Workforce Management Strategies in a Disaster Scenario

Andjelka Kelic
Adam Turk
Systems Engineering and Analysis, Sandia National Laboratories,
PO Box 5800, MS 1137, Albuquerque, NM 87185-1137
akelic@sandia.gov, alturk@sandia.gov

Abstract

A model of the repair operations of the voice telecommunications network is used to study labor management strategies under a disaster scenario where the workforce is overwhelmed. The model incorporates overtime and fatigue functions and optimizes the deployment of the workforce based on the cost of the recovery and the time it takes to recover. The analysis shows that the current practices employed in workforce management in a disaster scenario are not optimal and more strategic deployment of that workforce is beneficial.

Keywords: telecommunications, labor, operations, PSTN, disaster, workforce, strategy

Introduction

Detailed models of the voice telecommunications network have been constructed to explore the effects of the loss of particular pieces of the infrastructure on congestion and the ability of users to make voice calls (O'Reilly, et al[2006], Jrad, et al [2005]). Follow on work has included the repair and maintenance functions that keep the network operating and respond in the event of a large scale failure. The repair and maintenance models that were developed are used to explore workforce management strategies in a disaster scenario, such as a hurricane, where large portions of the network are out of service and the workforce is subject to overtime and the associated fatigue.

This paper explores different workforce strategies and their effects on the time and cost of recovering the network in a disaster scenario. The model used in this study was described in detail in "Telecommunications Operations Resiliency: Labor Shortages and the Voice Network" [Kelic, et al., 2007] and is summarized briefly here along with changes that were made to focus on workforce management and disaster recovery.

Overview of Model Structure

The model was modified from the earlier work to isolate the effects of workforce management strategies on the cost and ability of the workforce to recover the network in a disaster scenario. In the course of the normal functioning of the voice telecommunications network, equipment fails at a particular rate such that a small fraction (on order of 0.25%) is damaged at any given time. The failure rate of that equipment is the nominal failure rate that is present in the network. In order to repair failures, replacement material is typically required. For the purposes of isolating

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000

workforce management strategies, all material constraints and nominal failure rates were removed from the model.

The disaster operations model consists of two interconnected systems: the network infrastructure state, and the worker model. One of the key elements of interest in the model is the amount of functioning equipment in the network. The model includes the effects of fatigue on the workforce which results in an increase in repair time and thus an increase in the amount of damage in the network at any given time. The model also tracks the ongoing cost of the workforce in ‘cost units’ where a normal work hour is equivalent to one cost unit and an overtime hour is equivalent to 1.5 cost units.

Modern telecommunications systems comprise a great variety of specialized equipment and skills. We balanced the competing demands for parsimony and accuracy by defining four kinds of critical equipment: switches, frames, transport elements, and local copper loops; and two kinds of repair workers: network operations center (NOC) workers and field technicians.

Workers do not provide repairs for damage in the network; they provide repairs for damage that is recognized in the network. While this distinction may appear trivial, it is precisely the gap between occurrence of damage and recognition of damage that is addressed by the majority of network operations’ center (NOC) support software and processes. Parameters that control the behavior of this gap provide a model of the tools and software platform in use by the repair organization.

Tracking the operating state of the network is the driving goal of the model. Failures are repaired at varying rates depending on the type of network element and the status of the worker resources. Repairs can be performed either remotely or not, depending on the nature of the failure. The network infrastructure state portion of the model represents the state of the network in response to failure and repair.

The worker portion of the model captures the activities of the human resources involved in operations. Workers arrive and leave, are dispatched on tasks, become fatigued as individuals or overloaded as groups. The state of the available workers can limit the rate at which repairs are initiated and completed, and the number of failures in the network impacts the dispatch requirements of the workers.

Details of the model structure can be found in the following sections.

Network Infrastructure

The network infrastructure portion of the model drives the worker dispatch and repair functions. In a normally operating and maintained network, critical components fail at a particular rate. Some of the damage can be repaired by operations center workers (see the variables “front end close time” and “average front end close rate” in Table 1) and other damage requires a field technician. For the purposes of isolating the cost and duration of repair operations for a disaster scenario, routine failures were removed from the model, and an exogenous failure component was added.

After they fail, components become a part of the pool of “damaged infrastructure equipment” and “infrastructure unreported damage.” These failures can only be repaired once they are noticed, either by network operations workers through monitoring equipment, or from customer notification. The rate of damage being noticed is tracked in the variable “damage reporting rate.” Once damage in the network is noticed, it goes from “infrastructure unreported damage” to “perceived damage” and can then be repaired.

The amount of damage in the network determines the workers necessary for the repair. The rate of repair of failures is dependent on the type of component and on the state of the worker. The default values for constants in the model can be found in Table 1 and are further discussed in the section on baseline model runs. The worker segment of the model is described in the next section.

Workers

Failures in the network cannot be repaired until they are noticed and the appropriate resources are dispatched. The worker dispatch portion of the model tracks the human resources associated with repair and is shown in Figure 1. Two different categories of workers are tracked: operations center workers and field technicians.

