




INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Optimal Routing and Assignment of Consultants for Energy Education, Inc.

Junfang Yu, Randy Hoff,

To cite this article:

Junfang Yu, Randy Hoff, (2013) Optimal Routing and Assignment of Consultants for Energy Education, Inc.. INFORMS Journal on Applied Analytics 43(2):142-151. <https://doi.org/10.1287/inte.1120.0656>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2013, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Optimal Routing and Assignment of Consultants for Energy Education, Inc.

Junfang Yu

Department of Engineering Management, Information and Systems, Southern Methodist University,
Dallas, Texas 75205, yuj@lyle.smu.edu

Randy Hoff

Energy Education, Inc., Dallas, Texas 75225, rhoff@energyed.com

Energy Education, Inc. (EEI), a US management consulting firm, specializes in implementing energy conservation programs for schools, universities, and large churches. Similar to many consulting firms, travel expenses are among its largest budget items. Managing consultant travel for minimum cost and in a manner that meets client needs is critical. Typically, a subject matter expert at the company produces a consultant routing and assignment schedule using a labor-intensive, time-consuming, manual process; the schedule produced is usually far from optimal. The objective of our research is to minimize the total cost of consultant travel and staffing. Our models use a cluster-first, route-second methodology. We developed a set-covering binary integer programming heuristic to cluster clients based on geographic location. The relaxed consultant routing and assignment problem is formulated as a mixed-integer linear programming model using cluster locations and demand with consultant skills and availability.

In a recent 12-week period, the results of our research reduced EEI costs by 24 percent and provided several qualitative benefits. We conducted sensitivity analysis to provide EEI with improved decision analytics for additional modification of its existing processes and business routines.

Key words: vehicle routing; employee assignment; cluster analysis; consulting.

History: This paper was refereed. Published online in *Articles in Advance* January 14, 2013.

Energy Education, Inc. (EEI) is a 25-year-old consulting firm that focuses solely on the design, implementation, and service of energy conservation programs for educational institutions and large churches. It uses a high-touch model (i.e., consultants typically interact with clients face to face). Consultants implement nearly 100 percent of their engagements at the client site, frequently during unoccupied hours (i.e., nights, weekends, holidays, extended breaks). EEI consultants interact personally with all levels of client employees, using specialized implementation skills that cannot be easily replicated or for which other consultants cannot be easily cross-trained. Therefore, consultants travel to almost all client assignments and incur significant travel expenses for flights, hotels, rental cars, and meals. In 2010, each consultant spent an average of \$1,000 per week in flight costs.

Managing consultant travel at minimum cost and in a manner that meets client needs is one of EEI's most important objectives. Our research enabled the

firm to reduce its travel costs by 24 percent and focus on specific areas for additional future cost savings. This paper reports on the process, methodology, and insights gained from improving EEI's consultant routing and assignment process.

The Problem and Study Purpose

Consulting firms desire a solution to their consultant routing and assignment problem to help them make dramatic improvements to these key metrics; however, no current methodology or repeatable process that a firm can apply directly is available. In this paper, we discuss our development of new algorithms for assigning and routing traveling consultants. Our solution incorporates integer programming models and heuristic techniques. In using our model to create a weekly schedule, we incorporate guidelines concerning supply, demand, capacity, flow balance, safety, employee quality of life, and feasible travel schedules. The model's objective is to minimize the total cost of employee travel and the variable expenses of the

consultant population. Although routing and assignment problems have been studied and solved using a variety of methods, EEI's problem has several unique requirements because of the distributed nature of the clients and consultants, variation in client needs, and combinations of skill sets available.

Manual Approach to Consultant Scheduling

At EEI, scheduling was a highly manual process that a subject matter expert (SME) at the company's headquarters completed. SMEs used their knowledge of client and consultant locations and their understanding of reasonable travel schedules to allocate consultant supply to client demand. This manual scheduling process used a greedy method of visit allocation. It was a labor-intensive process that required significant experience and more than 16 hours of dedicated effort each week. Therefore, cross-training EEI staff was difficult. Figure 1 presents the steps in the high-level decision process EEI used in its manual consultant scheduling.

