Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

SIMULATION BASED ANALYTICS FOR EFFICIENT PLANNING AND MANAGEMENT IN MULTIMODAL FREIGHT TRANSPORTATION INDUSTRY

Parijat Dube
João P. M. Gonçalves
Shilpa Mahatma
Francisco Barahona
Milind Naphade

IBM T. J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598, USA

Mark Bedeman

IBM Global Business Services
Avenida Diagonal, 571
Edifici "L'illa"
08029, Barcelona, SPAIN

ABSTRACT

The multimodal freight transportation planning is a complex problem with several factors affecting decisions, including network coverage, carriers and their schedules, existing contractual agreements with carriers and clients, carrier capacity constraints, and market conditions. Day-to-day operations like booking and bidding are mostly done manually and there is a lack of decision support tools to aid the operators. These operations are governed by a complex set of business rules involving service agreements with the clients, contractual agreements with the carriers and forwarder's own business objectives. The multimodal freight transportation industry lacks a comprehensive solution for end-to-end route optimization and planning. We developed analytics for trade lane managers to identify and exploit opportunities to improve procurement, carrier selection, capacity planning, and business rules management. Our simulation based analytics tool is useful for managing business rules and for doing what-if analysis which can lead to better resource planning, cost management, and rate negotiations.

1 INTRODUCTION

The multimodal freight transportation involves movement of freight through air, water and land. Planning of multimodal freight transportation is a complex problem involving different operations like transport, warehousing, distribution, and freight forwarding. Figure 1 shows a typical intercontinental multimodal freight network between Singapore and NYC. The network consists of five legs involving ocean, air and land routes with transit ports, regional distribution centers and local distribution centers.

The operations of freight industry are tied to several external factors like network coverage, carriers and their schedules, existing contractual agreements with carriers and clients, carrier capacity constraints, market conditions, weather conditions etc. Many planning and operational decisions have to be made under uncertainty associated with weather, market, and available capacity. The freight logistics industry has not taken the advantage of the digital revolution with much of the day-to-day operations still being done manually. These operations are governed by a complex set of business rules involving service agreements with the clients, contractual agreements with the carriers and forwarders own business objectives. The multimodal freight transportation industry lacks a dynamic and efficient decision support solution for end-to-end route optimization and planning.

A freight transportation network typically has three main stake holders: (i) Shippers, (ii) Freight forwarders, and (iii) Carriers. Freight forwarders are primarily responsible for end-to-end supply chain management of the freight transportation. The carriers are either already selected by the shippers or the freight forwarders also provide carrier selection services to the shippers. In the former case, the shippers typically have established contracts with the carriers with pre-negotiated rates. The freight forwarders

978-1-4799-7486-3/14/$31.00 ©2014 IEEE

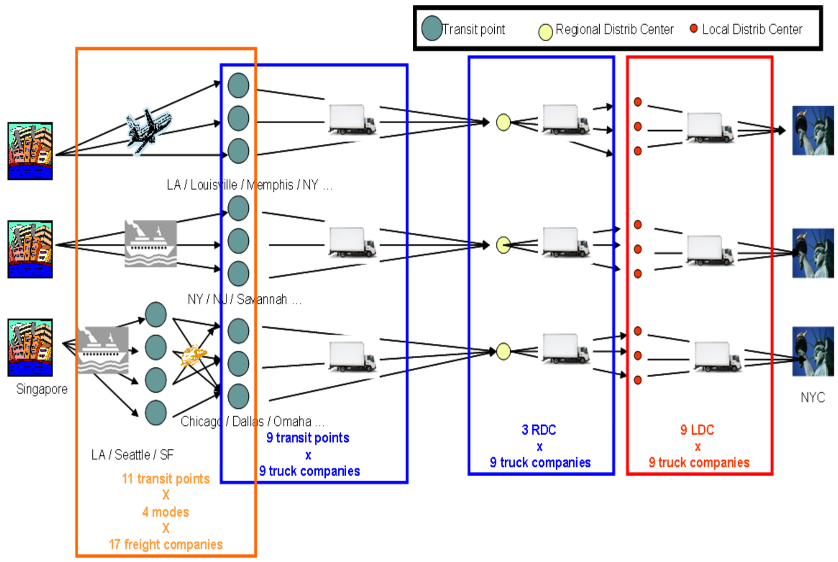

Figure 1: A multimodal freight transportation network between Singapore and NYC.

operate on ”trade lanes” which are characterized by a set of origin-destination (O-D) port pairs restricted to some geography, e.g., Mainland China to North-West Europe which includes O-D port pairs like Shanghai-Felixstowe, Xiamen-Rotterdham, Yantian-Belfast etc. Freight forwarders offer different types of services including less-than-container load shipments, full container load shipments, breakbulk, project forwarding, partial and full charter services, freight management services, bundled solutions, kitting and labeling etc.

