduces the savings by only 1% (from 18% to 17%). We should note that the change in cost might have been different if were able to solve the model simultaneously over the average data of five days, instead of just two. But we were unable to do that as that model becomes very large and difficult to solve with the current computational capabilities we have.

We do a similar robustness analysis for the weeks of the quarter, where we solve the model for the average weekly data. The results for the first five weeks are illustrated in Fig. 5. To find stable routes over the different weeks, we solve the model

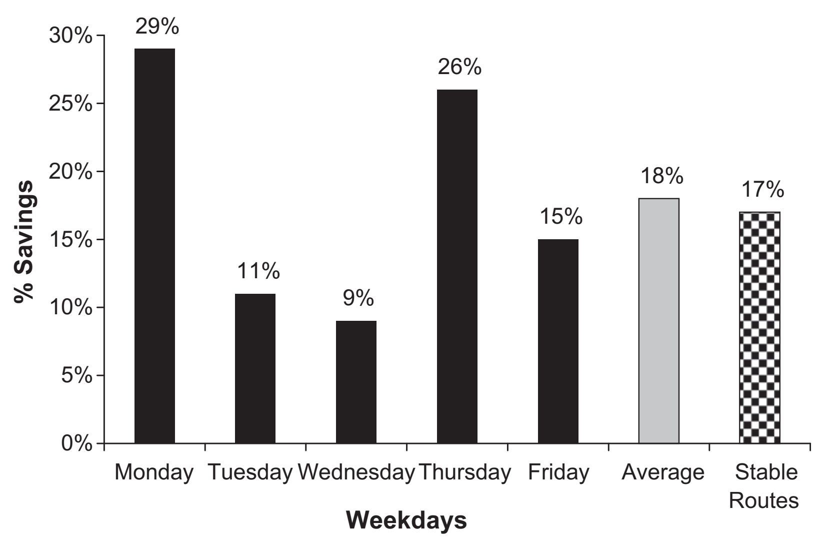

Fig. 4. Robustness analysis over the weekdays. The savings fluctuate significantly as different days have different demand structures.

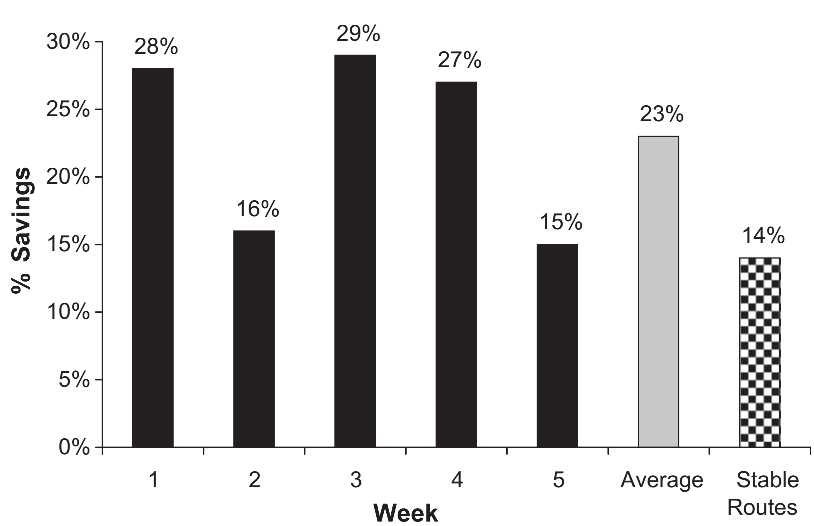

Fig. 5. Robustness analysis over the weeks of the quarter. The savings fluctuate significantly as different weeks have different demand structures.

simultaneously over two different weeks. For that, we chose the 1st and 12th weeks of the quarter since they are far away from each other and have significantly different demand patterns. The overall average savings is 23%. We see that the savings across the weeks fluctuates but the intensity of the fluctuation is less than the fluctuation we have seen across the days of the week in Fig. 4. Forcing to have the same routes on both weeks reduces the savings by 9% (from 23% to 14%).