Step 1. Gather all the new information for the current scheduling period, including data on consultant supply. The SME determined consultant supply by the amount of time each person is willing and able to work, excluding days for personal leave, holidays, training, or other company meetings. Client demand was guided by a separate process that factored in client size, age, performance level, and special requests. Consultant supply must equal or exceed client demand. If demand exceeded supply, a management review was required to artificially reduce priorities for client visits. Fixed visits were also included for clients who requested a specific consultant with a special skill during a given week.

Step 2. Assign fixed visits to each consultant's travel calendar. The location of a fixed visit becomes an anchor point for the consultant, who will likely be scheduled to visit other clients in the same location to save travel costs. Thus, this simple process of

assigning fixed visits gave anchor points for a few consultants. Once this step was completed, the SME indexed consultant availability and reduced it on an individual level.

Step 3. Assign consultants to home area visits. The SME manually matched consultants who have the required skills to clients in their home area. These visits support client needs at the lowest cost because consultants do not require air flights and usually do not require hotel accommodations. After assigning these visits, the SME indexed and reduced consultant availability. Following this step, some consultants have no capacity (i.e., they have no available time left).

Step 4. Allocate consultants with remaining capacity for available visits near their assigned fixed visits. The logic was that because a consultant had to travel to a distant area, the sunk cost of the airline ticket might cover other visits in that vicinity. Once again, the SME indexed and reduced consultant availability and removed consultants with no remaining capacity from consideration for this schedule.

Step 5. Manually match all remaining visit demands with the remaining consultant supply. Although this step attempted to minimize travel difficulty for each consultant, it was not mathematically optimized in any way. To meet demand, consultants often were required to take one or more flights per day.

Step 6. Review calendar. Management and peers at EEI reviewed the entire calendar to verify the validity of the completed consultant schedule set. If any member of the review team felt that a travel schedule was infeasible, the review team asked the SME to rework the flow. This often had a cascading effect because visits were shuffled among consultants.

A New Approach: Cluster First, Route Second

Based on our research, we defined a two-stage process model (see Figure 2) that automates and optimizes the process of routing consultants in a manner that




Figure 1: The diagram shows the six steps of the manual scheduling process for consultants.




Figure 2: The consultant assignment and routing heuristic has two stages. The flowchart shows the input to each stage and the output from that stage.

(1) minimizes the number of flights and consultants required by EEI each week, and (2) fully meets client demand. In Stage 1, we establish clusters of clients, using preferred airports as centroids. The clustering methodology uses a set-covering integer linear programming model to group and aggregate client demand by skill type. We assign clients to clusters in a mutually exclusive fashion that ensures no overlap in support area. In Stage 2, we use a mixed-integer linear programming (MILP) model to create weekly assignments and routes for each consultant.

The consultant scheduling problem requires decisions regarding the assignment of consultants in support of client cluster demand. We constrain the categories into three primary groups: cluster demand, consultant availability, and consultant travel flow. Cluster demand requires all required visits, including fixed visits by a specific consultant, to be filled during the scheduling period. Consultant availability requires conditions such that (1) consultants can only fill demand if they are proficient in the necessary skills, and (2) consultants cannot be assigned more visits that they are capable of filling. Finally, consultant travel constraints are used to organize the flow of visits such that consultants start their travel from the farthest point away from home and work toward home during each scheduling period in the most efficient manner. The model minimizes the total direct cost for all consultants within the scheduling week, including travel from and to consultants' homes and among clusters.

Clustering Clients

Many routing and assignment problems such as the traveling salesman problem, are NP-hard (i.e., nondeterministic, polynomial-time hard) with solution times that increase exponentially as additional variables are introduced. For this reason, clustering a set of data into homogeneous groups can greatly simplify a model's solution and reduce its computing time. Using clustering, we reduce the mathematical model from a huge number of clients to a small number of clusters, thus sharply reducing the problem size. Many techniques have been used to cluster clients for routing problems. One promising heuristic technique (Mastrogiannis et al. 2011, Toregas et al. 1971) formulated an integer linear program set-cover model as the basis for clustering. This technique used an m -row, n -column, zero-one matrix, with specific costs assigned to each column. The clustering heuristic defined a method for partitioning the columns into homogeneous clusters and gave a rule for selecting the best column for the each cluster. The researchers (i.e., authors) defined an iterative process for selecting the prime cover for each row attribute.

