Vehicle Routing Challenges in the Automotive Supply Chain

Robert R. Inman,
Chief Data and Analytics Office
General Motors Company
Warren, MI 48092, USA
robert.inman@gm.com

Rana Afzali-Baghdadabadi and Baiyang Liu
Global Purchasing and Supply Chain
General Motors Company
Warren, MI 48092, USA
rana.afzali@gm.com, baiyang.liu@gm.com

Abstract

This paper describes an open and vital vehicle routing research opportunity. We describe the pervasive vehicle routing problems (VRPs) automakers use for three different links in their supply chains. Our research methodology was to analyze an automaker's logistics problems to identify the most relevant VRP characteristics and search the VRP literature for contributions addressing those characteristics. Arguably the most challenging automotive VRP is parts collection for which we detail the key problem characteristics of capacitated trucks, periodic schedules and stochastic volumes. We provide industrial data on the routes, supplier volumes and locations. This shows the large scale of automotive VRPs and characterizes the randomness in the daily shipping volumes. We observe that the daily volumes are stochastic with generally positively skewed distributions and positive autocorrelation. Since the large-scale automotive VRPs would be computationally intractable for the existing methods on capacitated VRPs with stochastic volumes, research applied to the automotive industry typically assumes deterministic volumes. Since we observed meaningful randomness in the supplier volumes, this identifies a need for research on periodic capacitated VRPs that are both stochastic and large-scale. This paper's data will serve as a test-bed for future research on this important problem facing the automotive industry.

Keywords
Milkrun, capacitated, periodic, stochastic, data

1. Introduction

There are hundreds of automotive assembly plants globally which together produced over 95 million motor vehicles in 2018 (International Organization of Motor Vehicle Manufacturers, 2019). Each of these large automotive assembly plants receives hundreds of different part numbers daily (and many parts multiple times per day) to assemble into vehicles. Coordinating the delivery of hundreds of parts every day – in a reliable and economic manner – poses a huge logistical challenge. Milk-runs help solve this challenge. The term milk-run, derived from the delivery system of milk distributors delivering fresh milk to the homes of subscribers on a regular basis, is the colloquial term for a collection or distribution route that visits multiple locations in a tour. Milk-runs are a staple for just-in-time inbound parts supply to automotive assembly plants. Nemoto et al. (2010) observe that milk-runs are becoming the standard system for just-in-time parts supply. A survey of the trade literature uncovers reports of milk-runs used by Toyota (Ludwig and Williams 2016), Nissan (Ludwig 2015), Honda (Newton 2011), Volkswagen (Ludwig 2017b), Porsche (Ludwig 2010), Magna Steyr assembling vehicles for BMW, Daimler and Peugeot (Ludwig 2013), Fiat (Kumar 2014), Scania (Danby 2015), Skoda (Hogg 2015a), Chery Automotive (Hogg 2015b), Jaguar Land Rover (Ludwig 2014), Volvo (Ludwig 2017a), and Ford (Automotive Logistics 2009). Designing effective milk-runs is vitally important because of their pervasive use, on a large scale, on a daily basis, worldwide.

Figure 1 shows where milk-runs fit in the automotive supply chain. For inbound parts deliveries to the assembly plant, some suppliers, such as supplier 4, can justify dedicated shipments for one of two reasons. First it could be that supplier 4's volume is great enough to nearly fill a truck by itself in which case there is no savings by adding it to a route with another supplier. Or second, supplier 4 could be so isolated that the distance to another supplier is so great that it is less expensive to send a dedicated truck to supplier 4 than combining with another supplier. For many other suppliers, such as suppliers 1, 2, and 3, the milk-run deliveries are preferred.

The reason for inbound milk-runs' utility for the automotive industry is not hard to understand. In pursuit of lean production, automakers prefer daily shipping. As Iyer et al. (2009) point out, one of Toyota's philosophies is frequent deliveries so that suppliers send at least one shipment per day. Receiving every part every day, or even more frequently, benefits automotive assembly lines by reducing the inventory inside the plant, reducing the floor space needed in the plant, reducing scrap (e.g., if a supplier's process drifts out of control and produces out-of-specification parts, then all the parts in the pipeline must be scrapped after recognizing the quality issue), reducing misplaced parts, eliminating double handling if delivered directly to the line or in-sequence, and reducing the feedback time after quality concerns. Frequent parts delivery is a key lean manufacturing principle. But for suppliers whose shipping volumes are small (i.e., less than a truckload of parts a day), daily deliveries could become expensive if the truck only delivered this single supplier's parts. It would be wasteful to send a nearly empty truck to the assembly plant every day. For suppliers with *much* less than a truckload per day, daily deliveries with a dedicated truck would be *prohibitively* expensive. Collection routes mitigate the expense of daily deliveries for small volume suppliers because a single truck can collect parts from multiple suppliers on a frequent basis.

As depicted in the right part of Figure 1, automakers' supply chains also include the outbound delivery of finished vehicles from the assembly plant to dealers. For finished vehicle deliveries, milk-runs are effective for the journey's final leg to deliver vehicles to the dealers. There may be a rail or ocean leg from the assembly plant to a rail yard or port, but the journey's final leg is almost always via a car-hauler truck. These car-haulers can typically carry five to twelve vehicles depending on the truck design, vehicle size, and road weight and size restrictions. If a single dealer does not have enough deliveries to fill a car-hauler truck, the car-hauler will deliver vehicles to two or more nearby dealers. Automakers strive to deliver vehicles quickly because the vehicle storage yard has limited capacity, vehicles waiting to be delivered incur a carrying cost, and most importantly a customer may be waiting. Hence vehicles usually must be shipped within a few days of production. Some dealers, such as dealer D (analogous to supplier 4), can justify dedicated shipments because of their volume or isolated location. For many other dealers, such as dealers A, B and C, the milk-run deliveries are preferred. To serve as a test bed for research, Appendix A provides data on actual vehicle deliveries to dealers.

In addition to inbound parts collection routes and outbound vehicle delivery routes, automakers have a third supply chain to manage – repair parts. Automakers not only sell new vehicles, they also service the vehicles they sold. The automakers' dealer networks provide repair and maintenance service. Automakers typically maintain regional warehouses and deliver repair parts from the warehouses to dealers. The schematic for the repair parts supply chains

© IEOM Society International

would be the same as that shown in Figure 1 replacing "*Automotive Assembly Plant*" with "*Regional Parts Warehouse*." Like the assembly plant, the parts warehouse receives shipments from the suppliers and delivers parts to the dealers. Since the volume of parts needed for repairs is much less than for assembly, the shipments of parts from suppliers to the parts warehouses are only as needed in accordance to an inventory stocking policy and are not typically performed by milk-runs. However, automakers do use milk-runs to deliver repair parts from their warehouses to dealers. For example, Schittekat and Sorenson (2009) provide a tool to generate routing solutions for delivering spare parts to dealers that includes designing milk-runs as a sub-problem.

