

Paper 201-2008

Using SAS/OR[®] to Improve Service Completion at a Fortune 100 Company

Chuck Kelly and Jinxin Yi, SAS Institute Inc., Cary, NC

Philip Easterling, SAS Institute Inc., Houston, TX

ABSTRACT

The services group at a Fortune 100 company provides maintenance, warranty, and repair services for customers who request assistance with their appliances or equipment. The company handles millions of service orders annually, with about one-half of the orders requiring parts to complete the service. A significant number of those orders cannot be completed on the first service attempt because the required parts are not available on service technicians' trucks. Incurring an additional cost for each follow-up visit, a cost which is not passed through to the customer, the company estimates its annual revisit cost to be over \$150 million.

In an effort to increase the first-visit service completion rate, we developed a solution that optimizes both technician assignment and service part assortment. This solution relies on several analytical technologies, including SAS/OR, SAS[®] Inventory Optimization, and SAS[®] Simulation Studio. Tests that use historical data revealed that this SAS solution can improve the first-visit service completion rate by about 9%. This improvement yields a projected annual savings of over \$15 million in revisit cost with a minimal increase in inventory-carrying cost while maintaining current customer service levels.

This paper describes the details of the business problem and the comprehensive solution used to achieve the projected savings.

INTRODUCTION

The services group at a Fortune 100 company, hereafter referred to as ABC to preserve anonymity, provides in-home services to customers who call in for help when their appliances or equipment need repair or maintenance. This group employs approximately 10,000 service technicians who serve about 50 separate geographic zones across the country. Each zone is divided into multiple service areas. A service area is a geographic region defined by one or more zip codes. Usually, technicians are assigned to work within a single service area, although they are sometimes dispatched to calls in areas adjacent to their assigned service area.

ABC handles millions of customer service orders annually. Types of orders include repair, warranty service, and scheduled maintenance. Each service order is assigned to a specific service category determined by the type of appliance or equipment that needs service. ABC defines about 30 different service categories, including cooking appliances, laundry appliances, recreational equipment, and heating and air conditioning. Service technicians are assigned to cover one or more categories, based on their individual skills and abilities and their geographic location.

When responding to a customer order, the goal is to complete service on the service technician's first visit. The ability to meet this goal is greatly affected by the availability of spare parts on the technicians' trucks. Of over 10 million orders that ABC responds to annually, about one-half require parts to complete service. Several issues affect the likelihood that a technician will have the appropriate part for a service visit. These include the number of assigned service categories, the wide variety of makes and models to be serviced, services that require multiple quantities of a given part, limited truck capacity, and the infrequency of individual part demand. Each of these issues complicates the process of building an effective part assortment for each technician.

Besides decreasing customer satisfaction and the capacity to serve additional customers, multiple service attempts also increase the cost to deliver service. A revisit cost is incurred for each subsequent visit after the first. ABC estimates that this cost alone is over \$150 million annually. The goal of our work is to reduce this cost by increasing the first-visit service completion rate.

To achieve this goal, we developed a mixed-integer optimization model to determine the assignment of technicians to service categories. We also created a method for calculating an effective service part assortment for each technician. Running a simulation model with historical data, we project a 9% improvement in the first-visit service completion rate. This improvement yields a projected annual savings of over \$15 million in revisit cost with a minimal increase in inventory-carrying cost and maintains or improves current customer service levels.

PROPOSED SOLUTION

To improve service completion at ABC, we applied knowledge and expertise from several different product areas. The proposed solution relies on analytical engines and principles from the following software components: SAS/OR, SAS Inventory Optimization, SAS Simulation Studio, and SAS® High-Performance Forecasting. Figure 1 displays the dependencies between the high-level components of the solution. The primary components of the solution are the modules that calculate the technician assignment and the parts assortment.