We use client clustering based on geographic location as our approach to relax the problem structure and reduce the number of variables and constraints needed. Our model uses a set-cover binary integer programming heuristic. The goal of our analysis was to find the minimum number of airportcentric clusters that will cover all clients so that we could aggregate client demand. Appendix A shows the details of the model for clustering clients.

Finally, we conduct a manual review of each client to assign a best fit. In EEI's process, this step required the SME to choose the best natural cluster for a client who could potentially be assigned to multiple groups. The natural cluster considers nearby towns and cities, or physical barriers such as rivers, mountains, and interstate highways. Our model requires assigning clients to mutually exclusive cluster groups.

Routing and Assigning Consultants to Clusters

The process of routing and assigning consultants to clusters is similar to a vehicle routing problem (VRP), which provides a solution method for distributing goods from a depot to multiple customers through a known set of routes. The solution of a VRP is a minimum-cost routing schedule that fulfills customer

demand and satisfies all other constraints. Beck et al. (2003), Cheng and Wang (2009), Bard et al. (2002), Toth and Vigo (2002), and Diaz (2005) researched vehicle routing or documented the similarities between vehicle routing and various scheduling applications. Their research identified specific areas of overlap that contribute to the performance of solution methodologies. The multiple-depot vehicle routing problem (MDVRP) is a specific and relevant type of VRP. In this method, a company may have several depots from which it can serve customers. Because these depots and customers are commonly intermingled, a solution that uses segmentation and multiple models is not feasible. The objective of the MDVRP is to service all customers while minimizing the number of vehicles and travel distance. The main difference between our consultant routing and assignment problem and an MDVRP is the skill requirement of each visit and the unique set of skills specific to each consultant. Therefore, our model has additional constraints to address these differences.

We formulate the consultant routing and assignment problem as an MILP model using the clustered clients, as we describe above. In our formulation, each cluster represents a group of clients; these clusters have a specific demand for a number of skills. Each consultant has a specific skill set and availability. Consultants must fully satisfy all cluster demands by starting at their home location, traveling to various clusters, and then returning home at the end of the period. A small percentage of EEI consultants are independent consultants. EEI hires these contract consultants as needed on weekly basis. The model's objective is to minimize the total flight cost for all consultants and the variable labor cost for all assigned contract consultants in a scheduling week. Appendix B shows the model formulation for routing and assigning consultants to clusters.

Implementation Methodology

The SME role is required to plan the consulting demand for each applicable client, including numbers of visits (shifts), type of visits (skills), and specific consultants (fixed visits). The SME also maintains each consultant's skill-set matrix and captures consultant availability for the planning week in terms of number of available shifts. The client cluster is formed using a set-cover model. The client clustering structure does not change each week; it changes only if the mix or location of a considerable number of clients changes.

Based on the client clusters obtained (see Figure 3), the client-based consulting demands are converted to cluster-based demands. The flight cost for each possible origin-destination pair is adjusted based on the airfares during the planning week, as does the labor cost for each potential contract consultant.

We prepare the demand-and-cost data and other supporting data in a format needed by our solver interface (note that we use CPLEX) and solve it within one hour. For this implementation, we did not develop any specific software for the project. However, we have plans to develop software to automate the processes listed above. For each consultant, the output includes a weekly assignment and a travel route that starts at the consultant's home and has a list of consecutive clusters to be visited. The assignment and route include the number of visits (shifts) for each type (skill) and cluster.

For each cluster, the cluster-based consultant assignments are allocated back to client based within the cluster. Because no additional cost is considered within a cluster, any feasible allocation will serve the purpose. The allocation process always begins with fixed visits, if any are present. The SME reviews the final assignment schedule and, for specific business reasons, may apply manual adjustments that do not




Figure 3: The graphic displays the number of clients in each cluster in descending order.

incur additional cost and are within a predefined tolerance. The assignment schedule is posted on each consultant's calendar.

The SME's use of mathematical programming is still a manual process because we are waiting for funding to allow us to fully implement and automate the scheduling. We summarize the steps that the SME performs in developing the weekly schedule.