We also do a "worst-case" scenario analysis in the following way: We take the 90th percentile of the demand and supply requirements of every customer and supplier and solve the model over that data. Obviously we should note such a demand and supply pattern is extremely unlikely. Nevertheless, the results will be useful to see the robustness of the model. Since we are feeding the model with very high demand and supply data, we test the model with different maximum volume utilization values. As noted in the problem formulation, the current value required is 70%. The results are illustrated in Table 2. In each of the maximum volume utilization rates, the savings with the high demand and supply rate are significantly ($\sim$ 10%) less than the savings with the average data. We also see that as we increase the max truck fill rate, the savings increase in both cases.

Table 2
Max truck volume utilization on average and 90th percentile data.

Truck fill rate (%)	% Savings	
	Average data (%)	90th percentile (%)
70	25	14
80	27	16
90	29	20
100	32	21

9.3. Utilizing the crossdock

One alternative setting that extensively utilizes the crossdock may be appealing to Bosch/ChP, because in such a system, there will be dedicated trucks between the crossdock and Bosch/ChP and all the routing will be made at the crossdock and this will be an easy-to-manage system. As clearly indicated in Fig. 6, forcing the flow of customers and suppliers, which do not use their old routes in the new solution, to go through the crossdock reduces the savings we obtained before as this restricts the solution space. When we solve the model with this restriction, we see that the savings are actually reduced to 21%, as opposed to the 25% we had before. These savings are at the 70% max truck fill rate. If we allow 100% truck fill rate, the savings goes up by only 3% (to 24%), whereas without the restriction, the savings were much higher (32%).

9.4. Inventory considerations

In addition to the savings in transportation costs, this new system will result in reduced inventories for some of the customers, suppliers and also for Bosch/ChP because the new system assumes that deliveries and pickups will be made every day. This means an increase in the frequency of deliveries and pickups for some customers and suppliers, which results in reduced inventories since shipments will be made in smaller batches. We should also note that although smaller batches reduce inventories, there may be increased costs due to the additional number of batches being dealt with. These costs may relate to administration costs or costs of checking each batch.

The potential reduction in inventories is illustrated in Fig. 7, where the graph on the left is for the case where a customer's demand occurs only twice a week, namely on Mondays and

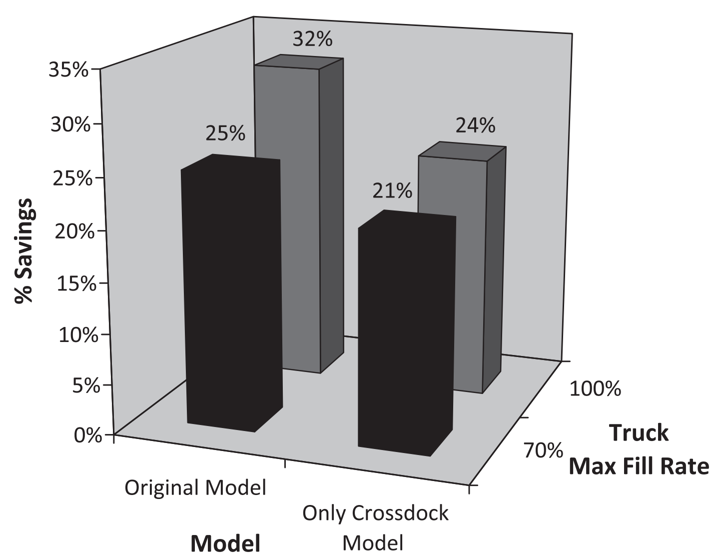

Fig. 6. Forcing the flow of customers and suppliers, which do not use their old routes, to go through the crossdock reduces the savings at both truck fill rate limits.