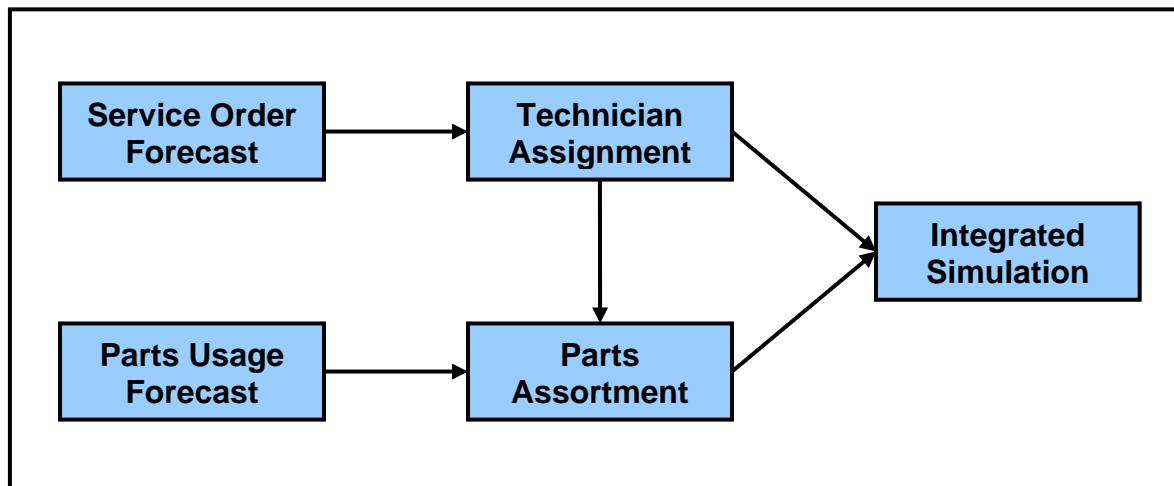


Figure 1. Solution Component Dependencies

The solution requires two separate forecasts as input: a monthly service order forecast for each service category within each service area and a monthly forecast of part usage within each service area. These two forecasts provide input to the technician assignment and parts assortment modules, respectively. The technician assignment module also provides input for the parts assortment module. Specifically, a technician's part assortment is based on a technician's assigned service categories. Finally, the integrated simulation takes output from both modules to simulate system behavior by using historical data. The simulation provides summary results and performance metrics for evaluating the impact of the solution.

We focus our discussion on the technician assignment, parts assortment, and integrated simulation components; details about the forecasting components are beyond the scope of this paper.

TECHNICIAN ASSIGNMENT

One way to increase the likelihood that a technician will have the required parts for a service order is to limit the number of service categories that he is required to cover. The goal of the technician assignment module is to effectively manage the assignment of categories to technicians while providing sufficient coverage to meet the expected service call volume. Preferably, the workload will be fairly balanced, with no technicians having too little utilization. Also, there is a limit on the amount of expected overtime that might be required to achieve feasibility due either to limited skill availability or to variability in the distribution of incoming call volumes by service category during a given month.

We tried several different model formulations before arriving at the one that best met the desired results. The formulation that we chose balances between minimizing the number of categories that a technician can be assigned to and minimizing the number of total technicians that are required to cover the anticipated demand. In this way, we reduce the assignment of technicians who will have minimal utilization. The optimization model is designed to handle the expected call volume, but we did not build in allowances for any redundancy. In most service areas, not all technicians are planned to be fully utilized; thus, there is generally excess technician capacity available to cover random peaks in demand without adversely affecting response time.

Customer service orders fall into one of about 30 different service categories. For each service area, the number of service orders within a month is forecasted for every service category. Assuming a steady arrival of calls within a month, we use the number of working days within the month to determine the expected call load for each day. Then, using the expected number of visits per service order and the expected number of service visits that a technician can

handle per day, we estimate the average number of full-time equivalent technicians that are needed for each service category per day in each service area. Often, the daily technician requirement for many categories is less than one, except for only the most popular service categories such as laundry appliances and refrigerators. As might be expected, the requirements can differ based on the season and geographic location.