- (1) Collect weekly demand data for each client and each visit type, including fixed visits.
- (2) Aggregate client-based demand to cluster-based demand for both fixed and nonfixed visits.
- (3) Collect the availability data for all consultants.
- (4) Update the cost data for flights and contract consultants.
- (5) Manually run CPLEX to generate the cluster-based schedule, including the routing information, among the clusters.
- (6) Manually disaggregate the cluster-based schedule to the client-based schedule for each consultant.
- (7) Publish each consultant's schedule.

Model Input

We collected data over a scheduling period starting with the week of February 7, 2011 (identified as week 1) and ending with the week of May 16, 2011 (identified as week 15); this equates to 15 scheduling weeks. We removed three weeks (week 2, week 11, and week 14) from consideration, because we considered them atypical because of artificially low consultant supply numbers. The model significantly improved cost and performance for these three weeks; however, we could not extrapolate the results over a longer period. Final data analysis yielded a nonconsecutive 12-week period that we used as the basis for all analyses and comparisons.

The typical scheduling week averaged 229 visits required by clusters. Figure 4 shows the variation in demand by week.

The consultant supply was stable and exceeded demand each week. Many consultants are employed as independent contractors to EEI; as needed, they provide scheduling flexibility. These independent contractors are paid on a weekly basis. If they are asked to make a single visit in a given week, they receive their full compensation and are treated as variable costs. Therefore, balancing capacity and the use of independent contractors is critical. Consultant capacity approached 300 shifts per week in each scheduling period.

Fixed visits, which occur when a client demands a specific consultant to fill a need, are an important demand type to monitor. Increasing fixed visits reduces scheduling flexibility and increases cost. Although the average demand for fixed visits is nearly 45 per week, these visit requests are highly variable and based on specific needs. Figure 5 shows the number of fixed visits demanded by clusters each week.

Results and Comparison to Actual Schedules

We used the routing and assignment model for all 12 weeks in the scheduling period. Because of the model's size and complexity, we could not achieve optimality in any model run. To obtain our solutions, we used a CPLEX version 12 solver running on a Linux server with a dual six core Intel Xeon 3.4 GHz processor and 96 GB memory. Figure 6 illustrates a typical model run. We scheduled the model to run for 20 hours. The objective function improved significantly in the first several hundred seconds of run time. At the one-hour (3,600-second) mark, it was




Figure 4: Weekly cluster demand varies from week to week.




Figure 5: The graph displays the demand for fixed visits during each week that we modeled.




Figure 6: The results from week 3 demonstrate how the solution approaches its optimality over time.

\$19,106; at the end of the 20 hours, it was \$18,952—only 0.8 percent less. Therefore, the model made no improvement for the last 16 hours of run time. After comparing the trade-off between run time and objective function improvement, we determined that ending each run at the 3,600-second mark achieved an effectively balanced solution.

Table 1 shows a comparison of flight costs from actual schedules and the model. Using the SME (i.e., manual) method, actual flight costs for the 12-week

period were \$134,660; using the MILP model, they were \$74,882—a 44 percent reduction.

Table 2 compares costs for independent contractor consultants. Independent contractors are variable costs because no costs are incurred if they are not used. The model reduced direct labor costs of these consultants by 15 percent from \$323,350 to \$273,630.

The combined flight and labor cost reduction for the 12-week period was \$109,498, a 24 percent reduction in direct costs. This annualizes to savings

	Actual (\$)	Model (\$)	Savings (%)
Week 1	11,996	6,264	48
Week 3	10,721	5,313	50
Week 4	12,355	5,849	53
Week 5	11,036	6,947	37
Week 6	11,266	7,273	35
Week 7	10,831	5,744	47
Week 8	10,210	5,420	47
Week 9	10,737	5,783	46
Week 10	11,600	6,637	43
Week 12	10,503	6,265	40
Week 13	11,886	6,073	49
Week 15	11,519	7,312	37
Total	134,660	74,882	44

Table 1: The table shows flight costs from the actual schedule and the model, and the resulting savings percentages.

