

Selecting Telecommunication Carriers to Obtain Volume Discounts

Joris van de Klundert

Mateum, Postbus 616, 6200 MD Maastricht, The Netherlands, joris@mateum.nl

Jeroen Kuipers

Department of Mathematics, Maastricht University, Postbus 616, 6200 MD Maastricht, The Netherlands,
kuipers@math.unimaas.nl

Frits C. R. Spieksma

Department of Applied Economics, Katholieke Universiteit Leuven, Naamsestraat 69, B-3000 Leuven, Belgium,
frits.spieksma@econ.kuleuven.ac.be

Maarten Winkels

Department of Mathematics, Maastricht University, Postbus 616, 6200 MD Maastricht, The Netherlands,
m.winkels@math.unimaas.nl

During 2001 many European markets for mobile phones reached saturation. Hence, mobile phone operators have shifted their focus from growth and market share to cutting costs. One way of doing so is to reduce spending on international calls, which are routed via network operating companies (carriers). These carriers charge per call-minute for each destination and may use a discount on total business volume to price their services. We developed a software system that supports decisions on allocating destinations to carriers. The core of this system is a min-cost flow routine that is embedded in a branch-and-bound framework. Our system solves the operational problem to optimality and performs what-if analyses and sensitivity analyses. A major telecommunication services provider implemented the system, realizing two benefits: it has structured the business process of allocating carriers to destinations and cut the costs of routing international calls.

Key words: industries: communications; networks: graphs, flow algorithms.

History: This paper was refereed.

The telecommunications industry has changed dramatically over the last decade largely because of two parallel developments. First, within the European Union, the telecommunication markets have been privatized. Consequently, a monopoly market has become competitive. Second, on a global level, the advances in digital technology, in particular the introduction of mobile phones, have increased the variety of services available and financial turnover. Handling this increased complexity in the market has been challenging since the late 1990s.

New telecommunication service providers (telcos), with different business concepts have appeared. Many have used aggressive marketing techniques to gain market shares. The competition has been and still is fierce. In the Netherlands, for instance, the number of mobile operators (telecommunications

service providers offering mobile-phone-based services) increased from one, the previously state-owned KPN, to six in 2001. For some of their services, mobile operators use networks that are either operated (and owned) by the mobile operator itself or operated by another telco. Each mobile operator owns a radio network of radio-frequency antennas that communicate with the mobile phones and are connected to the wired net. All mobile calls are routed to their destinations via the wired net. Because most mobile operators do not own a wired net, they rely on other telcos' wired networks. We refer to such telcos as *carriers*.

Initially, many mobile operators' first priority was to gain market share. Since roughly 2001, however, in a mature market, many telcos have focused on their financial positions. Some have made enormous investments in networks and technology. For instance,

they have spent billions of euros on acquiring universal mobile telecommunications system (UMTS) frequencies in Great Britain and Germany. Many telcos therefore seek to improve their operational results and hence to reduce operational costs. Indeed, many telcos have set up companywide cost-reduction programs.

One way of reducing operational costs is to decrease spending on calls routed via carriers' wired networks. A major provider of telecommunication services asked us to work on its carrier-selection problem for international calls—calls that generate traffic on carrier networks across the Dutch border.

Since 2001, the market has remained turbulent, and some telcos and carriers have been more successful than others. Price is an important competitive weapon in this market, and low prices depend on low costs, which emphasizes the importance of low cost. We developed a model for reducing operational costs for international calls and implemented an effective solution method for it in a software system.

The Problem and Previous Work

Our client operates a countrywide mobile radio-frequency network, which is connected to another telco's wired network. Every international call is routed through this wired network. From the wired network, the calls are routed to their destinations (geographical entities, such as countries or cities, characterized by unique number sequences). The final goal of a call may be a phone connected to the wired network, or in the case of a mobile destination, the radio frequency network of our client or some other mobile operator. In the Netherlands, international calls are all routed to a single central hub where the wired national network is connected to the international telecommunication networks of the carriers. Examples of such carriers are Teleglobe, Canadian Overseas Telecommunications Corporation (COTC), Versatel, and France Telecom.