We describe FR8NET, a decision support system developed at IBM Research for multimodal freight planning and operations. FR8NET has several features for empowering freight forwarders with data and decision support, including

- *Data Integration and Analysis Engine:* Provides data management services including data gathering from multiple heterogeneous information sources, data cleansing, normalization and data model building.
- *Booking Decision Support:* Provides decision support to booking operators at booking time by automated creation of feasible multimodal routes and calculation of their end-to-end cost, transit time and risk. This helps in maximizing network yield, enforcing compliance with strategic planning and leveraging market volatility.
- *Analytics Support:* Provides decision support to trade lane managers (TLMs) and product lane managers (PLMs) to identify and exploit opportunities to improve procurement, carrier selection, capacity planning, and business rules management.

The paper is organized as follows. Section 2 describes the deployment architecture of FR8NET including the data integration and analysis engine. Section 3 covers the booking decision support tool and its various features. The analytics support tool and its two components, Historical Analysis Module and Business Rules Management Module are presented in Section 4. Section 5 provides different analytics views offered by FR8NET. Related work is covered in Section 6. Finally we conclude in Section 7.

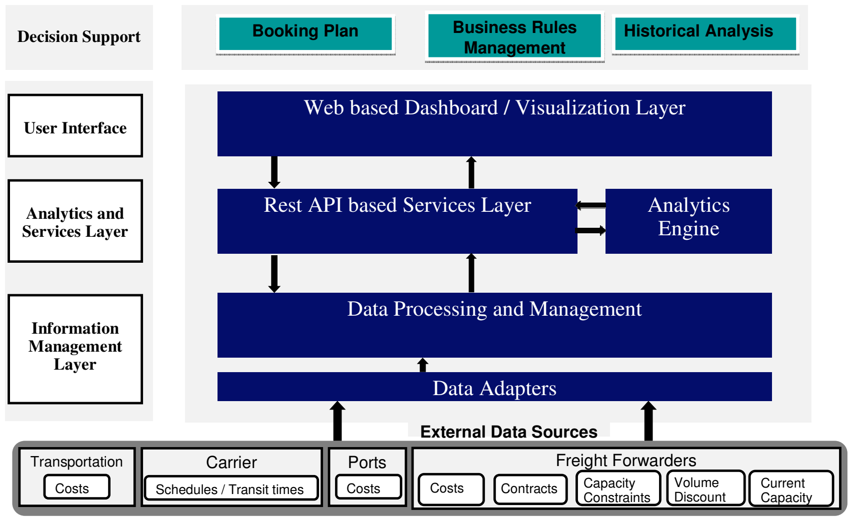

Figure 2: Deployment architecture of FR8NET.

2 ARCHITECTURE

FR8NET provides a platform that can support a set of data-driven analytics and allows them to function at scale in an operational environment. Figure 2 shows the deployment architecture of FR8NET.

FR8NET is a multi-layered architecture comprising of: (i) Information management layer, (ii) Analytics and services layer, and (iii) User interface. The Information management layer enables ingestion and processing of data. It comprises of data adapters and data processing engine. Given a varied set of data sources, this layer has data adapters to import data from these sources. These sources could be an API, a relational database or any other external feed. Regardless of the source, the data adapter layer abstracts the ingestion of data from varied sources. The processing layer consists of: (i) Data validation engine, (ii) Data transformation engine, and (iii) Data cleansing engine. Data validation engine enables qualitative and quantitative evaluation of incoming data. This is critical to analytics models since the quality of analytical results is a function of selecting only the validated set of data. The validation engine applies validation rules to validate the quality of the data and to check for missing values etc. Data transformation engine enables converting the raw incoming data into a standard data format defined by FR8NET. Data cleansing engine applies data cleansing rules to remove bad / missing data or outliers.

The Analytics and services layer consists of: (i) Common data model, and (ii) Analytics engines. The common data model has comprehensive set of tables and views to support analytics and visualization. The analytics engines layer enables plugging in various set of analytics models using the plug and play architecture.

The user interface layer consists of a set of predefined portals defined to render the analytics output to the user. The predefined portals are: (i) Spatial visualization charts, and (ii) Business intelligence reports.

3 BOOKING DECISION SUPPORT

The booking decision support tool was developed to support booking agents who have to respond with viable options to a request for freight transportation from a client. The tool takes as input a set of requirements that describe the transportation request and outputs a set of alternative routes.

The current version of the tool was designed for ocean transportation of full container loads. It accommodates up to three ocean legs. One of those legs is usually an intercontinental leg on a large vessel from a major port in one continent (e.g., Shanghai in Asia) to a major port in another continent (e.g., Rotterdam in Europe). The other two legs usually consist of smaller vessels going from a smaller port to a major port (or vice-versa) in the same continent (these are known as feeder legs). In addition to the ocean legs, a route may contain up to two truck legs. Truck legs are needed when the route includes transportation from an inland origin to a port and from a port to an inland destination. The different combinations of ocean and truck legs provide the tool with the capability to recommend routes for the following service types: Port to Port, Port to Door, Door to Port, and Door to Door.