MODEL FORMULATION

The technician assignment problem is formulated as a mixed-integer optimization model for each service area. For this formulation, *Technician* represents the set of technicians and *Category* represents the set of different service categories. *TechSkills* is the indexed set of service categories that each technician is qualified to serve. *CatDemand_c* specifies the number of technicians required for each service category per day.

The decision variable *useTech_t* is a binary variable that indicates whether or not a technician is assigned to any service category. A value of 1 indicates that the technician is used, and a value of 0 indicates that the technician is not used. The binary decision variable *techCat_{t,c}* indicates that technician *t* is assigned to category *c*. A value of 1 indicates an assignment. Technicians can be assigned to multiple service categories, but they can be assigned only to categories for which they are qualified.

The decision variable *techCatRate_{t,c}* is a continuous variable that indicates the percentage of a regular working day that technician *t* spends on service category *c*. Similarly, the continuous decision variable *techCatOTRate_{t,c}* indicates the overtime percentage that technician *t* works on a category *c*. Note that a technician's regular and overtime work rates cannot exceed 100% and 50%, respectively, and the objective function effectively minimizes technician overtime. Excessive technician overtime is an indicator that additional technician capacity is needed for one or more service categories in that service area. Likewise, an infeasible solution typically indicates insufficient technician capacity within the service area.

The model formulation is represented as follows:

$$\begin{aligned} \min \quad & \sum_{t \in \text{Technician}} \text{techCost} * \text{useTech}_t + \sum_{\langle t,c \rangle \in \text{TechSkills}} [\text{techRateCost} * \text{techCatRate}_{t,c} + \text{techOTCost} * \text{techCatOTRate}_{t,c}] \\ & + \text{maxTechCatCost} * \text{maxTechCat} + \text{maxTechOTCost} * \text{maxTechOTRate} \end{aligned}$$

$$\begin{aligned} \text{subject to} \quad & \sum_{c \in \text{Category}} \text{techCatRate}_{t,c} \leq 1 && \forall t \in \text{Technician} && (\text{rate}) \\ & \sum_{c \in \text{Category}} \text{techCatOTRate}_{t,c} \leq 0.5 && \forall t \in \text{Technician} && (\text{overtime}) \\ & \sum_{t \in \text{Technician}} [\text{techCatRate}_{t,c} + \text{techCatOTRate}_{t,c}] \geq \text{CatDemand}_c && \forall c \in \text{Category} && (\text{demand}) \\ & \text{techCatRate}_{t,c} \leq \text{techCat}_{t,c} && \forall \langle t,c \rangle \in \text{TechSkills} && (\text{bin1}) \\ & \text{techCatOTRate}_{t,c} \leq \text{techCat}_{t,c} && \forall \langle t,c \rangle \in \text{TechSkills} && (\text{bin2}) \\ & \text{techCat}_{t,c} \leq \text{useTech}_t && \forall \langle t,c \rangle \in \text{TechSkills} && (\text{bin3}) \\ & \sum_{c \in \text{Category}} \text{techCat}_{t,c} \leq \text{maxTechCat} && \forall t \in \text{Technician} && (\text{maxCat}) \\ & \sum_{c \in \text{Category}} \text{techCatOTRate}_{t,c} \leq \text{maxTechOTRate} && \forall t \in \text{Technician} && (\text{maxOT}) \end{aligned}$$

all variables are nonnegative; *useTech_t* and *techCat_{t,c}* are binary

The decision variable *maxTechCat* is a continuous variable that represents the maximum number of different service categories to which any single technician can be assigned. The cost *maxTechCatCost* associated with *maxTechCat* acts as a control on the diversity of technicians' assignments. The continuous decision variable *maxTechOTRate* indicates the maximum overtime percentage that can be assigned to any single technician. The purpose of this variable and corresponding constraint is to even out the distribution of overtime across the entire set of technicians.

The *rate* constraint for each technician ensures that the sum of the rates of a technician's category assignments does not exceed 100%. Similarly, the *overtime* constraint for each technician limits a technician's total overtime assignment to 50%. The *demand* constraint for each category dictates that the optimized solution provides technician assignments that meet the required level of coverage for each service category.