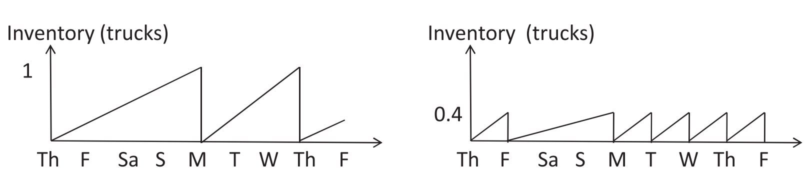

Fig. 7. Inventory reduction when the frequency of shipments is increased from twice a week (graph on the left) to five times a week (graph on the right).

Thursdays, whereas the graph on right is for the case where the same customer's demand, which is equal to 2 trucks per week, is spread to five weekdays. In the first case, the inventory is built from Thursday to Monday to meet the demand on Monday and from Monday to Thursday to meet the demand on Thursday. In the second case, Monday's demand is built from Friday to Monday and every other weekday's demand is built during the previous day. Clearly the inventory amount in the second case is significantly lower (60%) than the first case, not only for Bosch/ChP but also for the customer. Thus, our solution methodology is aligned with the philosophy and current implementation of BPS and represents a win for all involved parties.

9.5. Other considerations

There are some other potential side benefits of this planning model and the new solution it generates. With this tool, Bosch/ChP will be able to offer different service options to its customers, such as delivering to their location, or allowing them to pickup from the crossdock instead of Bosch/ChP. The frequency of the shipments can also be adjusted based on customers needs. In addition, Bosch/ChP thinks that this kind of system-wide improvement will potentially increase their chances to have more demand from auto manufacturers that value continuous improvement and lean practices.

A key challenge in implementing the proposed solution is the perceived lack of control the customers feel over their inbound shipments, which they rightly view as being crucial to their operations. Another business consideration is how the delivery contracts should be re-priced to add the increased delivery costs for Bosch, and how this will impact negotiating for customer contracts.

10. Summary and extensions

This study was motivated by Bosch's desire to explore the opportunity to extend lean principles, which are applied internally within the framework of BPS, to the external logistics operations and it was carried out at Bosch's Charleston, SC plant.

In this project, we selected a subset of customers and suppliers to consolidate their shipments on the same routes. We also considered the use of a crossdock as a consolidation point. We developed a network flow model combined with vehicle routing from the crossdock and the plant to minimize the system-wise logistics costs. We identified the potential for significant cost savings over the current system based on our model assumptions.

We presented our findings to a wide range of people at Robert Bosch North America (RBNA). These include the plant managers and directors and managers of Customer Service, Logistics and Planning departments of two plants, including ChP. Moreover, we have presented it to the CEO and also to the Logistics Director of RBNA. They were happy with the overall outcome, the tangible results that showed real potential for cost savings and with the different design alternatives that we presented. The systematic analysis of the logistics processes that we performed was also highly appreciated.

The challenges mentioned earlier have prevented the immediate implementation of our proposed solution since it involves renegotiation of customer contracts. However, our analysis of the matching opposite and complementary flows has triggered an ongoing discussion on similar issues in RBNA's overall logistics operations, where both flows are within the purview of (different) RBNA locations. The benefits of supply chain coordination, which our study clearly revealed, has provided them with strong evidence for significant savings by overall coordination of logistics operations within RBNA's network.

We finally would like to point out two potentially beneficial methodological extensions. The first is the use of stochastic programming to handle the demand uncertainty. One important restriction would be the exponential number of discrete scenarios. Even with coarse characterization of demand and supply values, such as a three tier characterization as low-average-high, the number of scenarios for n nodes will be 3^n. With small n, the stochastic programming approach may be used, but for large n, it may be computationally not possible or practical.

The second extension is the use of a route-based restricted model. In that, we would consider only a set of routes instead of all possible ones in solving the problem. For that model, we need to generate a set of good routes using heuristics. We certainly want routes to have low costs. But also we want to generate routes that have daily stable loads as we will run the model every day. For this, one can build routes that have suppliers with negatively correlated shipments on the same routes. So, when the shipment of a supplier is high in one day then it is likely that the other suppliers supply will be low and the load on this route will be stable.