Carriers price their services by means of a price per call-minute for each of the destinations in their networks. To ensure that its customers have proper connections, our client has to select a carrier for each destination. For example, our client can select COTC as the current carrier for the calls to Mexico. Because of the network structure, our client does not need to

select a carrier for national calls. Carriers bill on an aggregated basis. At the end of the billing period, say each month, our client pays each carrier for each of the destinations—the price times the number of call-minutes sent via this carrier.

Our client wishes to route call-minutes to carriers so as to minimize the sum of the amounts of the carriers' end-of-year bills. It must, however, take two complicating factors into account.

First, having invested large amounts of money in developing the wired networks, carriers strive to use their capacity fully. To do so, they need estimates and insights concerning the amount of usage expected. In fact, carriers expect our client (and other telcos), to make statements and enter into agreements about how much network capacity it is going to use in the near future. The carrier then sets lower and upper bounds on the available capacity per time period. Thus, a contract with a carrier may contain agreements about lower bounds and upper bounds on the number of call-minutes our client will allocate to that carrier. In drawing up these agreements, our client should not rely on prices alone but also on forecasts of the number of call-minutes.

Second, some carriers offer volume discounts. They distinguish several volume intervals for which they give lower and upper thresholds in the number of call-minutes and set different prices. At the end of the year, the carriers determine the total number of call-minutes received over all destinations and the appropriate interval. Subsequently, they base their bills on the prices set for that interval. To minimize its total costs, our client has to take into account the appropriate end-of-year intervals. During the year, it is not known which interval our client will end up in, because this depends on an unknown number of future calls and the carriers it selects to route these calls. However, these decisions have a major impact on its profitability.

We developed optimization software to help our client to solve the problem, which we call the carrier-selection problem under volume discounts (CSPV). The CSPV is essentially a procurement-optimization problem. We can view the carriers as suppliers (or vendors) and the call-minutes for the different destinations as the products to be procured. We can find literature on related problems in procurement

optimization, an area that has been quite active recently. In the CSPV, suppliers use a joint business-volume discount. Further, the inventory costs or ordering costs do not play a role because call-minutes are not physical items and we do not need to order them or store them.

The CSPV deals with the tactical issue of from whom to order. Various authors have studied such vendor-selection problems (for instance, Xu et al. 2000). Katz et al. (1994) and Sadrian and Yoon (1994) discuss a procurement problem that is related to the CSPV. In their case, which arises in the context of purchasing telecommunication network hardware, the discount intervals are stated in terms of business volume, that is, dollar value. Thus, the vendor's discount depends on the total dollar value of purchases from that vendor, whereas in our case, it depends on the total number of units purchased from that vendor (carrier). Their problem also differs from the CSPV in another respect: they distinguish between purchases on a commitment basis and purchases on an as-ordered basis. They stress the importance of sourcing flexibility and wish to explicitly model the fact that they do not want to purchase all future items via committed contracts. Likewise, they explicitly consider the number of vendors for each item and consider constraints on the number of these vendors and the percentages of their requirements awarded to each of the vendors.

Degraeve and Roodhooft (2000) and Degraeve et al. (2000) view vendor-selection problems from the perspective of total cost of ownership. Following this approach, they use accounting techniques to identify all relevant costs associated with a vendor. Based on this information, they developed a mathematical program to decide which suppliers to use, when to order, and how much to order.

Rosenthal et al. (1995) consider the problem of vendor selection with bundling. *Bundling* refers to selling a bundle of items, where the price of (types of) items depends on other (types of) items in the bundle. Sarkis and Semple (1999) simplify and improve on their mixed integer linear programs (MILPs) which they solved using LINDO.

Crama et al. (2004) consider a problem in which the discount structure is exactly the same as that in the CSPV (they call it a total quantity discount). Their

problem is motivated by an application in which a company is buying raw materials for an industrial process. However, in their model, they deal with the additional complexity that forecasts concern outputs, whereas the prices are for inputs, and there are several ways to transform inputs into outputs. Thus, not only do they have a procurement problem, they also have a production problem, which makes their problem much more complex than the CSPV.