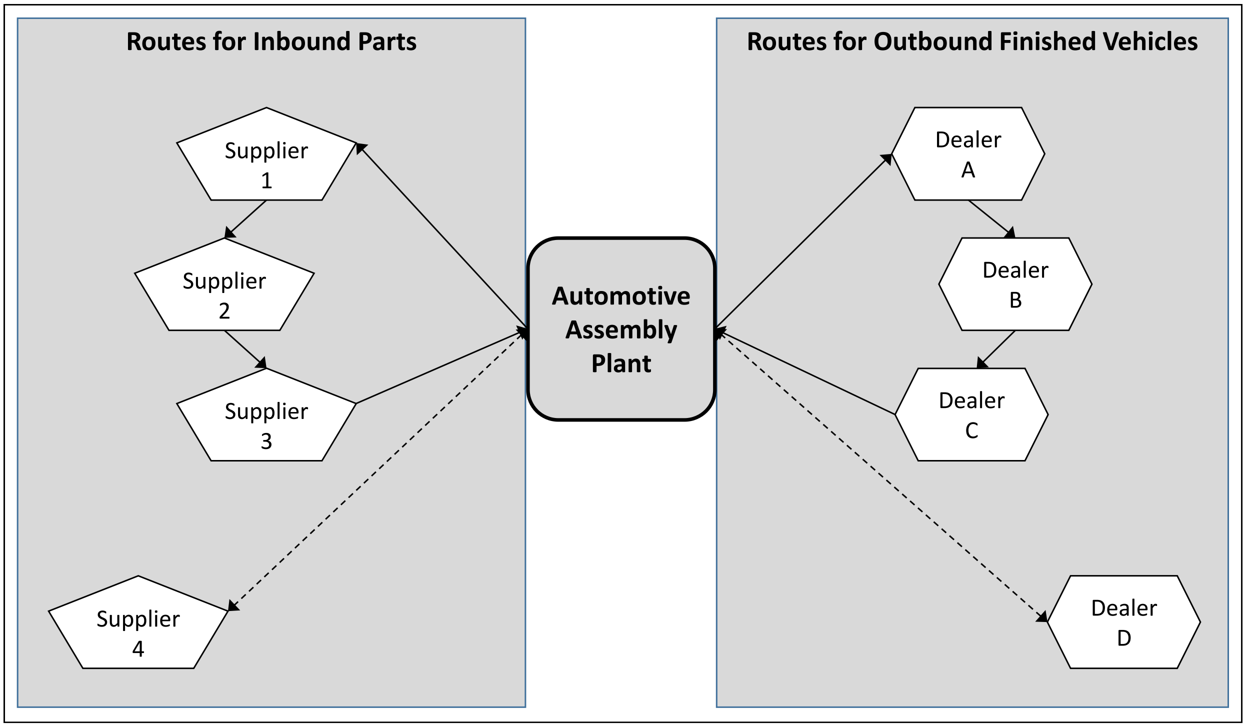

Figure 1. Inbound and outbound routes for the automotive supply chain

Milk-runs are one type of vehicle routing problem (VRP). Bertsimas (1992) analyzed the capacitated VRP with stochastic demands. Laporte et al. (2002) studied the capacitated VRP with stochastic demand and provided an algorithm assuming Poisson or normally distributed customer demand. Sungur et al. (2008) provided a robust optimization approach to capacitated VRPs with stochastic demand that minimized the transportation costs while satisfying the demand in a bounded uncertainty set. Dong and Turnquist (2015) optimize the pickup frequency at suppliers and routes to minimize the total logistics costs. They show significant savings by allowing the flexibility to have different pickup frequencies instead of requiring all suppliers to have the same pickup frequency.

The milk-runs for the three automotive supply chain links (parts collection for assembly, finished vehicle delivery to dealers, and spare parts delivery to dealers) have several characteristics in common. All three have capacitated vehicles, all three seek to minimize cost which depends on travel time and distance, all three have a strict time limit on the route duration. All three are *capacitated* VRPs. However, the parts collection routes for automotive assembly are also *periodic* and *stochastic*. The periodic characteristic signifies that the same route is run repeatedly on a fixed period, and the stochastic characteristic signifies that the volumes to be collected during each tour are random. Unlike vehicle delivery routes (or package delivery) where the route is redesigned for every tour, the inbound routes remain in place for months at a time and must collect random amounts of material from the suppliers on that route. Since the periodic and stochastic nature of inbound parts milk-runs make them more challenging than the repair parts or finished vehicle delivery milk-runs, the rest of this paper focuses on inbound automotive parts collection routes.

The paper proceeds as follows. Section 2 describes the key characteristics of inbound parts collection routes for automotive assembly plants. Section 3 provides data from an automaker's milk-run suppliers to serve as a test-bed for further research and points out the meaningful randomness in daily volumes. Section 4 describes existing research on inbound automotive milk-runs (that assume deterministic volumes) and highlights the opportunity for research on large-scale VRPs explicitly comprehending the randomness observed in the data.

2. Description of Parts Collection Routes in the Automotive Industry

Milk-runs have clear advantages and are widely used for inbound parts collection routes in the automotive industry. Boysen, Emde, Hoeck, and Kauderer (2015) review the automotive industry inbound logistics literature and highlight the importance of milk-runs. De Moura and Botter (2016) demonstrate the cost effectiveness of milk-runs over direct shipments for supplying parts to an automotive assembly plant. They point out their additional advantage of reducing congestion at the plant by reducing the number of truck deliveries (that also eases the scheduling of unloading docks). We now highlight key characteristics of the VRP for automotive parts collection.

2.1 Large Scale

The automotive parts collection VRPs are large scale. For example, Ludwig and Hogg (2016) report that truck maker Scania uses milk-runs for inbound parts to its truck assembly plant in Brazil that cover 160 suppliers, using 50 trucks per day with an average of 100 daily pickups. Battani, Boysen, and Emde (2012) report that a BMW plant in Germany receives material from about 600 suppliers on more than 400 trucks a day. For a BMW plant in the US, Coia (2013) reports that they receive about 740 trucks a day. Nemoto et al. (2010) describe the Toyota plants in Thailand that use milk-runs to collect parts from about 120 suppliers using 600 trucks on about 50 collection routes for each plant. The number of suppliers to an automotive assembly plant is well over 100 requiring many parts collection routes.