A route includes a set of transportation legs. Each leg is described by its origin, its destination, the type of transportation, the type(s) of container(s) allowed, and time information. The time information available depends on the type of transportation. In the case of the ocean legs, the specific schedules consisting of departure date and arrival date are available. In the case of truck legs only estimates of the travel times are available. The time information for all legs in a route is combined with dwell times at ports in order to compute an estimated departure time from the route's origin, an estimated arrival time at the route's destination, and an estimated transit time for the whole route.

The cost of a route is the sum of the transportation rates for each leg and additional charges such as terminal handling charges, war risk charges, etc. Both the transportation rates and the additional charges may depend on the type of container to be used in the shipment. Therefore, the type of container (e.g., 20 ft or 40 ft container) is one of the inputs to the booking tool. Some carriers offer volume discounts, which are typically applied based on the annual volume shipped by a client. In that case, the calculation of the cost for a particular shipment request depends on the number of containers already shipped by the client on that carrier during that year. The transportation rates might also depend on the commodity to be shipped and on the existence of specific contracts between the shipper and the carrier.

For a particular transportation request, the number of routes that can fulfill that request is limited by the constraints imposed in the request. The booking tool currently supports the following constraints:

- Time window - each route has to fit within a time window specified by the user.
- Total transit time - the estimated transit time of each route must be smaller than a maximum transit time specified by the user.
- Include port - each route has to go through a particular port specified by the user.
- Exclude port - each route must not go through a particular port specified by the user.
- Include carrier - the ocean transportation in each route must be provided by a particular carrier specified by the user.
- Exclude carrier - the ocean transportation in each route must not be provided by a particular carrier specified by the user.

The time window is specified by the user by providing the following dates (see Figure 3):

- Cargo Ready Date (CRD) - the date when the cargo is available for shipment.
- Earliest Ship Date (ESD) - the earliest date when the shipment can depart.
- Latest Ship Date (LSD) - the latest date when the shipment can depart.
- Earliest Delivery Date (EDD) - the earliest date when the shipment can arrive at the destination.
- Latest Delivery Date (LDD) - the latest date when the shipment can arrive at the destination.

Depending on the business needs, the user may provide only a subset of the above dates. For example, the user may provide only the Cargo Ready Date and the Latest Delivery Date. In this case, the routes generated by the booking tool must depart at or after the Cargo Ready Date and must arrive at or before the Latest Delivery Date. It should be noted that if both the Cargo Ready Date and the Earliest Ship Date

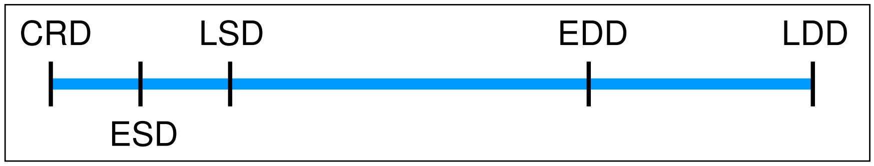

Figure 3: Dates used to define time window.

are provided by the user, the routes must depart at or after the latest of those two dates. If only one of them is provided, the routes must depart at or after that date.

Amongst the routes that satisfy the constraints of a transportation request there are usually some that are preferable than others from a business perspective. The booking tool currently includes three metrics for evaluating the routes:(i) Cost, (ii) Transit time, and (iii) Tier.

The estimation of the first two metrics (cost and transit time) was described above. The third metric, tier, classifies a route based on the ocean carrier used. The user may prefer certain carriers over others and therefore can attribute a higher tier level to the preferred carriers. The decision on the tier of each carrier depends on the business needs and can be based on many different aspects of the carrier. For example, it can be based on the payment terms provided by the carrier or the percentage of the time that the shipments on the carrier arrive on time. It can also be based on a combination of several aspects of the carrier.

The user of the booking tool chooses a criterion for selection of the best routes based on the three metrics available. The options are: (i) Minimum cost, (ii) Minimum transit time, and (iii) Higher tier carrier. Whatever the criterion selected by the user, the booking tool outputs the details of the three best routes. With this information the user can decide which route (or possibly which routes) to use for the shipment.

In Figure 4, we present a diagram of the booking tool. The tool connects to databases where all the data needed is stored. The tool also provides a graphical user interface where the user enters the information about the transportation request and where the output (i.e., the best routes found) is displayed. The central square in the diagram illustrates the sequence of steps in the tool to find the best routes. The route enumeration module corresponds to the construction of the feasible routes, the metric computation module corresponds to the estimation of the three metrics, and the optimization module corresponds to the selection of the best routes to present to the user. In the route enumeration module, an enumeration algorithm is used that basically constructs feasible routes one at a time by selecting transportation legs from the database that when put together satisfy all the constraints specified by the user. The output of this module is a set of feasible routes, i.e., a set of routes that satisfy all the constraints. In the metric computation module, the three metrics described above (cost, transit time, and tier) are computed for all the routes generated in the route enumeration module. Finally, in the optimization module the best routes are selected from the above set of feasible routes. In the current version of the tool, the optimization module consists simply of sorting the feasible routes according to the criterion selected by the user (see previous paragraph). For example, if the user selects the criterion of minimum cost routes, then the feasible routes are ordered according to increasing cost and the tool outputs the first three routes of the sorted list, i.e., the three cheapest routes in the list.