All of the authors mentioned above solved small- to medium-sized problems, concerning around 10 suppliers and less than 100 products, or in the case of Katz et al. (1994), 500 products. Our client stated explicit performance measurements regarding problem size that greatly exceeded these numbers. They wanted routinely to solve instances with 10,000 destinations (products), up to 15 carriers, and up to five discount intervals per carrier. Hence, we started the project with a feasibility study to determine the problem sizes that we could realistically solve in an acceptable amount of time and the corresponding technological requirements in terms of hardware and software.

What caused our client to look for optimization software for the CSPV was the state of affairs at the time, the year 2001. Carrier-switching software systems had become available. These systems can digitally assign a carrier to the appropriate destination without manual intervention, and our client had implemented such a system. Thus, it had become very easy for telcos to route calls to the right carriers once they knew the optimal selections. Our client needed a way to find the right carriers. Because of the market circumstances, our client wanted to solve the CSPV, and it had the IT tools installed to use a software system that would solve the CSPV.

Our client's yearly expenses for carriers exceed 10 million euros. Moreover, it had realized that its policy of assigning destinations to carriers that offered the lowest price per call-minute (assuming that the yearly total amount ends up in the highest interval of that carrier) could yield a nonoptimal solution, if at all feasible. Solving the CSPV, and solving it to optimality would contribute to profitability.

The number of destinations distinguished by the carriers was growing quickly (to well over 10,000), and carriers might use up to five price intervals. Selecting carriers manually was an enormous, if not

unmanageable, task. Worse, when the forecast changed or a carrier changed its prices, our client needed to reconsider its selections. It needed to automate the process to cope with these developments.

Problem Definition

You can gain a deeper understanding of the problem by studying a simple example (Table 1). In this example, we have two destinations and two carriers. Carrier B has a price per call-minute of 0.25 for Destination 1 and 0.35 for Destination 2, whereas Carrier A’s price depends on the total number of call minutes it handles: if this number is less than 1.5 million call-minutes, the prices are 0.40 for Destination 1 and 0.60 for Destination 2; otherwise they are 0.20 and 0.40. The forecast for both destinations is one million call-minutes.

One solution is to select Carrier B for both destinations, thereby routing 2,000,000 call-minutes via Carrier B. This results in a total cost of 600,000. Alternatively, one can select Carrier A for both destinations, which results in an equally high total cost of 600,000 because the combined forecasts of 2,000,000 exceed the threshold of 1,500,000 of Carrier A’s second price interval. We can find a more subtle solution by dividing the call-minutes for Destination 2 equally between the two carriers and choosing Carrier A for all call-minutes to Destination 1. This results in routing 1,500,000 call-minutes via Carrier A, yielding prices from the second interval, and routing 500,000 call-minutes via Carrier B. The total cost of this solution is 575,000.

Presumably, part of the example’s complexity stems from the fact that we had to divide the forecast for a single destination over two carriers to obtain an optimal solution. On the other hand, it is possible to

construct hard problem instances whose optimal solutions do not have that property.

The CSPV considers at a given time (the run date) the following inputs:

- (1) Actuals for each destination and for each carrier, the amount of call-minutes that the carrier has received for the destination since the start of the year;
- (2) Forecasts for each destination and for each month between the run date and the end of the year, the expected number of call-minutes;
- (3) Prices for each destination and for each interval of every carrier, the carrier’s current price for one call-minute of capacity for that destination;
- (4) Intervals for each carrier, the lower and upper thresholds of annual call minutes for which various prices are valid;
- (5) Monthly lower and upper bounds imposed by carriers on the total capacity available each month;
- (6) Penalty costs for various destinations for one call-minute of capacity of a carrier for that destination, which make it possible to forbid certain carrier-destination combinations and allow consideration of the carriers’ quality of service in the selection process.

These inputs are the parameters of an integer-programming formulation of the CSPV (appendix). The CSPV is to select one or more carriers for each of the destinations and to decide how many call-minutes to send via each of the selected carriers, while respecting the lower bounds and the upper bounds, so as to minimize the yearly procurement costs. The costs result from prices and penalties.