2.2 Drive Time Constraint

A key constraint for the automotive parts collection VRP is the driver time limit. Many countries limit the number of hours truck drivers can drive each day and week. For example, in the US, the US Department of Transportation (2019) requires that property-carrying drivers may drive a maximum of 11 hours after 10 consecutive hours off duty. For milk-runs, the route duration is not just driving time. Milk-run drivers must pick up materials from multiple suppliers, and each supplier stop takes time to unload empty containers and load full containers. The drive time constraint can limit the number of stops or constrain the supplier locations on the same route. To avoid assigning a team of drivers for a route, every effort is made to design routes that can be completed within the legal driver time limits.

2.3 Capacitated

Trucks have finite capacities which must be considered when designing routes (de Moura et al., 2016). In addition, weight must be considered as both the trucks and roads have weight limits.

2.4 Periodic

Automotive parts supply VRPs are periodic. Automakers design inbound milk-runs to be in place for several weeks at a time (De Moura et al., 2016; Ludwig and Williams, 2016). Given the problem's scale with hundreds of trucks arriving daily, rescheduling the trucks and drivers daily would require an extremely flexible workforce and large on-call fleet of trucks. There are several reasons for this. The first reason is that it is costly and time consuming to design and source routes. Since most automakers contract delivery from independent trucking companies, there is a lengthy process of issuing requests-for-quotes (RFQ) for each route. Once the RFQ is broadcast, trucking companies have time to develop bids, which are then reviewed by the automaker who chooses the best. The process of designing and then sourcing routes takes many days if not weeks. Hence the routes will be in place for many weeks. A second reason for periodic routes is the automakers' lean inventory strategy. Liker (2017) describes Ford Motor Company's lean "every part every day" philosophy towards supplier parts delivery and Rechtin (2002) goes further reporting that Ford ideally wanted every part to be delivered every shift. This strategy of a steady regular flow of parts in turn leads to periodic deliveries. A third reason for periodic routes is to simplify material management in the plant. Since there are limited receiving docks and limited material handlers at the plant, automakers schedule the deliveries using time windows in a manner similar to airlines scheduling their arrivals to make use of their finite gates. The schedule eliminates contention for scarce dock time and wasteful traffic congestion and queuing time. The schedule also enables the material handlers at each dock to stage the empty containers and know ahead of time where each truck's parts will go (De Moura et al., 2016). A fourth reason is to facilitate returning empty parts containers. Automotive suppliers usually ship components in sturdy returnable containers to protect components during transportation and reduce the consumption of cardboard for environmental conservation. After installing the components, the assembly plant sends the empty returnable containers back to the correct suppliers (Boysen et al, 2015; De Moura et al., 2016). This constraint argues for consistent routes that visit the same suppliers every tour to minimize the handling of empty containers and make sure that they get back to the correct supplier. Finally, running a fixed schedule of routes simplifies the logistics and material handling and production control. With a fixed schedule, it is easy to know when a delivery is late and which parts are impacted. Automotive assembly plants are extremely risk averse regarding parts deliveries because a missing part is so expensive. Malucci (2006) reports that if an automotive assembly line shuts down for lack of a single part, it could cost the manufacturer $20,000 per minute. For all these reasons, the vast majority of inbound automotive milk-runs are designed to be on a periodic schedule for many iterations.

2.5 Stochastic

In addition to being large scale, capacitated and periodic, parts collection routes for automotive assembly are stochastic. Although existing literature on inbound automotive VRPs assumes deterministic volumes at suppliers, this assumption is probably due in large part to the computational burden of solving a large scale stochastic capacitated VRP. The next sub-section displays supplier and route volume data for an automotive assembly plant quantifying the amount of randomness facing industry.

3. Industrial Data for Automotive Parts Collection Routes

To elucidate the nature of the stochastic process facing inbound automotive milk-runs, we provide descriptive statistics from a large automotive assembly plant that operated 16 milk-runs to daily collect parts from 47 suppliers. The data presented is in terms of fractions of truckloads using a standard US trailer with 3,860 cubic feet of capacity. Since the components' containers typically don't allow full utilization, the practical capacity is somewhat less than this. But instead of imposing a safety factor on the data, we provide the liquid cube. This paper reports actual shipping volumes. The routes described were developed based on forecast data before this actual data occurred. So, these routes may well not be optimal for the observed data. Nevertheless, they are useful benchmarks and the route data provides insights into the nature of the milk-run route design problems facing the automotive industry. The period of record spanned 83 production weekdays. Table 1 shows descriptive statistic for the milk-run routes collecting parts from suppliers and delivering to the automotive assembly plant each day. Many routes have a wide range of requirements, one route had a day of zero volume, and the skewness varies from positive to negative.

In addition to the descriptive data on the existing routes, we provide the data for each of the suppliers. Table 2 displays the descriptive statistics for each supplier's daily shipping volume (as a fraction of the standard truckload). To enable use as a test-bed for VRP research, we provide the latitude and longitude of each supplier relative to the plant. Table 2 shows considerable variability in the daily supplier volumes. There are several reasons for this. Modern mixed-model automotive assembly plants not only build many different models, they build them in many different configurations which require different parts. Very few part numbers are required on every vehicle built in a plant. The vast majority depend on the model and configuration. Hence the required volume varies day by day for the vast majority of parts. Although forecasts for the number of vehicles that will be produced in a day are accurate, forecasting the configuration (and hence the specific parts) is much less accurate. Moreover, the sequence of ordered vehicles built in the plant is scrambled due to rework loops and parallel operations. In addition, the granularity of containers causes variability. Suppliers ship only full (fixed capacity) containers of parts. A container may hold 48 units of a particular part. If the average part usage was 72 per day, even a perfectly known and constant 72 per day, the shipping schedule would be 48 one day and 96 the next. Moreover, material handling issues such as scrap, misused or stolen parts, and misplaced parts introduce additional variability and uncertainty. For these reasons the quantity of a part to be shipped varies over time and is uncertain. Of course, other plants and OEMs will have different levels of variability.