	Actual (\$)	Model (\$)	Savings (%)
Week 1	29,640	28,915	2
Week 3	21,545	13,740	36
Week 4	32,090	22,800	29
Week 5	26,840	25,190	6
Week 6	26,715	21,445	20
Week 7	26,895	23,020	14
Week 8	23,795	20,095	16
Week 9	24,240	22,090	9
Week 10	29,315	23,650	19
Week 12	28,645	25,000	13
Week 13	27,515	24,185	12
Week 15	26,115	23,500	10
Total	323,350	273,630	15

Table 2: The table shows consultant variable costs from the actual schedule and the model, and the resulting savings percentages.




Figure 7: Total costs from both the actual and the model are compared (note that F&C represents flight and consultant).

of \$474,490 as a result of routing and assignment improvements. Figure 7 shows the savings for each week modeled.

Sensitivity Analysis and Additional Benefits

Sensitivity analysis, which can be used to improve management decision making or simulate business process changes, is also a benefit of using the model. For this project, we conducted sensitivity analysis on two key business areas involving seven modeling scenarios. We modeled each scenario independently and compared the impact by ranking additional cost savings achieved on a percentage basis (see Table 3 and Figure 8). In the next two paragraphs, we describe the scenarios. Note that the scenarios listed in Table 3 and Figure 8 are shown in parentheses.

The first business area we analyzed was the impact of personnel and training decisions EEI can make in the future. In the first scenario (maximum training), we looked at decisions on the cost savings associated with expanding consultant training and skills development. The second scenario (new hire) involved




Figure 8: The sensitivity analysis graph shows the percentage of savings achieved for each scenario modeled.

determining the best cluster to hire one additional consultant to achieve maximum cost reduction. In the third scenario in this category (relocation), we analyzed quantifying the impact of relocating a consultant from a cluster of excess visit supply to a place of higher visit demand.

The second business area involved the impact of modifying EEI's internal business rules and policy decisions. We tested four specific scenarios in this area. The first scenario we tested (max flights = 6) addressed increasing the number of flights a consultant can complete each week. In the second scenario (availability = 12), we modeled increasing the maximum number of training visits a consultant could complete each week. Third (use excess capacity), we tested using consultants who are used only partially to fulfill future visit needs in upcoming weeks. In the fourth scenario (eliminate fixed), we analyzed the impact of eliminating fixed visits where the client requests a specific consultant and (or) specific date.

Table 3 and Figure 8 show the results of the sensitivity analysis on week 8 costs. The most significant impact was the elimination of fixed visits. Eliminating the 41 fixed visits in week 8 increased savings by 28.8 percent. This change would have saved \$7,225 in additional costs or \$176 per fixed visit. The model realized the \$7,225 savings by reducing staffing levels by one consultant, reducing consultant variable costs, and lowering flight costs.

Other benefits of reducing the number of airline flights include reduced exposure to airline disruption, increased sense of consultant ownership to clients

Test	Scenario	Total cost (\$)	Savings (\$)	Savings (%)
	Original model	25,515		
P&S	Maximum training	25,229	286	1.1
P&S	New hire	25,201	314	1.2
P&S	Relocation	25,347	168	0.7
BR	Max flights = 6	25,449	66	0.3
BR	Availability = 12	25,057	458	1.8
BR	Use excess capacity	24,979	536	2.1
BR	Eliminate fixed	18,290	7,225	28.3

Table 3: We conducted sensitivity analyses for seven business scenarios. This table shows the costs, savings, and percentage saved (note that P&S represents personnel and staffing and BR represents business rules).

in their home geography (i.e., consultants will be assigned for visits to these clients whenever possible), increased productivity because of reduced time on flights, and general morale improvement in field consultants because of less time away from home. Another indirect benefit of using a model for consultant routing and assignment is a reduced dependency on the SME for scheduling and logistics decisions. This improves cross training and company flexibility.

Conclusions and Insights Gained

The heuristic model formulations in this paper are flexible and robust and can be easily manipulated to represent many possibilities encountered in real-life situations. The model has applications in a wide range of settings, including tactical, planning, and strategic environments. The model gives consulting firm management a method for calibrating and quantifying the impact of current decisions on future operations by running multiple what-if scenarios.