4 ANALYTICS SUPPORT

FR8NET provides analytics support for trade lane managers and product managers. There are mainly two analytics modules in FR8NET: Historical Analysis Module (HAM) and Business Rules Management Module (BRMM).

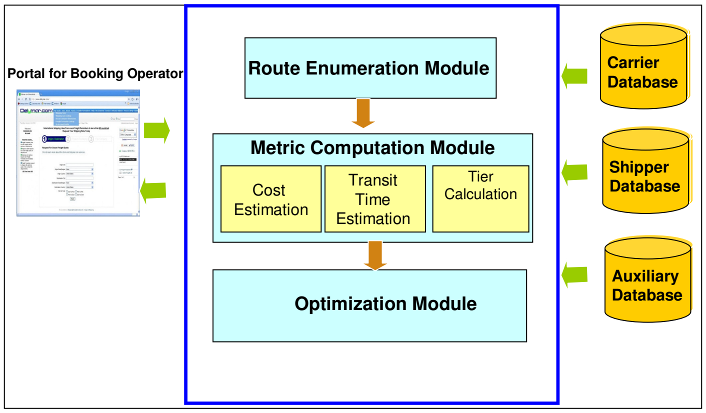

Figure 4: Diagram of the booking tool.

4.1 Historical Analysis Module

This module analyzes past transactions on multimodal shipping choices for the purpose of

- Identifying hot-spots in O-D pairs for different metrics;
- Identifying opportunities for possible cost-savings (missed opportunities);
- Understanding tradeoffs between transit time and cost which can be exploited in future negotiations with shippers/carriers;
- Identifying inconsistencies between local decisions and global business goals;

FR8NET provides a rich UI for HAM. We next describe three different use cases of HAM.

4.1.1 Cost-Savings Analysis

While it is rational to select cheapest carrier but often due to other constraints like limited procured capacity on cheaper carriers, longer transit times on cheaper carriers, or business rules governing carrier selection, booking operators end up selecting other costlier carriers. HAM can be used to quantify the cost-savings that could have been possible on historical transactions by the freight forwarder. For each past transaction, we select the cheapest carrier available at that time and use this to estimate any possible cost savings. Observe that, this is a theoretical bound as it assumes that there is always enough capacity available on the cheapest carrier and that the transit-delay on the cheapest carrier is acceptable to the customer. If the transit-time requirements are strict, the analysis can only restrict to selection of the cheapest carrier that also satisfies the transit-time and quantify any possible cost savings. The insights from this analysis can be helpful during capacity procurement on different carriers and when negotiating price with carriers.

4.1.2 Delay-Cost Tradeoffs Analysis

If the customers are tolerant of additional delay in transit-time, then freight forwarders can provide more value to customers by providing them cheaper carriers options. For each past transaction, we can identify the cheapest carrier whose transit-time is within x days more than the one selected and thus can quantify the total savings possible by exploiting the tradeoff between transit-time and cost.

4.1.3 Hot-Spot Analysis

HAM provides distribution of the total trade-lane traffic volume among different O-D pairs constituting the trade lane and among different carriers operating on the trade lane based on historical data. This information can be used for hot spot analysis which is aimed at identifying top O-D pairs or carriers for different route performance metrics, e.g., top O-D pairs by traffic volume, by cost contribution, or by potential cost-savings. Similarly, top carriers by traffic volume and by their cost contribution can be identified. Hot spot knowledge is useful during carrier capacity procurement when volume discounts can be negotiated with carriers over hot O-D pairs. Also this knowledge helps to speed-up turnover time of bids by filtering out insignificant O-D pairs and concentrating on optimized bids for hot spots.

4.2 Business Rules Management Module

Business transactions of freight forwarders, like any other firm are governed by business rules. These rules govern the selection of carriers and routes by the forwarders during booking and bidding. Business rules are typically set by trade-lane managers and/or product managers and are not changed for long periods. This is partly due to lack of predictive analytics to quantify the sensitivity of forwarder's revenue to different business rules and to changes in them. Traditional revenue management models revenue as a function of demand and price charged. However the price charged is dictated by the underlying business rules governing the choice of carriers or other service providers. Thus there is a need for analytics to clearly understand the repercussions of different business rules on the forwarder's revenue and performance. While much of the predictive analytics have traditionally been focused around market demand estimation, there is a lack of understanding of how different business rules effect the revenue. This knowledge can be exploited to dynamically change business rules with changing market conditions, demand and carrier prices for revenue optimization. Also this knowledge can be used for better negotiations with carriers during tender management.