The sets of constraints, *bin1*, *bin2* and *bin3*, are used to ensure the integrity of the binary variables. For instance, if a technician *t* is assigned to work at any rate on category *c*, the corresponding *techCat_{t,c}* variable has a value of 1; otherwise, it has a value of 0. Similarly, the *useTech_t* variable is forced to have a value of 1 when any corresponding *techCat_{t,c}* variable is positive; otherwise, it has a value of 0 at optimality, since the objective is minimization.

The objective function attempts to maintain a balance among the costs of using technicians, assigning technicians to multiple service categories, and requiring technician overtime. The objective function costs alone are not tied to any business-related values. To help achieve feasibility, technician overtime is allowed; however, the cost for maximum technician overtime rate, *maxTechOTcost*, is set to the highest value of all objective function coefficients to place a heavy penalty on excessive overtime assignments. The overtime cost for a technician, *techOTCost*, is set to a higher value than the standard rate, *techCostRate*, to ensure that technicians' standard rate is 100% before they are assigned any overtime. The cost for maximum technician categories, *maxTechCatCost*, is used to control the diversity of technician assignments.

MODEL SOLUTION

The mixed-integer optimization problem was solved individually for each service area to determine the technician assignments for all technicians within the zone. The problem was defined and solved by using the OPTMODEL procedure and the MILP solver, recent additions to the SAS/OR product. PROC OPTMODEL provides an algebraic modeling language that makes it very simple to define and solve optimization problems. The syntax is very intuitive and mirrors the description of the problem formulation in the previous section.

Consider an excerpt of the statements that define the technician assignment problem:

```
var useTech{Technician} binary;
var techCat{TechSkills} binary;
var techCatRate{TechSkills} >= 0 <= 1;
var techCatOTRate{TechSkills} >= 0 <= 0.5;
var maxTechCat >= 0;
var maxTechOTRate >= 0;

min obj = sum{t in Technician} techCost*useTech[t]
          + sum{<t,c> in TechSkills}
            (techRateCost*techCatRate[t,c] + techOTCost*techCatOTRate[t,c])
          + maxTechCatCost*maxTechCat + maxTechOTCost*maxTechOTRate;

con rate {t in Technician}:
    sum{c in Category: <t,c> in TechSkills} techCatRate[t,c] <= 1;

con overtime {t in Technician}:
    sum{c in Category: <t,c> in TechSkills} techCatOTRate[t,c] <= 0.5;

con demand {c in Category}:
    sum{t in Technician: <t,c> in TechSkills}
      (techCatRate[t,c] + techCatOTRate[t,c]) >= CatDemand[c];

con bin1 {<t,c> in TechSkills}: techCatRate[t,c] <= techCat[t,c];
con bin2 {<t,c> in TechSkills}: techCatOTRate[t,c] <= techCat[t,c];
con bin3 {<t,c> in TechSkills}: techCat[t,c] <= useTech[t];

con maxCat {t in Technician}:
    sum{c in Category: <t,c> in TechSkills} techCat[t,c] <= maxTechCat;

con maxOT {t in Technician}: sum{c in Category: <t,c> in TechSkills}
    techCatOTRate[t,c] <= maxTechOTRate;
```

This excerpt contains the variable declarations and the objective and constraint definitions. By defining the optimization model separately from the data, it is very easy to adjust the model formulation with minimal changes in the statements. The solution of the technician assignment problem is written to SAS data sets for use by the parts assortment module.

PARTS ASSORTMENT

After the technician service assignments are determined, we use this information to drive the calculation of the parts assortment for each technician. Part usage is forecasted for all parts (per service category) within each service area. With this forecast and the technician assignments, the parts assortment module calculates a part assortment mix for each individual technician.

Each technician has his own designated truck in which to store spare parts for use on service calls. Since these trucks have a limited storage capacity, we must consider each individual part that can be added to the assortment mix. Some parts are excluded from consideration due to ABC's business rules. For instance, some parts might be too large or too expensive to stock, and other parts are not stocked due to safety reasons.