We formulated a straightforward MILP for the CSPV. A typical approach to solving the resulting model (1)–(9) would be based on solving the linear relaxation that arises when relaxing the integrality constraints. We would then use a branch-and-bound procedure to solve the problem to optimality. An alternative solution approach is based on a related, yet different approach. Because the integer variables are 0–1 variables used to model the fact that we choose exactly one interval for each carrier, the problem that remains after fixing an interval for each carrier is a min-cost flow problem. Specialized algorithms from network flow theory allow one to solve min-cost flow problems much faster than a standard linear-programming (LP) solver does (Ahuja

Destinations	Carrier A price		Carrier B price	Volume forecast
	Total volume interval [0–1,500,000)	Total volume interval [1,500,000–∞]		
Destination 1	0.40	0.20	0.25	1,000,000
Destination 2	0.60	0.40	0.35	1,000,000

Table 1: Example instance of the CSPV with two carriers, of which one uses two discount intervals.

et al. 1993). Hence, we can use an alternative solution approach for the model:

Step 1. Enumerate all possible ways of selecting a single interval for each carrier.

Step 2. For each of these ways, solve a min-cost flow problem, and store the solution found.

Step 3. Output the best solution.

Steps 1 and 2 may imply the solution of an enormous number of min-cost flow instances. To reduce the number of min-cost flow problems to be solved, we added an extra check at the beginning of Step 2 to see whether the combination of intervals was immediately infeasible (for example, because the sum of their upper bounds was smaller than the total number of call minutes). Still, enumeration of all feasible combinations of the integer variables may appear to be more time consuming than the branching process that results from solving the model by branch-and-bound. On the other hand, we can solve large-scale min-cost flow problems faster by using dedicated algorithms than by using standard LP solvers.

In a preliminary computational study, we started by solving instances of the min-cost flow problem by using state-of-the-art LP solvers as well as by using dedicated min-cost flow routines. The test instances were real-life instances of moderate size (five carriers, 10 months, and about 5,000 destinations). Apart from some initial problems we encountered with numerical stability (which resulted from scaling), dedicated min-cost flow routines turned out to be much faster on our test instances than state-of-the-art LP solvers. The latter typically took several minutes to solve the problem instances, whereas the min-cost flow routines took several seconds on an 800 MHz, 128 RAM, Pentium. A likely explanation for the difference in computation times lies in the size of the problems, which contained several hundreds of thousands of variables. Such a size brings out the added value of a dedicated solver versus a more general LP solver. Another explanation is that the double-scaling algorithm dealt much better with the small differences in the prices than the simplex-based LP approach. In addition, memory usage appeared to be more problematic for the LP-based approach than for the alternative.

We completed this preliminary study (including the modeling phase) in close collaboration with our

client's financial department within three months of the project's start. Based on the outcome of this feasibility study, our client decided to use the min-cost flow enumeration approach and to use IGSys'tems' CS2 code for solving the min-cost flow problem. CS2 is a double-scaling algorithm (Goldberg 1997) available at <http://www.cs2.com>. The computational results were good enough to develop software for optimizing actual carrier selection as well as for decision-support purposes. In particular, the CS2 code solves the problem quickly enough to perform scenario studies online (that is, while waiting for the results).

Discussion

Before and during our development of the software application solving the CSPV, a number of issues came up.

Frequency

The client's managers questioned how frequently to solve the CSPV. Important in addressing this question is the extent to which the input parameters are uncertain. For instance, if the (perceived) quality of the forecasts is high, we do not need to solve the CSPV as frequently as we would if the forecasts were poor. For some of the input parameters, the realization can be different from the input. Uncertainties are present in the monthly forecasts and in prices over the year. As time passes the number of call-minutes actually routed for a particular month may differ from the number of call-minutes forecasted for a particular destination. In addition, the forecast itself can change over the year. For example, the forecast made in January for the number of call-minutes to some destination in December typically differs from the forecast made in November. Currently, our client produces a monthly forecast for each destination.

Most carriers announce their prices for a three-month period in which the prices will remain fixed. Although these numbers do not usually change drastically, it is very hard to predict the exact prices after such a period. We estimate future prices by assuming that the current price will be valid until the end of the year.