For example, suppose a forwarder operates on two trade-lanes, T_1 and T_2. Carrier A offers local volume discounts on T_1 while Carrier B offers global volume discounts. Let the volume discount structure be as follows:

- *Carrier A*: 50% for any TEU beyond 100K TEUs on T_1,
- *Carrier B*: 40% for any TEU beyond 100K TEUs globally,

where TEU stands for Twenty-foot Equivalent Unit and is a standard unit of shipping capacity. 1 TEU is equivalent to storage capacity of a 20x8x8 ft container. Let C_{A,T_1} and C_{A,T_2} be the cost per TEU of carrier A on T_1 and T_2. Similarly define C_{B,T_1} and C_{B,T_2}. Let the demand be 120K TEUs on T_1 and 80K on T_2 and there be two business rules. Rule-1 says always go with Carrier A while Rule-2 says go with Carrier-B. We next calculate forwarder cost under two business rules.

- (Rule-1) $100K * C_{A,T_1} + 20K * C_{A,T_1} * 0.5 + 80K * C_{A,T_2} = 110K * C_{A,T_1} + 80K * C_{A,T_2}$.
- (Rule-2) $100K * C_{B,T_1} + 20K * C_{B,T_1} * 0.6 + 80K * C_{B,T_2} * 0.6$(assuming first 100K are for T_1) $= 112K * C_{B,T_1} + 48K * C_{B,T_2}$.

Depending on the values of $C_{A,T_1}, C_{A,T_2}, C_{B,T_1}, C_{B,T_2}$ either Rule-1 or Rule-2 results in less cost. In a scenario where these cost changes over time, the business rule choice should also change to save cost.

BRMM in FR8NET manages business rules governing booking choices by:

- Identifying business rules guiding operators to select costly/slower carriers, and
- Quantifying possible cost-savings/time-savings by simulating effect of changes in business rules on carrier selection.

In particular, it offers features to identify business rules resulting in reduced cost/time-savings by analyzing the business rules governing carrier choices in the past transactions. It also provides a simulation engine to simulate cost/time-savings possible by changing those business rules both on past transactions and on hypothetical future business scenarios. Simulation engine can also be used for sensitivity analysis of revenue to business rules. It provides an interactive UI for dynamically changing the business rules and simulating their effect on the cost/time-savings. Once a good candidate set of business rules have been identified, the user can trigger rule amendment in FR8NET.

The amendments to business rules can be due to changes in parameters of individual rules and/or changes in priority among rules in scenarios with more than one rule. BRMM interacts with the booking tool to generate feasible routes for a given booking request. From the feasible routes, we can select the best route under the given rule scenario. In case of multiple qualifying routes one can randomly select one of the routes. Average cost and transit-time under different rule scenarios can then be compared.

5 ANALYTICS VIEWS

The inputs to analytics view are the set of trade lanes, the customer type and the set of carriers. By specifying the appropriate set of inputs TLMs can restrict analytics domain. Once the analysis type (Historical Analysis or Business Rules Analysis) is specified, the analytics are executed over the input data and results are displayed. FR8NET offers several filters in the analysis view including filtering by Port of Lading (POL), Port of Discharge (POD), container type, time window etc. This can be used to drill down into the analytics results and identify performance bottlenecks.

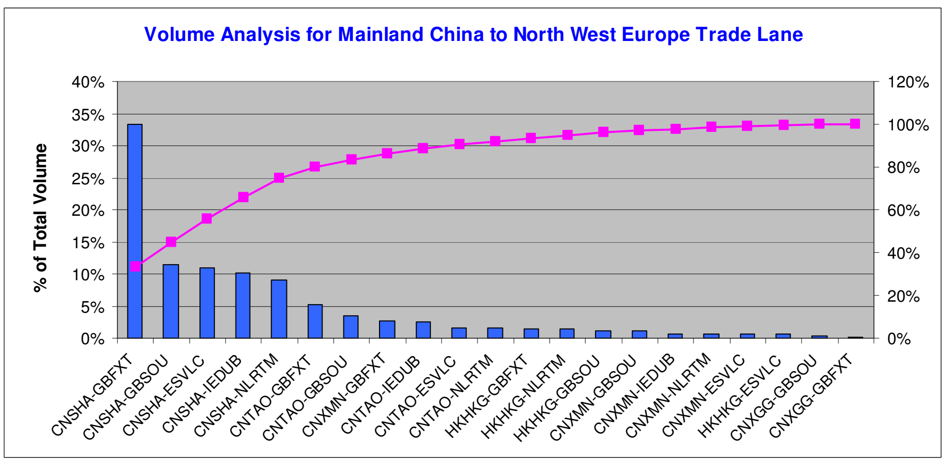

Figure 5: Traffic volume distribution over POL-POD pairs for a trade lane.