Most suppliers had significant randomness in their daily volumes. Almost all had positive skew which makes assuming that they are distributed normally (or with any symmetric distribution) underestimate the risk. These volumes are liquid cube truckloads that assume there is no wasted space after loading containers into the truck trailer and that all the containers can be stacked to the trailer's ceiling. Using a 20% rule of thumb for wasted space, in the 83-day 16-route data there were 93 cases of exceeding the truck's capacity. In other words, even though the highest average route liquid cube was 77% (from Table 1) we expect at least one route to violate its truck's capacity every day. This industrial periodic capacitated VRP is especially challenging because of random volumes with positively skewed distributions. Besides the day-to-day variability in supplier volumes, many suppliers' volumes are autocorrelated over time. A practical recourse to exceeding a truck's capacity today (or this period) is to carry as much as we can today and pick up the rest next period in addition to next period's volume using the regularly scheduled route. If supplier volumes are negatively autocorrelated, this recourse strategy would often be effective because it would mean that above average volumes are likely to be followed by below average volumes. In addition to the practical impact of autocorrelated volumes, a common theoretical assumption in stochastic modelling is that is that the volumes for different days are independent, and that the volume in a particular period is independent of its volume in other time periods. Since each milk-run contains at least two suppliers, once designed we are more concerned with the autocorrelation in each route's volume than in the individual supplier volumes. Figure 2 displays the histogram of lag-1 autocorrelations for the existing routes alongside the lag-1 autocorrelations for the suppliers on those routes. The dashed lines in Figure 2 depicts the threshold for statistical significance at the 95% confidence level. Positive autocorrelation was widespread; 9 out of 16 routes and 19 out of 47 suppliers had statistically significant positive autocorrelation. This positive route autocorrelation increases the risk of expedited freight. If a route exceeds its capacity today, we may attempt to not expedite today hoping to catch up tomorrow by adding today's excess to tomorrow's shipment. But since the route volumes are positively autocorrelated, if the route's volume is high today it is likely to be high tomorrow as well.

Table 1. Descriptive statistics for existing milk-runs

Route	**Min**	**Max**	**Mean**	**Median**	**Standard Deviation**	**Skewness**
A	0.50	0.73	0.54	0.54	0.03	4.81
B	0.54	1.06	0.77	0.78	0.09	0.12
C	0.45	1.09	0.72	0.80	0.15	(0.93)
D	0.52	0.81	0.59	0.59	0.03	3.43
E	0.14	0.82	0.29	0.16	0.25	1.42
F	0.56	0.91	0.69	0.70	0.06	0.45
G	0.42	0.61	0.47	0.45	0.05	1.24
H	0.60	0.88	0.65	0.65	0.03	4.95
I	0.59	0.96	0.69	0.68	0.04	3.01
J	0.40	0.68	0.45	0.44	0.04	3.14
K	0.48	0.70	0.52	0.52	0.03	4.68
L	0.39	0.55	0.42	0.41	0.02	2.32
M	0.26	0.39	0.28	0.28	0.02	4.13
N	0.63	0.96	0.70	0.70	0.04	2.96
O	0.71	1.05	0.77	0.77	0.04	4.81
P	0.43	0.63	0.51	0.48	0.06	0.58

4. Designing Parts Collection Routes for Automotive Assembly Plants

Designing parts collection routes for automotive assembly are stochastic periodic capacitated routing problems. For most automakers, the period is daily or more frequently if the volume justifies it. Campbell and Wilson (2014) described the broad applicability of the periodic VRP (of which inbound automotive milk-runs are one circumstance) and the literature addressing it. They highlighted the need for new solution methods that explicitly deal with application's specific constraints and objectives. Importantly, their review found little work that explicitly addresses stochasticity. This deficiency limits the applicability of prior research to the automotive industry where stochastic demand is a key element of milk-run design. Shedl and Straus (2011) compare a genetic algorithm to a simulated annealing approach to determine periodic routes for dairy delivery. In their work, the number of routes is given, and design routes for each day of the week that specify which customers to receive a delivery and the sequence of that day's tour. They model each customer's demand as an independent normal distribution and simulate the performance of alternative solutions. Similarly, Moghaddam et al. (2012) highlight the need for more consideration of capacitated VRPs with stochastic demand. The data in Table 2 provides a realistic test-bed of the randomness in automotive inbound milk-runs that will help researchers test new approaches for automotive applications.