In ABC's service chain, technicians' inventories are typically replenished one-for-one on a weekly or semi-weekly basis. In other words, each part is replenished after it is used, and deliveries are made once or twice per week. For the purposes of this research, we assumed that replenishment occurs only once per week. Thus, each technician's inventory needs to cover at most one week's worth of expected part demand. We developed the following methodology to solve the parts assortment problem.

ASSORTMENT METHODOLOGY

The first step is to establish a candidate part list from which to work because the size of the entire spare part list for ABC is quite large. When building a part assortment, ABC's practice has been to consider stocking only parts that have substantial usage. In an effort to create a more effective part assortment, we created a candidate part list that consists of the set of truck-eligible parts that were used on service calls within the service area during the preceding calendar year. For all parts in the candidate list, we use the parts usage forecast to derive an estimate of the expected weekly demand for each part within the context of the service categories for which the part is used.

Next, in order to work toward creating a demand estimate for each technician, we determine the technician coverage per service category. Based on the technicians' assigned service categories and the service category call-volume percentage within the service area, we calculate the total number of full-time *virtual* technicians (the sum of the individual technician rates) covering each category within the service area. With this information, we can apportion the expected part demand across the service area to individual technicians. We refer to this demand as the technician's *demand-to-cover*.

To calculate a technician's demand-to-cover for each part, we first determine the demand of the part for a single *specialized* technician, a technician who works full-time on that part's single service category. This value is calculated by dividing the part's weekly demand by the number of virtual technicians assigned to work on the part's service category. Then, to calculate the demand-to-cover for each technician, we multiply the demand per specialized technician by the percentage of time that the technician is assigned to work on the service category.

We create a candidate part list for each individual technician by selecting the set of parts that were used on service calls within their service categories inside their service area within the preceding year. If the technician is assigned to work on multiple service categories, we summarize the data to establish a single quantity for demand-to-cover for each unique part in his candidate list. With the demand-to-cover for each part in a technician's candidate list, we can build the technician's part mix.

The first step in building the technician's part mix is to determine the optimal inventory amount (for each part) that balances the revisit cost and the inventory holding cost. If we keep n units of a part in inventory for an anticipated demand D , then the expected inventory shortage is:

$$E(\max\{0, D - n\})$$

and the expected excess is:

$$E(\max\{0, n - D\})$$

Thus, the total expected cost is equal to the sum of the revisit cost and inventory holding cost. It is represented by the following expression:

$$\begin{aligned} f(n) &= \text{revisitCost} * E(\max\{0, D - n\}) + \text{unitHoldingCost} * E(\max\{0, n - D\}) \\ &= \text{revisitCost} * [E(D - n) + E(\max\{0, n - D\})] + \text{unitHoldingCost} * E(\max\{0, n - D\}) \\ &= \text{revisitCost} * E(D - n) + (\text{revisitCost} + \text{unitHoldingCost}) * E(\max\{0, n - D\}) \\ &= \text{revisitCost} * E(D) - \text{revisitCost} * n + (\text{revisitCost} + \text{unitHoldingCost}) * [n * CDF(n) - E(D) * CDF(n - 1)] \end{aligned}$$

where $E(D)$ is the expected demand for the part and $CDF(n)$ is the cumulative distribution function valued at n . The $unitHoldingCost$ is equal to the unit cost of the part multiplied by the holding cost rate.

To minimize the total expected cost, we need to find $n = S^*$ that minimizes $f(n)$. The optimal inventory amount S^* is the smallest integer that satisfies the following inequality:

$$CDF(S^*) \geq revisitCost / (revisitCost + unitHoldingCost)$$

We refer to the quantity on the right side of the inequality as the *critical ratio*.

Each technician's demand-to-cover is used as the expected demand. Because of the low volume of demand for most parts, we use the Poisson distribution to model demand. Logically, if there is no capacity constraint on a technician's inventory, we would stock S^* units of each part because carrying more than S^* units only increases the total cost. Therefore, we consider S^* to be the *maximum inventory level* for that part, and the technician's assortment will contain no more than S^* units of each part. Note that S^* might vary from part to part as a function of the expected demand and the part cost.