When EEI implemented the solution in its consultant scheduling process, it realized travel cost savings and improved employee deployment. These savings have nearly achieved the projected 24 percent travel cost reduction; however, because overhead savings (i.e., savings related to the overhead costs of the manual method relative to the application of the solution described above) have not yet been realized, additional work is underway to gain user acceptance and full adoption of the solution. The biggest implementation challenge has been integrating various legacy databases in an automated fashion. EEI currently pulls information from five databases to gather comprehensive information about clients, consultants, and travel. Because these databases are constantly updated, the SME must make manual queries each week to populate the model, a time-consuming process; thus, the SME's weekly work effort has not yet been reduced. Although EEI is in the process of automating its data integration, the project will take several years because of budget constraints and other factors.

Another implementation challenge relates to the adoption of this new process by the consultants. They have traditionally been given much autonomy in promising dates for follow-up visits to clients.

As needed, the SME would adjust the schedules to accommodate these visit dates. Sensitivity analysis of this project provided EEI with the information needed to eliminate this common practice in favor of more structured routing and assignment flows, resulting in some resentment and resistance from the consultants. In retrospect, consultants should have been more involved in the data collection and model development process to ensure their acceptance and ownership of the process changes.

In summary, this research significantly improves consultant planning models and algorithms for EEI. The new process makes a significant impact on the firm's profit, staff utilization, and overall productivity, while improving consultant quality of life and safety.

Appendix A. Model for Clustering Clients

Notation

M = number of client locations.

P = number of potential airportcentric cluster locations.

c_j = hub rating or cost associated with cluster j .

a_{ij} = 1 if client i is within a cluster radius of cluster j ;
0 otherwise.

x_j = 1 if cluster j is used; 0 otherwise.

$$\text{Min } \sum_{j=1}^P c_j x_j \quad (\text{A1})$$

$$\text{subject to } \sum_{j=1}^P a_{ij} x_j \geq 1, \quad \text{for all } i = 1, 2, \dots, M, \quad (\text{A2})$$

$$x_j \in \{0, 1\}, \quad \text{for all } j = 1, 2, \dots, P. \quad (\text{A3})$$

The objective function in Equation (A1) minimizes the total cost of airport hub clusters. Equation (A2) requires the number of clusters covering client i to be greater than or equal to 1. The radius of each cluster is determined based on how a flight-versus-drive decision is made. For this project, we determined the radius to be 300 miles for all clusters based on a survey of the firm's consultants, which we conducted. Equation (A3) requires the decision variables to be binary.

Appendix B. Model for Routing and Assigning Consultants to Clusters Notation

I = number of consultants to be assigned.

C = number of clusters of clients to be scheduled (determined by solving the model in Appendix A).

N = maximum number of trip legs allowed for a consultant.

S = number of skill levels.

r_{jk} = flight cost between clusters j and k .

q_{ik} = flight cost between consultant i 's home and cluster k .

t_{ijs} = number of s -type fixed visits needed from consultant i to cluster j .

u_{is} = 1 if consultant i has skill level s ; 0 otherwise.

d_{ks} = demand for skill level s visit by cluster k (i.e., number of visits).

a_i = number of shifts available from consultant i .

w_i = weekly wage for contract consultant i if he or she is assigned.

y_{ijkn} = 1 if consultant i travels from cluster j to k on the n th leg of trip; 0 otherwise.

y_{i0k1} = 1 if consultant i travels from home to cluster k on the first leg of trip; 0 otherwise.

$y_{ij(C+1)n}$ = 1 if consultant i returns home from cluster j on the n th leg of trip; 0 otherwise.

x_{ijs} = number of visits assigned to consultant i for cluster j for s type of visit.

$$\text{Min } \left\{ \sum_{i=1}^I \sum_{k=1}^C (q_{ik} + w_i) y_{i0k1} + \sum_{i=1}^I \sum_{j=1}^C \sum_{k=1}^C \sum_{n=2}^N r_{jk} y_{ijkn} + \sum_{i=1}^I \sum_{j=1}^C \sum_{n=2}^N q_{ij} y_{ij(C+1)n} \right\}, \quad (\text{B1})$$

$$y_{i0k1} \geq \sum_{l=1, l \neq k}^{C+1} y_{ikl2}, \quad \text{for all } k = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B2})$$

$$y_{i0kn} = 0, \quad \text{for all } n = 2, \dots, N; k = 1, 2, \dots, C, C+1 \text{ and } i = 1, 2, \dots, I, \quad (\text{B3})$$