5.1 Historical Analysis View

The Historical analysis view provides two views: (i) Aggregate Statistics, and (ii) Potential Missed Opportunity. The Aggregate Statistics view displays the distribution of aggregate traffic volume over the POL-POD pairs in the selected trade lane which can be used for Hot-Spot analysis. Figure 5 shows a chart from this view showing the distribution of the total traffic volume on a trade lane along its constituent POL-POD port pairs. Out of some 20 POL-POD pairs, 15 carry less than 5% of the total traffic, with 80% of the traffic accounted by top five POL-POD pairs. The Potential Missed Opportunity view provides cost-savings analysis (Figure 6) and tradeoff analysis (Figure 7). Figure 6 shows the distribution of total cost savings on the trade lane over different POL-POD pairs when the cheapest carrier was selected for each booking, assuming carriers have unlimited capacity. Observe that about 42% of the total savings is possible if the cheapest carrier operating on one particular POL-POD pair has enough capacity. Comparing to Figure 5, it is interesting to see that POL-POD pair accounting for majority of the savings is not the one carrying the majority of the traffic. Figure 7 shows the cost savings possible when the customers are tolerant of additional delay of x days compared to the transit time of the fastest carrier on different POL-POD pairs. This chart shows that up to 16% cost savings are possible on this trade lane for $x = 8$.

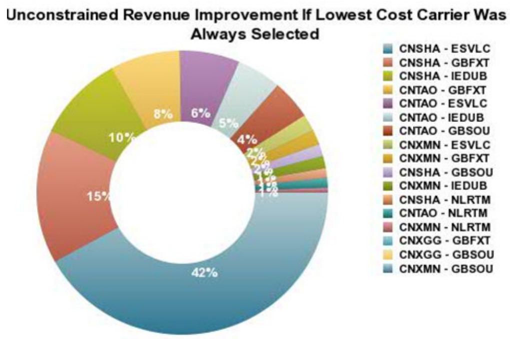

Figure 6: Cost-savings with cheaper carrier assuming no carrier capacity constraints.

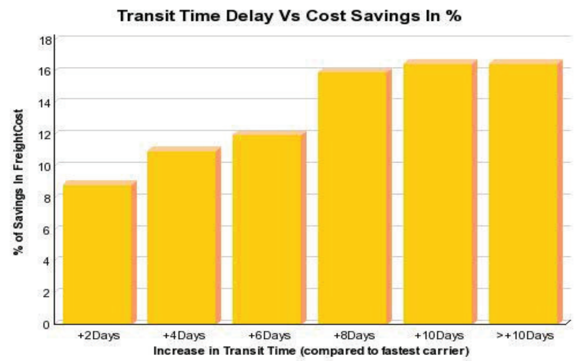

Figure 7: Cost-savings possible when customers are tolerant of additional transit times.

5.2 Business Rules View

This view provides interface to make changes in current business rules, specify load mix for hypothetical business scenarios, and simulate the effect of changes in business rules on the performance. A business scenario is specified by a booking request load mix and the set of business rules governing the booking choices. Currently the following business rules and their possible combinations are supported during route selection: minimum cost route, minimum transit-time route, route with higher tier carriers, include volume discounts offererd by carriers when calculating route cost, exclude specific carriers, and exclude specific ports. From the UI, the TLM can view the data associated with different rules. For example, current state tables corresponding to carrier tiers, volume discount parameters, available carrier capacity can be viewed and amended to simulate their effects.

We next describe the simulation engine used to simulate and compare different business scenarios. The main components of the simulation engine are demand specification, rules specification, feasible route generation, and optimization module. The demand specification component generates a load mix for simulation by specifying the aggregate load for each month and the percentage of total load carried by different POL-POD pairs. The load is in terms of a collection of booking requests for a given period of time. The load can be generated using a parametric model or from historical booking transaction data. The rules specification component is used to specify the set of business rules governing the business scenario.

The feasible route generation component basically consists of route enumeration module and metric calculation module of booking tool (see Figure 4). The simulation engine has a connector to the booking tool of FR8NET. Once the load mix is defined and the business rules are specified, the simulation engine calls the API of the route enumeration module for generating feasible routes for each booking request from the simulated load. The feasible routes generated at this step are not constrained by any business rules. Once the set of feasible routes are obtained, the metric computation module computes the cost, transit time and tier for each route. Next the optimization module is called along with the business rules (specified by the rules specification component) to filter feasible routes that are compliant to the simulated business scenario. From this subset of feasbile routes, one of the routes is selected based on some criteria, like minimum cost, minimum transit time, or higher tier. This is repeated for each booking request generated by the request geenration module. After each booking request is complete (i.e., is assigned a route along with the cost), the simulation completes.

After each simulation completes, the view gets populated with different charts showing the total and average shipment cost per TEU and average transit time under the simulated business scenario. Multiple business scenarios with different business rules can be simulated interactively and their results can be visually compared. Observe that, when simulating different business rule scenarios, the route enumeration module and metric computation module are called only once (for the first rule scenario) and the feasible set of routes are stored in a file or persisted in memory. For subsequent rule scenarios, the same file can be re-used as the feasible set of routes are not dependent on the specific rule scenario. Figure 8 shows business rules view charts for a hypotehtical scenario, with simulated load mix given by the pie chart and the cost metric under three different business rules, higher tier carrier, minimum cost route, and minimum transit time route. This functionality can be used to identify a candidate set of rules to accomplish business goals in future.