Once we have the maximum inventory level, the second step in building the technician's part mix is to determine how many units of each part should be stocked subject to the technician's truck capacity constraint. We use a *marginal analysis* approach to decide how many part units to stock. This approach is fairly straightforward.

For all candidate parts, we calculate the marginal contribution of each part unit up to the maximum inventory level. The marginal contribution is defined as the cost reduction (in the total cost) by stocking an additional unit of the part divided by the part's physical volume (length x width x height).

The cost reduction, or *incremental cost*, of the n th unit of a part is equal to $f(n) - f(n - 1)$. Using the expression for $f(n)$ defined above, the incremental cost is expressed as:

$$\begin{aligned} incrementalCost = & -1 * revisitCost \\ & + (revisitCost + unitHoldingCost) * [CDF(n) + (n-1) * PDF(n) - E(D) * PDF(n-1)] \end{aligned}$$

where $PDF(n)$ is the probability density function evaluated at n and $CDF(n)$ is the cumulative distribution function evaluated at n .

The following SAS statements show the logic to calculate the optimal inventory amount S^* , or *maximumInventory*, and the marginal contribution, or *ranking*, for each part in the technician's candidate list:

```
data candidatePartRanking;
  set candidatePartList;
  unitHoldingCost = unitCost*&holdingCostRate;
  criticalRatio = &revisitCost / (&revisitCost + unitHoldingCost);
  maximumInventory = quantile('POISSON', criticalRatio, demandToCover);
  pdf0 = pdf('POISSON', 0, demandToCover);
  do n=1 to maximumInventory;
    cdf = poisson(demandToCover, n);
    pdf1 = pdf('POISSON', n, demandToCover);
    incrementalCost = -1 * &revisitCost
      + (&revisitCost + unitHoldingCost)*(cdf + (n-1)*pdf1 - demandToCover*pdf0);
    ranking = incrementalCost / volume;
    output;
    pdf0 = pdf1;
  end;
run;
```

After sorting by the part ranking, we scan the list to add individual units of parts to the technician's assortment mix until the technician's truck capacity is reached or the candidate list is exhausted. Finally, we summarize the list of individual part units to create the technician part list and stocking quantities.

```
proc sort data=candidatePartRanking;
  by ranking;
run;
```

```

data candidateSelection;
  set candidatePartRanking;
  retain cumVolume 0;
  if (cumVolume + volume) >= &truckCapacity then stop;
  cumVolume+volume;
run;

proc sql;
  create table technicianAssortment as
  select distinct technician, part, max(n) as stockQuantity
  from candidateSelection
  group by part
  order by part;
quit;

```

The parts assortment methodology can be summarized in these steps:

- Establish the candidate part list.
- Derive the weekly demand for each candidate part.
- Determine the total technician coverage per service category.
- Determine the demand-to-cover of each part for each technician.
- Build the candidate list of parts for each technician.
- Calculate the marginal contribution of individual part units.
- Build the technician assortment subject to the truck capacity constraint.

INTEGRATED SIMULATION

To validate the effectiveness of our solution, we designed and developed a simulation model to emulate ABC's service history for a single month. Essentially, the integrated simulation "replays" historical service orders by using the optimized technician assignments and corresponding parts assortments. It compiles and outputs statistics on service performance. Because the optimized technician assignments are different from the historical ones, it is necessary to create a model that assigns individual service orders to appropriate technicians in order to evaluate the ability to complete service on the first visit.

SIMULATION MODEL

ABC's service history is used to create a queue of service orders to be processed. Since some service orders require multiple visits to complete, there is a queue of in-progress service orders when the simulation period begins. Service orders are processed in the order of their arrival as determined by the date that the orders are created.