$$y_{ijk1} = 0, \quad \text{for all } k = 1, \dots, C, C+1; j = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B4})$$

$$\sum_{j=1, j \neq k}^C y_{ijkn} \geq \sum_{l=1, l \neq k}^{C+1} y_{ikl(n+1)}, \quad \text{for all } n = 2, 3, \dots, N-1; k = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B5})$$

$$y_{i(C+1)kn} = 0, \quad \text{for all } n = 1, 2, \dots, N; k = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B6})$$

$$\sum_{j=1}^C \sum_{n=1}^N y_{ij(C+1)n} = \sum_{k=1}^C y_{i0k1}, \quad \text{for all } i = 1, 2, \dots, I, \quad (\text{B7})$$

$$\sum_{k=1}^{C+1} \sum_{n=1}^N y_{ijkn} \leq 1, \quad \text{for all } j = 0, 1, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B8})$$

$$\sum_{j=0}^C \sum_{n=1}^N y_{ijkn} \leq 1, \quad \text{for all } k = 1, 2, \dots, C+1 \text{ and } i = 1, 2, \dots, I, \quad (\text{B9})$$

$$\sum_{s=1}^S x_{iks} \leq a_i \sum_{j=0}^C \sum_{n=1}^N y_{ijkn}, \quad \text{for all } k = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B10})$$

$$\sum_{s=1}^S x_{iks} \geq \sum_{j=0}^C \sum_{n=1}^N y_{ijkn}, \quad \text{for all } k = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B11})$$

$$y_{ijjn} = 0, \quad \text{for all } n = 1, 2, \dots, N; j = 0, 1, \dots, C+1 \text{ and } i = 1, 2, \dots, I, \quad (\text{B12})$$

$$\sum_{i=1}^I x_{iks} = d_{ks}, \quad \text{for all } s = 1, 2, \dots, S \text{ and } k = 1, 2, \dots, C, \quad (\text{B13})$$

$$x_{iks} \geq t_{iks}, \quad \text{for all } s = 1, 2, \dots, S; k = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B14})$$

$$x_{iks} \leq a_i u_{is}, \quad \text{for all } s = 1, 2, \dots, S; k = 1, 2, \dots, C \text{ and } i = 1, 2, \dots, I, \quad (\text{B15})$$

$$\sum_{k=1}^C \sum_{s=1}^S x_{iks} \leq a_i, \quad \text{for all } i = 1, 2, \dots, I. \quad (\text{B16})$$

Objective Function for Routing and Assignment

The objective of the model, as defined in Equation (B1), is to minimize the total flight cost for all consultants and the variable labor cost for all assigned contract consultants.

The flight costs include the travel from and to consultants' homes and among clusters. Variable labor cost is the additional weekly expense incurred by scheduling a consultant who is employed as an independent contractor. The first component of the objective function minimizes the travel cost for the first visit the consultant makes; it includes the weekly wage because that wage is paid if the consultant is used for any client visit. The middle component of the objective function minimizes costs for the travel between clusters. The final component of the objective function minimizes costs for returning the consultant home at the end of the weekly scheduling period.

Constraints for Routing and Assignment

Constraints (B2) to (B3) define how each consultant will travel from home to various clusters and return home at the end of the period.

Constraint (B2) implies that if a consultant does not leave home on the first leg of a trip, that consultant will not make any more trips in the current scheduling period. Constraints (B3) and (B4) enforce that the first leg of a trip for any consultant must be the one in which that consultant leaves home. Constraint (B5) ensures that, starting from the second leg of a trip, a cluster will not be a starting cluster of a new leg of a trip if it is not the destination in the previous leg. Constraint (B6) enforces that a consultant will not travel on any additional leg of a trip after returning home. Constraint (B7) ensures that a consultant will return home if he or she leaves home for a trip. Constraints (B8) and (B9)

mean that a consultant travels from or to a particular cluster at most once in an entire trip. Constraint (B10) enforces the constraint that a consultant can only be assigned to a cluster if the cluster is one of that consultant's destinations. Constraint (B11) ensures that at least one visit is scheduled if a consultant stops over at a cluster. Constraint (B12) means that no leg of a trip is allowed if the origin and the destination are the same cluster.