© IEOM Society International

Table 2. Descriptive statistics for suppliers

Supplier	Min	Max	Mean	Median	Standard Deviation	Skew	Supplier Lat -Plant Lat	Supplier Long -Plant Long
A1	0.37	0.55	0.40	0.40	0.02	4.70	(4.41)	(7.53)
A2	0.13	0.19	0.14	0.14	0.01	4.81	(2.53)	(5.31)
B1	0.28	0.73	0.52	0.53	0.09	(0.33)	(7.82)	(3.35)
B2	0.18	0.39	0.25	0.25	0.04	0.73	(7.77)	(3.64)
C1	0.27	0.41	0.30	0.31	0.01	4.43	(8.66)	0.75
C2	0.17	0.68	0.42	0.50	0.14	(0.98)	(8.06)	1.75
D1	0.12	0.17	0.13	0.13	0.01	1.74	(16.96)	(14.06)
D2	0.22	0.33	0.24	0.24	0.01	3.80	(17.11)	(13.81)
D3	0.12	0.20	0.14	0.14	0.01	2.23	(16.81)	(14.54)
D4	0.06	0.11	0.07	0.07	0.01	1.43	(16.81)	(14.54)
E1	0.01	0.69	0.16	0.03	0.25	1.48	(0.16)	(1.70)
E2	0.12	0.18	0.13	0.13	0.01	4.28	0.89	(1.58)
F1	0.11	0.16	0.12	0.12	0.01	4.86	(0.75)	(0.91)
F2	0.44	0.75	0.57	0.59	0.06	0.16	(0.81)	(1.90)
G1	0.29	0.41	0.32	0.31	0.03	1.36	(0.24)	(2.41)
G2	0.08	0.11	0.09	0.08	0.01	1.21	(1.42)	(2.15)
G3	0.06	0.09	0.07	0.06	0.01	1.38	(1.03)	(0.98)
H1	0.07	0.10	0.08	0.08	0.00	4.05	(1.42)	(2.15)
H2	0.25	0.36	0.27	0.27	0.01	4.34	(0.24)	(2.41)
H3	0.13	0.21	0.15	0.15	0.01	1.36	(0.05)	(1.98)
H4	0.12	0.20	0.15	0.15	0.01	3.08	(0.91)	(2.80)
I1	0.05	0.08	0.06	0.06	0.00	3.59	(1.38)	(1.31)
I2	0.34	0.49	0.37	0.36	0.02	4.96	(0.99)	(1.07)
I3	0.19	0.39	0.26	0.26	0.03	0.89	(0.81)	(1.90)
J1	0.09	0.13	0.09	0.09	0.00	3.62	(7.26)	(0.29)
J2	0.22	0.36	0.27	0.27	0.02	1.67	(7.58)	1.18
J3	-	0.32	0.05	0.04	0.04	3.92	(6.48)	(3.67)
J4	-	0.07	0.04	0.05	0.02	(1.73)	(7.18)	(0.25)
K1	0.19	0.31	0.23	0.23	0.01	3.05	(0.47)	(0.06)
K2	0.27	0.39	0.29	0.29	0.01	4.77	(0.41)	0.75
L1	0.32	0.46	0.35	0.34	0.02	2.98	(1.63)	1.95
L2	0.06	0.09	0.07	0.07	0.01	1.27	(0.38)	0.40
M1	0.03	0.04	0.03	0.03	0.00	4.77	(0.39)	0.65
M2	0.13	0.21	0.15	0.15	0.01	4.08	(0.28)	0.27
M3	0.01	0.02	0.01	0.01	0.00	3.18	(0.22)	0.06
M4	0.08	0.12	0.09	0.09	0.00	4.73	(0.50)	0.76
N1	0.33	0.54	0.39	0.39	0.03	2.12	1.15	6.30
N2	0.16	0.23	0.18	0.18	0.01	1.64	0.81	4.21
N3	0.12	0.19	0.14	0.14	0.01	3.12	1.15	6.30
O1	0.30	0.43	0.32	0.32	0.02	4.77	0.67	4.08
O2	0.12	0.18	0.13	0.13	0.01	4.71	0.81	4.21
O3	0.27	0.44	0.32	0.33	0.02	3.23	0.97	2.95
P1	0.12	0.20	0.14	0.14	0.01	3.17	0.27	3.24
P2	0.17	0.25	0.19	0.18	0.01	2.43	0.46	2.49
P3	-	0.03	0.00	0.00	0.00	3.95	0.21	3.95
P4	-	0.12	0.04	-	0.05	0.73	1.05	4.22
P5	0.04	0.07	0.05	0.05	0.01	0.40	0.27	3.24
P6	0.08	0.12	0.09	0.09	0.00	3.94	0.27	3.24