To simplify the modeling logic, service orders are assumed to be completed in one or two technician visits. Service orders that do not require parts are assumed to be completed on the first visit. If a service order requires parts and the assigned technician has all of the required parts in inventory, the service order is also assumed to be completed on the initial visit. However, if all required parts are not available on the initial visit, the service order is put back into the queue to receive a second service visit. Scheduling of the second service visit is delayed one week to allow for the shipment of the required parts to the customer. According to ABC's company policy, the same technician does not need to provide service on the subsequent visit.

During the simulation, service orders are assigned to a technician based on the service order category and the availability of an appropriately skilled technician. Priority of technician selection is based on choosing the first available technician with the fewest assigned service categories. Consistent with ABC's standard procedures, this rule works to conserve availability of technicians who can serve multiple service categories. Thus, those technicians who serve fewer categories are scheduled first. If no suitable technician is available, the service order is left in the queue until a suitable technician becomes available. The simulation is run over the set of historical service orders that occurred during the selected month.

SIMPLIFYING ASSUMPTIONS

Certain simplifying assumptions were necessary:

- The technician service rate is constant and determined by the service category:
 - Service rate = 1 / (number of planned service visits per day for the service category).
 - This rate also serves as a proxy for length of service visits for each category.

- Technicians work only on assigned service categories in their designated service area.
- Technicians work five days per week.
- Technician vacation, holiday, and sick time is not considered.
- Technicians begin with a complete parts inventory, based on their part assortment mix.
- Technician part replenishment occurs 3 weekdays after part usage.
- Replenished parts are considered to be available for use on the day of their receipt.

IMPLEMENTATION

We used SAS Simulation Studio to build and run the simulation model. SAS Simulation Studio is a new addition to SAS/OR in SAS 9.2. It provides an interactive visual framework for building and running discrete event simulation models. Essential to our work is the ability to create custom objects with embedded logic for use within the simulation. SAS Simulation Studio makes it easy to save simulation statistics into SAS data sets for further analysis and reporting. Figure 2 shows the simulation model as it appears within SAS Simulation Studio.