Constraints for Cluster Demand

Constraints (B13) and (B14) ensure that each cluster demand of quantity of skills is fully satisfied by the available pool of available consultants. Constraint (B13) means that demand for a specific type of visit by a cluster will be satisfied. Constraint (B14) ensures that all fixed visits are scheduled accordingly.

Constraints for Consultant Supply

Constraints (B15) and (B16) require that only consultants with a specific skill can satisfy a specific demand, and that consultants have a finite capacity of availability during the period. Constraint (B15) is the skill constraint (i.e., a consultant who does not have a necessary skill will not be scheduled for a type of visit requiring that skill). Constraint (B16) enforces the capacity constraint for a consultant.

Acknowledgments

We extend a special thank you to Jeffery Kennington of Southern Methodist University. The time he spent with us on-site at a client location and his support throughout the study were invaluable. We would also like to thank Eli Olinick of Southern Methodist University for his review of this manuscript.

References

- Beck JC, Prosser P, Selensky E (2003) Vehicle routing and job shop scheduling: What's the difference. *Proc. 13th Internat. Conf. Automated Planning Scheduling, Trento, Italy*, 267–276.
- Bard J, Kontoravdis G, Yu G (2002) A branch-and-cut procedure for the vehicle routing problem with time windows. *Transportation Sci.* 36(2):250–269.
- Cheng CB, Wang KP (2009) Solving a vehicle routing problem with time windows by a decomposition technique and a generic algorithm. *Expert Systems Appl.* 36(4):7758–7763.
- Diaz J (2005) A compact model and tight bounds for a combined location-routing problem. *Comput. Oper. Res.* 32(3):407–428.
- Mastrogiannis N, Giannikos I, Boutsinas B (2011) Solving the set covering problem based on a new clustering heuristic. Presentation, Department of Business Administration, University of Patras, Greece.
- Toregas C, Swain R, ReVelle C, Bergman L (1971) The location of emergency service facilities. *Oper. Res.* 19(6):1363–1373.
- Toth P, Vigo D (2002) *The Vehicle Routing Problem* (Society for Industrial and Applied Mathematics, Philadelphia).

Verification Letter

Richard H. Roberson, CPA, Chief Financial Officer, Energy Education, Inc., 5950 Sherry Lane, Suite 900, Dallas, Texas 75225, writes: "Energy Education, Inc. (EEL) has a single goal with clients of reducing energy consumption through culture change. Achieving this goal requires a combination of practical experience in organizational behavior, technical expertise, the ability to effectively execute a large-scale, complex program and the ability to teach an organization to own and sustain their tailored program for the long haul. The company has served more than 1,100 clients in 48 states over 25 years and helped our clients save in excess of \$2.3 billion in energy costs.

"Our business model requires intense, personal interactions with each client and that generates significant travel expenses. In fact, employee travel costs are the second largest expense item for EEL after salaries. As an example, in 2010 the company spent nearly \$1,000 per consulting employee per week. There has also been a significant amount of management overhead expense associated with managing and updating travel schedules.

"The operations research project that Dr. Yu and Mr. Hoff completed has generated real savings and helped the company identify areas of further productivity gains. These include:

- (1) Reduction in consultant travel expenses of 24% through more efficient flight selection and employee routing
- (2) Reduction in overhead expense for scheduling decision making by subject matter experts
- (3) Improvement in selection criteria for new hire and employee training decisions.

"I verify the actual usefulness and benefits of the project work completed for the Optimal Routing and Assignment of Consultants project."

Junfang Yu is an assistant professor in the Department of Engineering Management, Information, and Systems at Southern Methodist University (SMU). Before joining SMU, he worked as a supply chain and logistics solution architect at i2 Technologies and as an assistant professor at Louisiana State University. He was also on the faculty at Shanghai Jiao Tong University. His research interests are in the areas of supply chain and logistics optimization and systems engineering.

Randy Hoff is CEO of Dallas-based Cenergistic. He holds a doctor of engineering degree in engineering management from Southern Methodist University, an MBA from Marquette University, and bachelor's degree in industrial engineering from Texas A&M University. He is a licensed professional engineer and certified energy manager.