© IEOM Society International

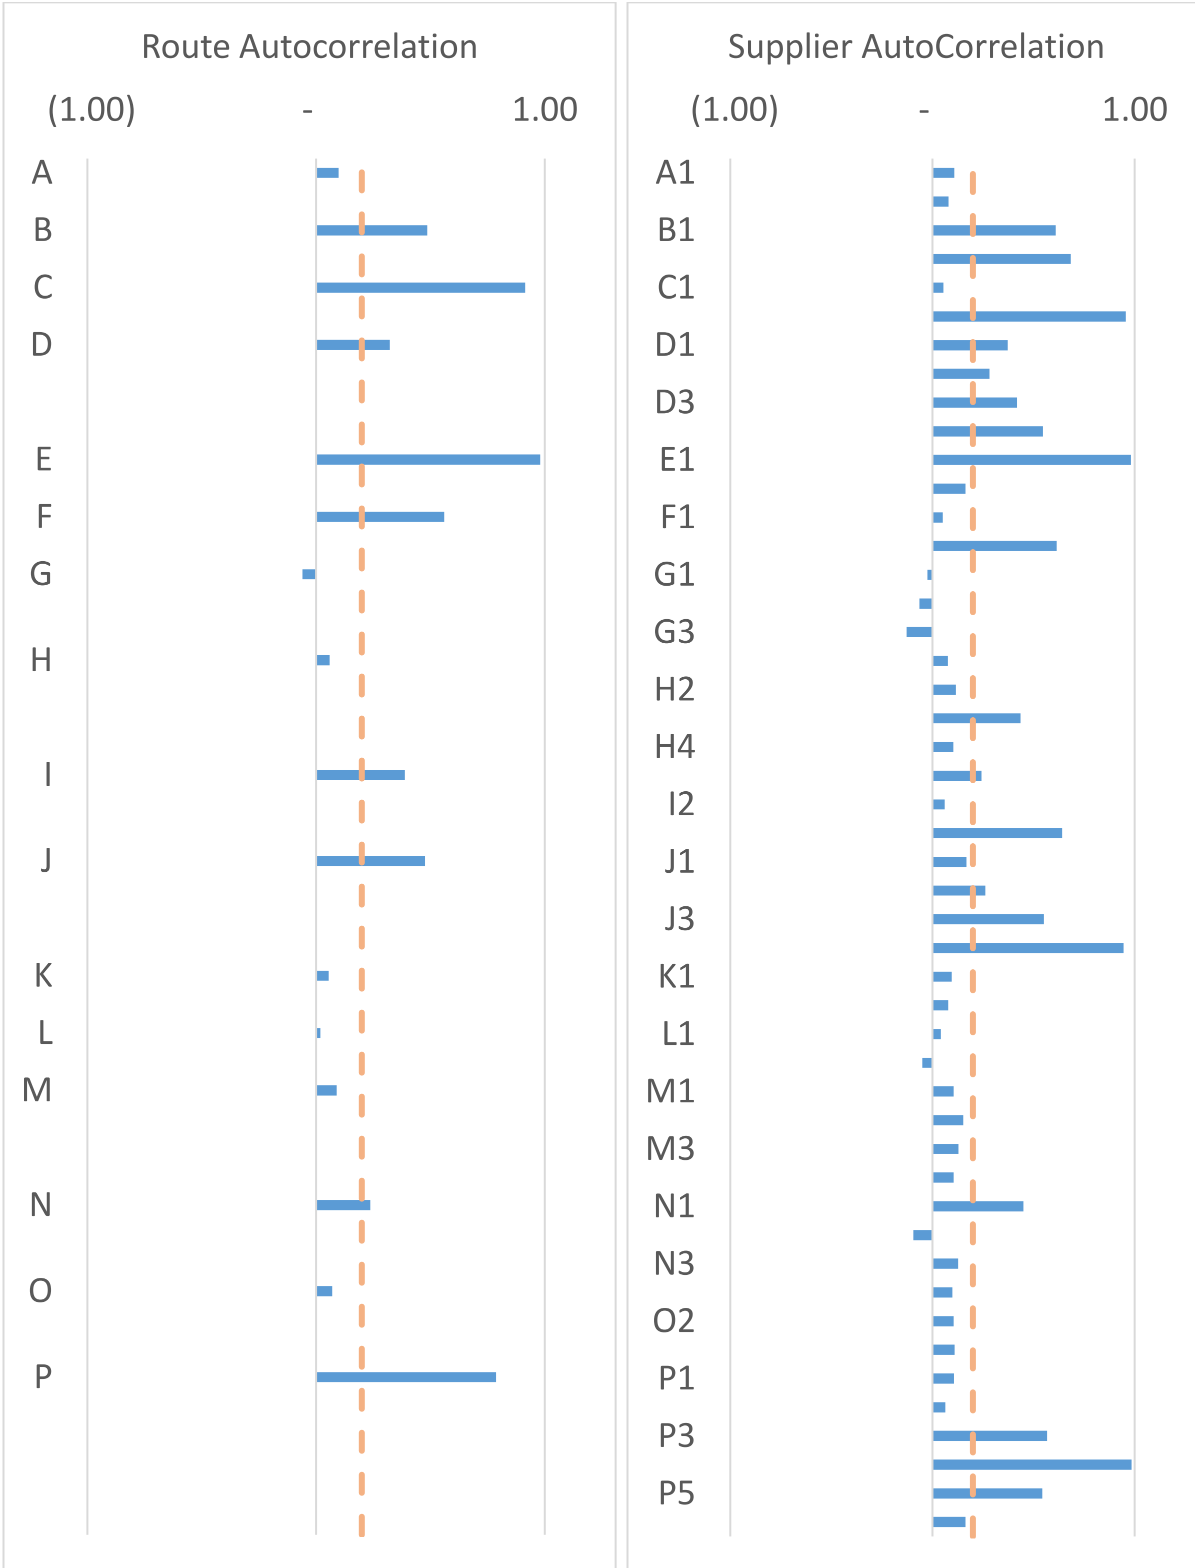

Figure 2. Autocorrelation within route and autocorrelation within supplier

Figure 3 displays a brief taxonomy of milk-run characteristics. The first split is between capacitated and un-capacitated milk-runs. The second split is if the route is periodic (i.e., if the route will be in place for several tours). The third split is if the amount of material collected (or delivered) is deterministic or stochastic. The shading roughly indicates the research devoted to those types of milk-runs. Darker shading indicates more research. The box for stochastic periodic capacitated milk-runs has no shading indicating relatively little research on this problem.

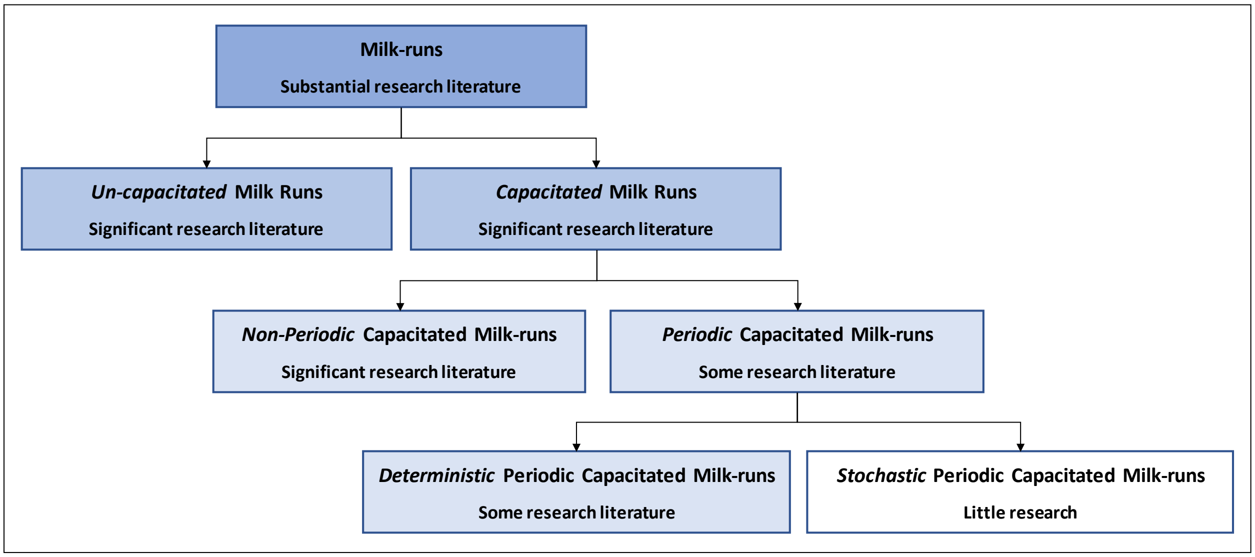

Figure 3. Taxonomy of milk-run problem characteristics with relative amounts of research

Theoretically much of the research on (non-periodic) vehicle routing problems with stochastic demand could apply. However, the computational complexity of many approaches makes application to the industrial-scale stochastic VRPs in the automotive industry intractable. For example, Jabali et al. (2014) point out that stochastic demand makes the VRP considerably more difficult to solve and that (at the time of their publication) the best algorithms can handle at most 3 routes and 50 suppliers (or customers). In practice, automotive assembly plants typically operate over a dozen milk-runs. Our data set had 16 daily milk-runs. Evidently due to the computational intractability of explicitly considering stochastic demand in the VRP when there are well over a dozen routes, applied work for automotive inbound collection routes typically assume deterministic demand when designing their periodic routes. The following papers address designing parts collection route design in an automotive setting – assuming deterministic volumes. Sadjadi, Jafari, and Amini (2009) use a genetic algorithm to design milk-runs for an automotive company assuming deterministic demand and non-static truck routes (i.e., they design a new route every day). Jafari-Eskandari et al. (2009) present a mixed integer problem for solving a deterministic milk-run VRP using data from an automotive assembly plant. Chuah and Yingling (2005) use column generation and tabu search to solve a deterministic milk-run design version VRP with time windows for suppliers of an automotive assembly plant. Alegre et al. (2007) optimize the periodic pick-up of parts for an automotive supplier assuming deterministic volumes. The above contributions on the deterministic version of the problem provide insights and good solutions. But if the suppliers in their problems had the same degree of randomness as those shown in Table 2, it seems that there is an opportunity to consider the stochastic nature of the VRP for automotive milk-run design. There is research on small scale *stochastic* capacitated VRP and there is research on large-scale *deterministic* periodic capacitated VRPs. There is a need for research on periodic capacitated VRPs that are *both* stochastic and large scale.