Clearly, all the forecast and price data required for an optimal solution are not available until the year is over. However, the firm must make decisions

during the year; the CSPV is a real-time problem. We and the client managers agreed that the software system should satisfy the following requirements: (1) it should solve the CSPV at any time optimally under the assumption that forecasts are perfect and prices remain fixed, and (2) whenever we have updates for the actuals, the forecasts, or the prices, we would solve the CSPV problem again and implement the new solution. This means that we solve the problem at least monthly.

Short-Term Use Versus Long-Term Use

Right from the beginning, we intended the solution software for the CSPV model to serve two purposes. First, we designed the software to solve the short-term procurement problem, CSPV. Second, we designed the software to support long-term contract negotiations. The software enables users to easily perform several types of scenario studies. They can easily modify forecasts, interval lower and upper thresholds, lower and upper bounds, and prices and solve the problem again, without affecting the actual carrier selection.

Quality of Carriers' Destination Pairs

Our client measures the quality of the carriers' connections with respect to each of the destinations. From a customer-service viewpoint, it may consider a carrier undesirable for a certain destination because of the low quality of the connection it provides for that destination. Prior to our involvement, the firm used these measurements to exclude certain carriers for certain destinations. In developing our application, we added the possibility of turning the quality measurements into a price (the penalty) that can be added to the price the carrier gives for that destination. Our client is thus able to manage the trade-off between the carriers' quality and price for various destinations. Our client implemented an organizational procedure for determining these penalties to avoid individuals' opinions influencing carrier selection.

Brokers

Brokers are parties that offer short-term capacity for certain destinations at low rates. Brokers are expected to play an increasingly prominent role in the international telecommunication network industry. However, our client has not yet decided to purchase capacity from them partly because the quality of the

connections brokers offer is unreliable. (Brokers do not own their networks, and at the time of offers, they may not have decided the network from which to procure capacity.) Moreover, the impact of accepting broker offers on the expected numbers of routed call-minutes, and therefore on the end of year discount rates, is uncertain at the time of acceptance. Therefore, it is hard to determine the financial consequences of accepting broker deals. Consequently, while scoping the project, we decided not to take brokers into consideration.

Software Application

The software system's core is the algorithm we described. It is embedded in a user-friendly multi-threaded windows application that can communicate its input and output to other software. This application is named BeCR (best cost routing). Because most of its users are familiar with spreadsheet-based data processing, BeCR extracts its input on actuals, forecasts, and so forth from Excel files. Some of these Excel files are generated by other software applications; others are maintained and kept up to date by the person responsible for BeCR. BeCR's output can be viewed within BeCR but can also be exported to Excel. The communication with the carrier-selection software is based on XML and is technically more involved than just encoding the solution of the problem. The application is file based and does not require human activity other than basic file management. It runs on a normal PC with acceptable response times, so it required no investments in hardware.

Benefits

Our first achievement in the project was gaining managers' recognition that allocating carriers to destinations (the CSPV) is a nontrivial problem for which mathematical effort is needed for optimal decision making. Moreover, considering the annual business volume involved, the managers readily realized the value of optimal decision making. Eliminating sub-optimality in the current settings and processes was important, and they expected major benefits in the future as market conditions changed (as carriers used more intervals, and as the importance of negotiating based on upper and lower bounds increased). BeCR has realized these benefits, but it is hard to calculate

precisely how much money the firm saved. However, it estimated that BeCR saves at least one percent of the yearly costs for carriers. (In fact, the instances we used to compute these improvements indicate that the short-term carrier selections alone save at least one percent.) The payback time for the firm's investment in the new system was months rather than years.

The client can also analyze the consequences of changes in input parameters, such as monthly lower and upper bounds, improving its ability to negotiate with the carriers. Investigating the carriers' flexibility and using this flexibility also contributes to a more efficient process.

Our client has also realized operational benefits. First, the carrier-selection process, which used to be extremely time consuming, is now much more efficient. Second, by modeling and formalizing this process, our client has improved its data integrity and reliability and gained greater control of the process.

Conclusion

By implementing BeCR, our client has reduced its costs for carriers and improved its selection process.

Our client can now reap the benefits of financial-procurement optimization and pick the higher hanging fruits of contract negotiation and accurate forecasting. Many players in the turbulent mobile-operator market have shifted their focus to the quality of services offered to customers, and therefore the quality of carriers' services is now a major issue.