Acknowledgements

This work has been supported by the Carnegie Bosch Institute at the Tepper School of Business, Carnegie Mellon University. We thank the institute and its President Mr. Johannes K. Elling for supporting this project and providing his valuable comments. We also would like to thank Dr. Sunder Kekre for his suggestions. Finally we thank all the managers and associates of the Customer Service, Logistics and Planning department of Bosch/ChP for providing much valuable input.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.ejor.2010.03.044.

References

Baldacci, R., Christofides, N., Mingozzi, A., 2008. An exact algorithm for the vehicle routing problem based on the set partitioning formulation with additional cuts. Mathematical Programming 115 (2), 351–385.

Berbeglia, G., Cordeau, J., Gribkovskaia, I., Laporte, G., 2007. Static pickup and delivery problems: A classification scheme and survey. TOP 15 (1), 1–31.

Bianchessi, N., Righini, G., 2007. Heuristic algorithms for the vehicle routing problem with simultaneous pick-up and delivery. Computers and Operations Research 34 (2), 578–594.

Cortés, C., Matamala, M., Contardo, C., 2010. The pickup and delivery problem with transfers: Formulation and a branch-and-cut solution method. European Journal of Operational Research 200 (3), 711–724.

Fukasawa, R., Longo, H., Lysgaard, J., Aragão, M., Reis, M., Uchoa, E., Werneck, R., 2006. Robust branch-and-cut-and-price for the capacitated vehicle routing problem. Mathematical Programming 106 (3), 491–511.

Gajpal, Y., Abad, P., 2009. An ant colony system (ACS) for vehicle routing problem with simultaneous delivery and pickup. Computers and Operations Research 36 (12), 3215–3223.

Gendreau, M., Potvin, J., Bräysy, O., Hasle, G., Løkketangen, A., 2007. Metaheuristics for the vehicle routing problem and its extensions: A categorized bibliography. In: Golden, B.L., Raghavan, S., Wasil, E.A. (Eds.), The Vehicle Routing Problem: Latest Advances and Challenges. Springer, Boston.

Golden, B., Raghavan, S., Wasil, E., 2008. The Vehicle Routing Problem: Latest Advances and New Challenges. Springer.

Hoff, A., Gribkovskaia, I., Laporte, G., Løkketangen, A., 2009. Lasso solution strategies for the vehicle routing problem with pickups and deliveries. European Journal of Operational Research 192 (3), 755–766.

Kerivin, H., Lacroix, M., Mahjoub, A., Quilliot, A., 2008. The splittable pickup and delivery problem with reloads. European Journal of Industrial Engineering 2 (2), 112–133.

Lysgaard, J., Letchford, A., Eglese, R., 2004. A new branch-and-cut algorithm for the capacitated vehicle routing problem. Mathematical Programming 100 (2), 423–445.

Made in Europe Magazine, Q., 2004. The Bosch Production System. Made in Europe Magazine Q4.
Mester, D., Bräysy, O., 2005. Active guided evolution strategies for large-scale vehicle routing problems with time windows. Computers and Operations Research 32 (6), 1593–1614.
Mitrovic-Minic, S., Laporte, G., 2006. The pickup and delivery problem with time windows and transshipment. Infor-Information Systems and Operational Research 44 (3), 217–228.
Prins, C., 2004. A simple and effective evolutionary algorithm for the vehicle routing problem. Computers and Operations Research 31 (12), 1985–2002.
Tarantilis, C., Kiranoudis, C., 2002. BoneRoute: An adaptive memory-based method for effective fleet management. Annals of Operations Research 115 (1), 227–241.
Womack, J., Jones, D., Roos, D., 1991. The machine that changed the world: The story of lean production. HarperPerennial.