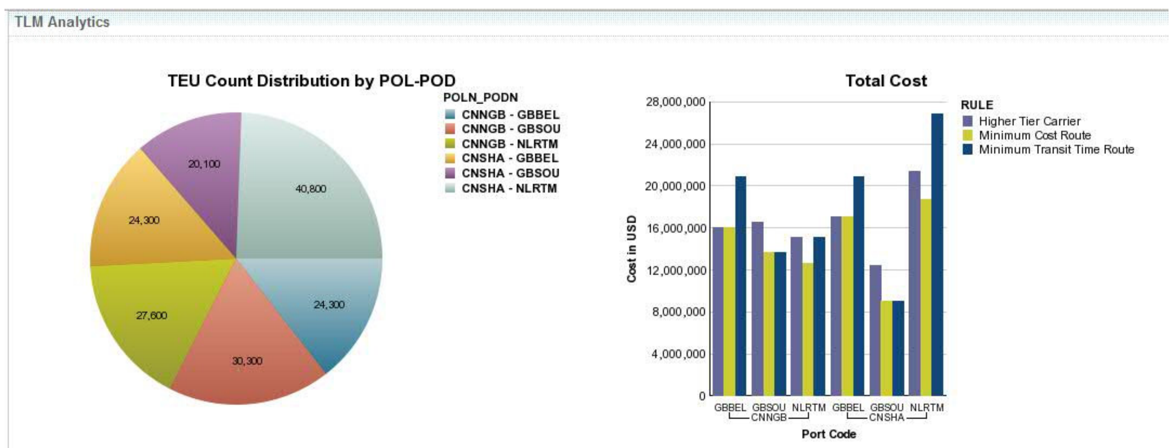

Figure 8: Cost comparison under different simulated business scenarios.

6 RELATED WORK

The development of efficient methods for the planning and management of multimodal freight transportation is a relevant and active area of research. This is explained in recent surveys and critical analyses of the current trends in multimodal decision support models (Caris, Macharis, and Janssens 2008; Crainic, Gendreau, and Potvin 2009; Caris, Macharis, and Janssens 2013; SteadieSeifi et al. 2014).

Simulation has been used in this context. Most of the literature on simulation applied to multimodal freight transportation systems seems to focus on the management of particular sub-systems of the transportation network. One example is the study of management of freight terminals (Arzani et al. 1996; Mazzucchelli, Recagno, and Sciutto 1997). There are, however, published studies more closely related to this paper which look at the whole multimodal transportation network. For example, Tavasszy, Smeenk, and Ruijgrok (1998) developed a decision support system that simulates the impact of policy measures on freight flows related to the Netherlands for a large number of products and transportation modes. A second example is given in Southworth and Peterson (2000) where the development of a model of a multimodal freight transportation network is described. The network was constructed to support the simulation of 5 million origin-to-destination freight shipments reported as part of the 1997 U.S. Commodity Flow Survey. Contrary to the above two examples, our simulation model focuses on the impact of specific business rules on the selection of the best routes for a set of shipments. Our study is an illustration of how simulation can be used in a support role for trade lane managers to help them select the best business rules that they can use to obtain the best outcomes in their business practice.

Among commercial offerings, IBM Operational Decision Management (ODM) Decision Validation Service (DVS) as described in Fu and Boyer (2014) provides a general framework for scenario-based simulation testing of business rules. However, it does not provide any plugin for multimodal freight transportation booking and analytics. Also it it not clear how to add external data sources to DVS.

7 CONCLUSION

Planning in the context of multimodal freight transportation is a complex task given the need for interaction and coordination between different entities. Probably due to its complexity, many of the planning tasks are still done manually and there is a need for automation and the development of efficient solutions.

This paper presents analytic tools developed at IBM Research to be used by trade lane managers. These tools aim at providing those managers with analytics capabilities to improve procurement, carrier selection, capacity planning, and business rules management. The analytic tools presented are based on simulation and consist of two modules: historical analysis and business rules management. The historical analysis module looks at past transactions and facilitates the identification of areas for operational improvement and cost savings. The business rules management module can look at the effect of business rules in past transactions and also simulate their effect on possible future business scenarios.

The tools presented in this paper were developed within a software architecture that is scalable and user friendly and therefore it is expected to lead to easy adoption by the freight transportation industry. These tools will empower decision support for end to end optimization and help improve the ability to plan better resulting in lower costs and higher margins.

In future, we intend to enhance the Business Rule Management module by adding functionality for sensitivity analysis of cost to business rules. Basically for each booking request, we can compare cost of final booking (after applying the business rule) vs. the best (for a given cost criteria, dollar amount, transit time etc) available (before applying the business rules) and the business rule(s) governing the final booking. By aggregating cost difference across different requests for each business rule we can identify the set of business rules degrading the aggregate cost-savings over all the requests. Having identified these candidate business rules, a further cost-benefit analysis can be done by accounting for the investment in changing the business rule and the resulting cost savings. This is important since changing business rules often involve some investment of time and money. The business rule(s) resulting in substantial Return on Investment (ROI) can be recommended for change to executives.