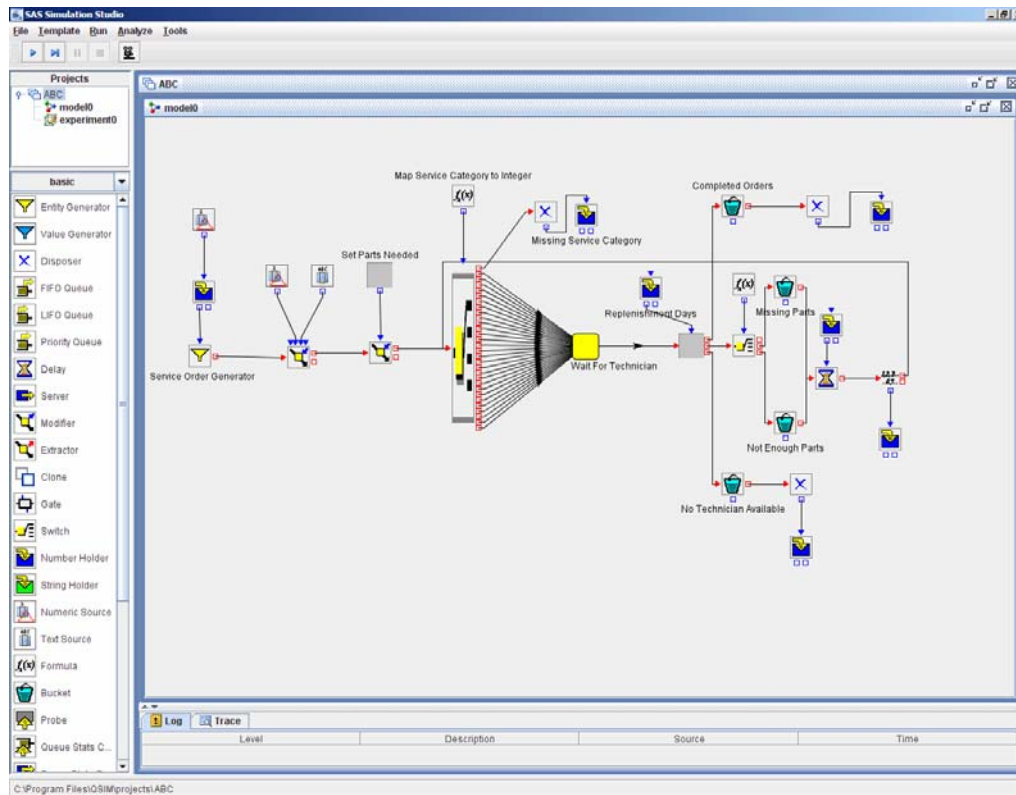


Figure 2. Integrated Simulation Model in SAS Simulation Studio

Since technicians are assigned to work in a specific service area, we run the simulation for each separate service area. The simulation determines how many service orders are completed on the first visit and how many orders require return visits. The simulation also provides the ability to distinguish between those service orders that failed to be completed on the first visit because the parts are out of stock or not stocked. This information is helpful in evaluating the quality of the inventory stocking quantities versus the set of parts included in the assortment mix. Additionally, the simulation provides an estimate of the service order response and cycle times.

By analyzing these outputs, we are able to compare our solution to ABC's actual performance. More importantly, the simulation provides insight about how we can adjust the technician assignment and parts assortment modules to achieve better solutions.

CONCLUSION

By combining technician assignment with parts assortment in a single solution, we are able to improve the technicians' part mixes by reducing the number of categories that each technician serves, versus the smaller improvement from simply changing parts assortment alone. We refer to this concept as *increasing technician*

specialization or decreasing technician category diversity. SAS Simulation Studio proved to be a unifying environment in which we could evaluate the impact of both modules simultaneously—giving us the opportunity to tune our optimization models and methods to generate better results.

Evaluating statistics compiled from running the integrated simulation model, we estimate a potential 9% improvement in ABC's first-visit service completion rate within the sample zone. With an additional part-carrying cost of approximately \$800,000, we project that our solution could save over \$15 million in revisit costs annually. Also, the projected average service order response time is reduced slightly from current ABC estimates. In addition, improving first-visit service completion frees up valuable technician capacity, which can allow ABC to process additional service orders without adding more technicians.

REFERENCES

- Gorman, M. and S. Ahire. 2006. "A Major Appliance Manufacturer Rethinks Its Inventory Policies for Service Vehicles." *Interfaces* 36(5): 407–419.
- SAS Institute Inc. 2006. *SAS Inventory Optimization 1.3: User's Guide*. Cary, NC: SAS Institute Inc.
- SAS Institute Inc. 2007. *SAS Simulation Studio User's Guide*. Cary, NC: SAS Institute Inc.
- SAS Institute Inc. 2007. *SAS/OR 9.1.3 User's Guide: Mathematical Programming 3.2*. Cary, NC: SAS Institute Inc.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the contributions of the following people to the research effort presented in this paper: Joe Katz, Hong Chen, Emily Lada, Phil Meanor, Skip Smith, Ellen Maxon, and Don Bennett.

The authors gratefully acknowledge the following people who reviewed and offered helpful suggestions for improving this paper: Anne Jones and Rob Pratt.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors as follows:

Chuck Kelly, Analytical Consultant
SAS Institute
100 SAS Campus Drive, Office T-3052
Cary, NC 27513
Work Phone: 919-531-7186
E-mail: Chuck.Kelly@sas.com

Jinxin Yi, Operations Research Specialist
SAS Institute
100 SAS Campus Drive, Office R-5304
Cary, NC 27513
Work Phone: 919-531-9276
E-mail: Jinxin.Yi@sas.com

Philip Easterling, Analytical Consultant
SAS Institute
4801 Woodway Drive, Suite 300E
Houston, TX 77056
Work Phone: 713-964-2619
E-mail: Philip.Easterling@sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. © indicates USA registration.

Other brand and product names are trademarks of their respective companies.