Designing parts collection routes for automotive assembly plants are large-scale stochastic periodic capacitated VRPs. The combination of large-scale and random volumes makes them extremely challenging – and under-researched. Although automotive parts collection routes are a narrow subset of more general vehicle routing problems, they are an enormously important subset. Hundreds of automotive assembly plants around the world operate multiple milk-runs daily. These milk-runs are critical to the efficient just-in-time global production of over 250,000 vehicles per day. There is abundant opportunity for research targeted at this rich area of industrial logistics networks.

References

Alegre, J., Laguna, M. and Pacheco, J., Optimizing the periodic pick-up of raw materials for a manufacturer of auto parts, *European Journal of Operational Research*, 179, 736-746, 2007.

Automotive Logistics, Ford and GM choose Ceva, Automotive Logistics, 15 July 2009.

Battani, D. Boysen, N., and Emde, S., Just-in-Time supermarkets for part supply in the automobile industry, *Journal of Management Control*, 24, 209-217, 2012.

Bertsimas, D. J., A vehicle routing problem with stochastic demand, Operations Research, 40, 574-585, 1992.

Boysen, N., Emde, S., Hoeck, M., and Kauderer, M., Part logistics in the automotive industry: Decision problems, literature review and research agenda, *European Journal of Operational Research*, 242, 107-120, 2015.

Campbell, A. N., and Wilson, J. H., Forty years of Periodic Vehicle Routing, *Networks*, 63(1), 2-15, 2014.

Chuah, K. H., and Yingling, J. C., Routing for a just-in-time supply pickup and delivery system, *Transportation Science*, 39, 328-339, 2005.

Coia, A., BMW's one and only in America, *Automotive Logistics*, December 3, 2013.

Danby, T., Maintaining traction through efficiency in Brazil, *Automotive Logistics*, 03 July 2015.

Dong, Z., and Turnquist, M.A., Combining service frequency and vehicle routing for managing supplier shipments, *Transportation Research Part E: Logistics and Transportation Review*, 79, 231-243, 2015.

Hogg, R., Skoda: Simply Clever logistics, Automotive Logistics, 06 July 2015a.

Hogg, R., South America conference: Chery hibernating to come back stronger, *Automotive Logistics*, 4 Nov 2015b.

International Organization of Motor Vehicle Manufacturers, 2018 World Motor Vehicle Production by Country and Type, accessed 6/16/2019, http://www.oica.net/category/production-statistics/2018-statistics/, 2019.

Iyer, A. V., Sehshadri, S., and Vasher, R., *Toyota Supply Chain Management* (McGraw-Hill, NY, NY), p. 79, 2009.

Jabali, O., Rei, W., Gendreau, M., and Laporte, G., Partial-route inequalities for the multi-vehicle routing problem with stochastic demand, *Discrete Applied Mathematics*, 177, 121-136, 2014.

Jafari-Eskandari, M., Sadjadi, S. J., Jabalameli, M.S., and Bozorgi-Amiri, A., A robust optimization approach for the milk run problem (an auto industry supply chain case study), *Int. Conf. on Com. & Ind. Eng.*, 1076-1081, 2009.

Kumar, R., Assembling a stable supply structure, *Automotive Logistics*, 17 Dec. 2014.

Laporte, G., Louveaux, F. V., and van Hamme L., An integer L-shaped algorithm for the capacitated vehicle routing problem with stochastic demands, *Operations Research*, 50, 415-423, 2002.

Liker, J. K., *Toyota Way: 14 Management Principles from the World's Greatest Manufacturer*, (McGraw-Hill, New York, NY), section 17.2, 2004.

Ludwig, C., Porsche balances in-and outsourced logistics, *Automotive Logistics*, 01 July 2010.

Ludwig, C., Logistics from concept to assembly at Magna Steyr, *Automotive News*, 2 Oct. 2013.

Ludwig, C., Dawn of a new era, *Automotive Logistics*, 24 June 2014.

Ludwig, C., Nissan North America part 3: The search for southern comfort, *Automotive Logistics*, 18 Dec. 2015.

Ludwig, C., and Hogg, R., South America Conf. 2015: Taking control in a crisis, *Automotive Logistics*, 28 Jan 2016.

Ludwig, C., and Williams, M., A connection worth more than the sum of its parts, *Auto. Logistic*s, 24 May 2016.

Ludwig, C., Volvo Cars China: A system built for growth, *Automotive Logistics*, June 29, 2017a.

Ludwig, C., Volkswagen Group Logistics: Come together, right now, *Automotive Logistics*, June 29, 2017b.

Moghaddam, B.F., Ruiz, R., and Sadjadi, S. J., Vehicle routing problem with uncertain demands: An advanced particle swarm algorithm, *Computers & Industrial Eng.*, 62(1), 306-317, 2012.

Malucci, L. I., Beneath the Surface, *APICS Magazine*, 16 (8), p. 12, 2006.

de Moura, D. A., and Botter, R. C., Delivery and pick up problem transportation - milk run or conventional systems, *Independent Journal of Management & Production*, 7, 3, 2016.

Nemoto, T., Hayashi, K., and Hashimoto, M., Milk-Run logistics by Japanese automobile manufacturers in Thailand, *Procedia Social and Behavioral Sciences*, 2, 5980-5989, 2010.

Newton, P., Honda to Overhaul Component Delivery System in Japan, *IHS Global Insight*, Feb. 21, 2011.

Rechtin, M., New Fiesta Sparks Assembly Line Changes, *Automotive News*, Jan. 14, 2002.

Sadjadi, S. J., Jafari, M., and Amini, T., A new mathematical modeling and a genetic algorithm search for milk run problem (an auto industry supply chain case study), *Int. J. of Adv. Manufacturing Technology*, 44, 194, 2009.

Schittekat, P., and Sorensen, K., OR Practice—Supporting 3PL Decisions in the Automotive Industry by Generating Diverse Solutions to a Large-Scale Location-Routing Problem, *Operations Research*, 57, 1058-1067, 2009.

Sungur, I., Ordóñez, F., and Dessouky, M., A robust optimization approach for the capacitated vehicle routing problem with demand uncertainty, *IIE Transactions*, 40, 509-523, 2008.

US Department of Transportation, Summary of Hours of Service Regulations, accessed 18 June 2019 https://www.fmcsa.dot.gov/regulations/hours-service/summary-hours-service-regulations.