REFERENCES

Arzani, G., M. Mazzucchelli, R. Nurchi, V. Recagno, and G. Sciutto. 1996. "Towards cost effective freight terminal management: a simulation tool". In *Proceedings of the 8th European Simulation Symposium*, 638–642.

Caris, A., C. Macharis, and G. K. Janssens. 2008. "Planning problems in intermodal freight transport: accomplishments and prospects". *Transportation Planning and Technology* 31 (3): 277–302.

Caris, A., C. Macharis, and G. K. Janssens. 2013. "Decision support in intermodal transportation: a new research agenda". *Computers in Industry* 64 (2): 105–112.

Crainic, T. G., M. Gendreau, and J.-Y. Potvin. 2009. "Intelligent freight-transportation systems: assessment and the contribution of operations research". *Transportation Research Part C* 17 (6): 541–557.

Z. Fu and J. Boyer 2014. "Using IBM Operational Decision Manager DVS simulation features for risk scoring analysis use cases". IBM Business Process Management Journal. Available online at http://www.ibm.com/developerworks/bpm/bpmjournal/1404_fu/1404_fu.html.

Mazzucchelli, M., V. Recagno, and G. Sciutto. 1997. "Integrated freight terminal simulation tool: design and implementation". In *Proceedings of the 1997 Transportation Systems Conference*, Edited by M. Papageorgiou and A. Pouliezos, IFAC Symposia Series, 587–592.

Southworth, F., and B. E. Peterson. 2000. "Intermodal and international freight network modeling". *Transportation Research Part C* 8 (1): 147–166.

SteadieSeifi, M., N. P. Dellaert, W. Nuijten, T. V. Woensel, and R. Raoufi. 2014. "Multimodal freight transportation planning: a literature review". *European Journal of Operational Research* 233 (1): 1–15.

Tavasszy, L. A., B. Smeenk, and C. J. Ruijgrok. 1998. "A DSS for modelling logistic chains in freight transport policy analysis". *International Transactions in Operational Research* 5 (6): 447–459.

AUTHOR BIOGRAPHIES

PARIJAT DUBE is a Research Staff Member at IBM T. J. Watson Research Center. He received his M.S. in Electrical Communication Engg. from Indian Institute of Science, Bangalore in 2001 and his Ph.D. in Computer Science from University of Nice-Sophia Antipolis in 2002 where he was affiliated to INRIA. He joined IBM T. J. Watson Research Center in 2002. His research interests include performance analysis, control, and optimization of computer systems, distributed computing, operations research, and game theory. His email address is <pdube@us.ibm.com>.

JOÃO P. M. GONÇALVES is a Senior Software Engineer at IBM T. J. Watson Research Center. He specializes in computational optimization and is interested in the application of optimization techniques to the solution of business problems. He received a Ph.D. in Industrial Engineering from Lehigh University in 2005. His email address is <jpgoncal@us.ibm.com>.

SHILPA MAHATMA is a Senior Research Engineer at the Industries research at IBM Research. She has been the chief architect on a number of key IBM initiatives. Her work has received KDD award and has been runner up for Edelman award. She has architected the solutions in Healthcare and Smarter Planet/Cities initiatives involving the use of advanced analytics to provide customers with strategic and tactical business insight. Her email address is <mahatma@us.ibm.com>.

FRANCISCO BARAHONA is a Research Scientist at the IBM T. J. Watson Research Center specializing on optimization. He has over 20 years of experience in optimization projects dealing with airlines optimization, transportation planning, facility location and network design. Before joining IBM he was a professor at the University of Waterloo in Canada. He has a Ph.D. in Operations Research. Dr. Barahona has been an Associate Editor for SIAM Journal on Optimization, Operations Research and Management Science, and EURO Journal on Computational Optimization. His email address is <barahon@us.ibm.com>.

MILIND NAPHADE is a Manager and RSM at IBM T. J. Watson Reserch Center. He is currently the Program Director of IBM Research THINK Labs. He has a PhD from University of Illinois at Urbana-Champaign and is a Senior Member of IEEE. He has deep technical expertise in machine learning, big data analytics, signal processing, video analysis and retrieval. His email address is <naphade@us.ibm.com>.

MARK BEDEMAN is currently Global Freight Logistics & Cargo Subject Matter Expert (SME), Travel & Transport, Freight Logistics, IBM Global Business Services, and is responsible for aligning industry experience and knowledge with the full range of IBM capabilities. He has worked with Accenture, working in the Supply Chain practice as a Logistics Industry specialist covering, outsourcing and 4pl. Mark is, ex Chairman, Council on Global Supply Chain, Conference Board of Europe and ex Director, European Logistics Users, Providers and Enablers Group. ELUPEG. He was an Interim Director Danzas for 2 years, managing a number of Board Level Projects, prior to which he spent 23 years with Exel Logistics, formally National Freight Company (NFC), and held many senior business unit managing director and strategy roles in the UK, Europe and International. His email address is <bedeman@us.ibm.com>