BeCR allows our client to deal with quality issues by excluding specific carrier-destination pairs or adding penalties to carriers' prices. It is using both options. BeCR serves as a starting point for subsequent models that deal with quality issues in ways suitable to their strategic priority. In fact, a key factor that will determine the firm's use of BeCR in the future is whether its output will continue to provide enough insight into the trade-off between expenses and quality of service.

Appendix

Integer Programming Formulation for the CSPV

We define continuous decision variables as follows:

x_{ijkt} = number of call-minutes sent via carrier i to destination j in month t for the price in interval k .

Further, we define the binary decision variable y_{ik} to be

$$y_{ik} = \begin{cases} 1 & \text{if the total number of call-minutes routed} \\ & \text{via carrier } i \text{ falls in interval } k, \\ 0 & \text{otherwise.} \end{cases}$$

Further, we define the following parameters:

p_{ijkt} = price per call-minute for destination j in interval k of carrier i in month t .

d_{jt} = forecast for destination j in month t .

LT_{ik} = lower threshold of interval k of carrier i .

UT_{ik} = upper threshold of interval k of carrier i .

LB_{it} = lower bound for the number of call-minutes routed via carrier i in month t .

UB_{it} = upper bound for the number of call-minutes routed via carrier i in month t .

Now, the model can be defined as

$$\text{Min } \sum_{ijkt} p_{ijkt} x_{ijkt} \quad (1)$$

$$\text{s.t. } \sum_{ik} x_{ijkt} = d_{jt} \quad \forall j, t, \quad (2)$$

$$\sum_k y_{ik} = 1 \quad \forall i, \quad (3)$$

$$\sum_{jt} x_{ijkt} \geq LT_{ik} y_{ik} \quad \forall i, k, \quad (4)$$

$$\sum_{jt} x_{ijkt} \leq UT_{ik} y_{ik} \quad \forall i, k, \quad (5)$$

$$\sum_{jk} x_{ijkt} \geq LB_{it} \quad \forall i, t, \quad (6)$$

$$\sum_{jk} x_{ijkt} \leq UB_{it} \quad \forall i, t, \quad (7)$$

$$x_{ijkt} \geq 0 \quad \forall i, j, t, \quad (8)$$

$$y_{ik} \in \{0, 1\} \quad \forall i, k. \quad (9)$$

Constraints (2) ensure that all forecasted call-minutes are routed via some carrier. Constraints (3) and (9) ensure that exactly one interval is selected for each carrier, while constraints (4) and (5) ensure that the prices corresponding to the selected interval are the prices paid. Finally, constraints (6) and (7) model the requirements, that in each month and for each carrier, the total number of call-minutes is not below the given lower bound and does not exceed the

given upper bound. Constraints (8) are nonnegativity constraints, and constraints (9) are the integrality constraints. The correctness of the model is not hard to verify.

You can see that the formulation uses a price for each destination for each interval of a carrier. Researchers usually assume that a percentage is involved, i.e., the quotient of the price in the k th interval and the price in the first interval is identical for all products. In the instances that we encountered, this was not the case. Thus, instead of having a set of prices as input for each carrier and a percentage for each interval, we have a set of prices for each interval of each carrier. Further, the model (1)–(9) does not explicitly take the actuals into account. However, we can easily modify the thresholds LT_{ik} and UT_{ik} ($\forall i, k$) to account for the call-minutes that have already been routed. Finally, in the presence of penalties, let a parameter pen_{ijt} denote the penalty per call-minute for destination j of carrier i in month t . Then, we can modify the objective function (1) to

$$\text{Min } \sum_{ijkt} (p_{ijkt} + pen_{ijt}) x_{ijkt}.$$

Min-Cost Flow Model

When given an interval for each carrier, the CSPN becomes a min-cost flow problem. In the model (1)–(9), assume that values for the y_{ik} variables are given. Then, we can derive the following formulation (the index k has disappeared because the intervals have been specified):