Appendix A: Descriptive Data for Outbound Milk-runs

Table A.1: Example of 3 days of vehicle deliveries to dealers within 200 miles of plant

Dealer	Vehicles	Relative Longitude	Relative Latitude	Dealer	Vehicles	Relative Longitude	Relative Latitude
1	2	-0.4521	0.4841	43	1	0.5190	2.4059
2	1	-0.2052	-0.7429	44	1	0.9414	2.1776
3	1	0.4337	-0.5537	45	1	-1.3984	-1.5906
4	2	-0.3857	-0.7162	46	1	0.3404	2.4875
5	1	0.5807	0.5846	47	1	0.6218	-2.4999
6	1	-0.5053	-0.7413	48	1	-1.6885	1.2026
7	1	0.0055	1.0225	49	1	1.7412	-1.1557
8	1	-0.4209	-0.9009	50	1	-1.1085	2.1319
9	1	-0.2056	1.0698	51	1	0.7179	2.4923
10	4	-0.2634	-1.1212	52	1	-0.7231	-2.4683
11	1	0.0451	-1.2386	53	1	-1.0848	-2.3430
12	1	0.0451	-1.2386	54	1	0.6928	2.6938
13	1	0.5994	1.0481	55	20	0.6209	2.7631
14	1	-0.5192	-1.1153	56	5	-1.9735	0.9855
15	1	-0.0587	-1.3527	57	1	-2.1174	-0.0319
16	2	-0.6485	-1.0948	58	1	-2.0977	0.6999
17	1	-0.6136	-1.2024	59	1	-1.7518	1.8252
18	1	-0.1038	1.5305	60	1	-0.6632	-2.8965
19	2	0.7471	1.3299	61	1	0.7449	2.9794
20	4	-0.0656	-1.7083	62	1	-2.1726	0.9149
21	2	0.4576	-1.7286	63	2	0.8761	2.9553
22	1	0.4576	-1.7286	64	2	-1.6940	-2.0822
23	1	0.4769	1.7762	65	1	0.7963	2.9977
24	1	0.0117	-1.8958	66	3	-0.1895	-3.2677
25	1	0.0117	-1.8958	67	1	-2.1001	1.5967
26	1	0.4502	1.8300	68	1	-1.9462	1.9378
27	1	0.5339	-1.8505	69	1	1.7698	2.3827
28	1	-1.2683	1.0436	70	1	-0.0157	3.4087
29	2	0.5312	2.0005	71	1	-0.2245	3.5303
30	1	-1.5300	0.3545	72	1	-0.8236	3.3315
31	1	-1.5693	0.3688	73	1	1.6195	2.8706
32	1	-1.3380	-1.1861	74	1	0.9129	3.4191
33	1	-1.6145	0.1745	75	1	0.1765	3.5956
34	1	-1.2780	-1.4270	76	4	1.1061	3.3353
35	1	-0.2500	-2.2418	77	1	-2.3428	1.7085
36	1	-1.0835	-1.7627	78	14	-2.1103	-2.2955
37	1	-1.0835	-1.7627	79	1	-0.2185	-3.8222
38	1	-0.7223	-2.1273	80	2	-0.2185	-3.8222
39	1	-1.7480	-0.3187	81	1	-2.1657	2.4220
40	1	0.3927	-2.3794	82	1	-2.6185	-1.4750
41	1	-1.0054	1.9734	83	1	-1.8900	-2.8783
42	1	-1.4490	1.4481	84	1	1.3668	3.5640

© IEOM Society International

Table A.1 displays three days' worth of production destined to dealers within 200 miles of the assembly plant. The dealers are sorted by distance from the plant. All the vehicles produced at this plant destined to these dealers would be delivered on a haul-away truck. The load factor is the number of vehicles that can fit on at haul-away carrier and typically varies from four to nine depending on the size of the haul-away carrier and the size of the vehicles. This data set gives an example of the outbound milk-run design faced every day. Dealers 55 and 78 are large enough to justify full truckloads just to those dealers. However, the other dealers would almost certainly receive their vehicles on a milk-run with other nearby dealers. And of course, some of these dealers may have deliveries on milk-runs with more distant dealers.

The vehicle delivery problem changes every day. When designing the delivery milk-run, there is a tension between delivering vehicles quickly and making efficient milk-runs (i.e., shorter tour length and with fewer stops). Automakers typically impose constraints on the how long vehicles can wait before being delivered so the vast majority of vehicles are delivered via milk-runs and not dedicated shipments. The vehicle delivery problem is similar to the package delivery where new routes are designed every day. This is dramatically different from inbound parts delivery where routes are in place for weeks or months at a time and the same route is run repeatedly.

Biographies

Robert Inman is a Technical Fellow in General Motors' Data and Analytics Office. He began his GM career at GM Research and has also worked in Customer Experience and Global Purchasing and Supply Chain. He has contributed to General Motors with projects in industrial engineering, supply chain, manufacturing, vehicle sales, service, customer experience, dealer inventory, and operations management. Robert has served as a University of Michigan Tauber Institute of Global Operations faculty advisor for summer intern projects at 3M, Verizon Wireless, Brose, Dell, Precision Castparts, IDEX Liquid Controls and Amazon. Robert is a member of IEOM. He is an INFORMS Fellow and earned his PhD in Industrial Engineering and Management Sciences at Northwestern University.

Rana Afzali-Baghdadabadi is a senior researcher at General Motors. She has worked in GM's Operations Research group and Global Purchasing and Supply Chain for four years. Rana has developed and implemented mathematical models for logistics, containerization, operations management and supply chain strategy. Her interests are in identifying opportunities to improve and advance operational and strategic tasks through advanced optimization skills such as linear/mixed integer programming, stochastic modeling, dynamic programming and simulation. She is a member of INFORMS and IEOM. Rana earned BS and MS degrees in Industrial Engineering from Sharif University of Technology and a PhD in Industrial Engineering from the University of Missouri - Columbia.

Baiyang Liu is a General Motors senior researcher. Baiyang focuses on helping GM's global purchasing and supply chain organization improve operational efficiency and strategic planning. During her five years at GM, Baiyang has applied her mathematical modeling expertise together with her industrial experience to improve GM's supply chain performance and cost. She has successfully implemented projects in logistics, containerization, operations management and procurement strategy. Before joining GM, she gained purchasing and supply chain experience at Tesco and Logitech. Baiyang is a member of INFORMS and Vice President of the INFORMS Southeast Michigan Chapter. Baiyang earned a BS degree in Logistics Engineering from Shandong University and an MS in Industrial & Operations Engineering from the University of Michigan - Ann Arbor.