$$\text{Min } \sum_{ijt} p_{ijt} x_{ijt} \quad (10)$$

$$\text{s.t. } \sum_i x_{ijt} = d_{jt} \quad \forall j, t, \quad (11)$$

$$\sum_{jt} x_{ijt} \geq LT_i \quad \forall i, \quad (12)$$

$$\sum_{jt} x_{ijt} \leq UT_i \quad \forall i, \quad (13)$$

$$\sum_j x_{ijt} \geq LB_{it} \quad \forall i, t, \quad (14)$$

$$\sum_j x_{ijt} \leq UB_{it} \quad \forall i, t, \quad (15)$$

$$x_{ijt} \geq 0 \quad \forall i, j, t. \quad (16)$$

To show how an instance of (10)–(16) gives rise to a min-cost flow instance, we construct a graph with

four sets of nodes: a node for each pair consisting of a destination j and a month t (called dm nodes), a node for each pair consisting of a carrier i and a month t (called cm nodes), a node for each carrier i (called carrier nodes), and finally a (single) source node s . The demand of the source node s equals $-\sum_{jt} d_{jt}$, the demand of a dm node equals $d_{jt} \forall j, t$, and all the other demands are 0. Arcs run from s to each carrier node i . With each such arc, we associate a lower bound LT_i and an upper bound UT_i ; these arcs have cost 0. Further, an arc runs from each carrier-node i to each cm node (i, t) . These arcs have lower bound LB_{it} and upper bound UB_{it} ; and also have cost 0. Finally, an arc runs from each cm node (i, t) to each dm node (j, t) that corresponds to the same month. These arcs have a lower bound of 0 and infinite upper bounds; the cost of the corresponding arc equals p_{ijt} . You can verify that a feasible flow corresponds to a feasible carrier selection and vice versa.

Acknowledgments

We sincerely thank all employees from our client's organization who were involved in developing and implementing BeCR. Also, we thank the reviewers for their insightful comments, which helped to improve the presentation of this paper.

References

- Ahuja, Ravindra K., Thomas L. Magnanti, James B. Orlin. 1993. *Network Flows*. Prentice Hall, Englewood Cliffs, NJ.
- Crama, Yves, R. Pascual, A. Torres. 2004. Optimal procurement decisions in the presence of total quantity discounts and alternative product recipes. *Eur. J. Oper. Res.* **159** 364–378.
- Degraeve, Zeger, Filip Roodhooft. 2000. A mathematical programming approach for procurement using activity based costing. *J. Bus. Finance Accounting* **27**(1–2) 69–98.
- Degraeve, Zeger, Eva Labro, Filip Roodhooft. 2000. An evaluation of vendor selection models from a total cost of ownership perspective. *Eur. J. Oper. Res.* **125**(1) 34–58.
- Goldberg, A. V. 1997. An efficient implementation of a scaling minimum-cost flow algorithm. *J. Algorithms* **22**(1) 1–29.
- Katz, Paul, Amir A. Sadrian, Patrick Tendick. 1994. Telephone companies analyze price quotations with Bellcore's PDSS software. *Interfaces* **24**(1) 50–63.
- Rosenthal, Edward C., James L. Zydiak, Sohail S. Chaudry. 1995. Vendor selection with bundling. *Decision Sci.* **26**(1) 35–48.
- Sadrian, Amir A., Yong S. Yoon. 1994. A procurement decision support system in business volume discount environments. *Oper. Res.* **42**(1) 14–23.
- Sarkis, Joseph, John H. Semple. 1999. Vendor selection with bundling: A comment. *Decision Sci.* **42**(1) 14–23.

Xu, Jiefeng, Leonard L. Lu, Fred Glover. 2000. The deterministic multi-item dynamic lot size problem with joint business volume discount. *Ann. Oper. Res.* **96**(1–4) 317–337.

The editor has received a confidential letter from an officer of the client's firm stating the following: "I hereby declare that the software application BeCR, as described in the paper 'Selecting Telecommunication Carriers to Obtain Volume Discounts' is in actual use since Q3 2001. Since then it has been used on

a monthly or bi-monthly basis for almost all carrier selection activities.

"BeCR has helped us tremendously in streamlining our carrier selection business processes and reducing direct purchasing costs. We estimate that cost savings are in the order of 1% per year, over the last 3 years. In view of the very short pay back period, it has been a very profitable project. Moreover, BeCR is also an enabler for improving our processes further so as to support current business priorities."