Operations center workers staff the network operations center and monitor equipment. These workers can fix problems with equipment that are software related – such as resetting a piece of equipment. Operations center workers do not require replacement components on order to solve a problem.

Field technicians travel to the physical site of the piece of equipment and repair physical problems such as splicing a cable or replacing a piece of equipment. If repair material is not available, field technicians will not be dispatched to the site of the problem.

Both categories of workers arrive on shift and then are dispatched to perform tasks. The rate and length of dispatch is dependent on the amount of damage in the network and in the case of field technicians, available replacement components. Workers dispatched for extended periods of time become fatigued and their productivity decreases.

The necessary workers are determined by the “perceived damage”, “workers required per unit of damage”, and the current “repairs in process”. The total number of necessary workers is then compared to the current “workers dispatched” to determine how many additional workers are necessary. The “repair time tolerance” represents how long a piece of the network equipment can wait before being repaired (in addition to the time it takes to repair that piece of equipment). Normally, field technicians require materials to conduct repairs and will not be dispatched if the material is not available. However, to isolate the workforce portion of the repair process, materials constraints were removed from the model. Operations center workers do not require material to perform their repairs, so will always be dispatched, even in the event of materials shortages. Operations center workers do not travel to other locations when they are dispatched; they are simply assigned to a problem.

Workers move through four states in the model as shown in Figure 1. They begin off-shift, go on-shift ("workers available"), are dispatched ("workers dispatched") and then either go off-shift again or return to available status depending on the length of the repair. Movement among states is governed by work schedules, workload, and fatigue. The model does not currently represent hiring additional workers.

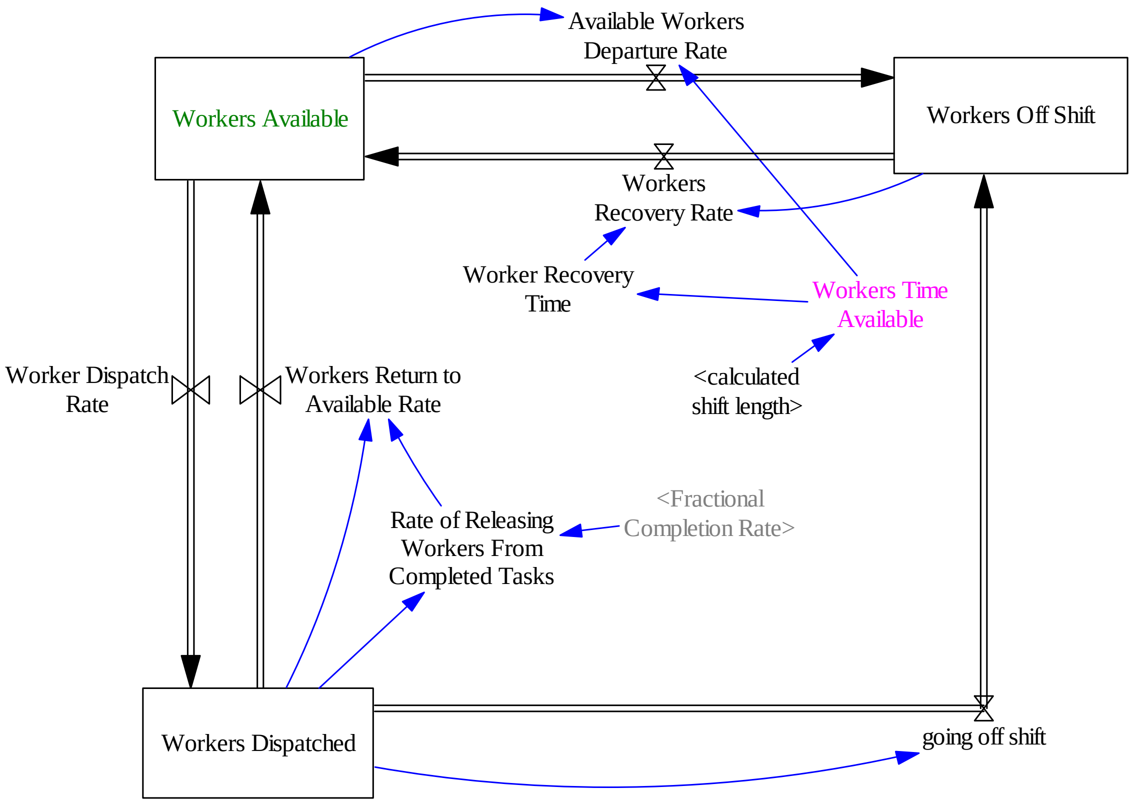

Figure 1: Worker Dispatch Model Component

Overtime and Fatigue

The model assigns workers to overtime when the amount of damage in the network increases beyond what can be repaired within the repair time tolerance by the total pool of workers as shown in Figure 2. Even when there is outstanding damage, workers will continue at the nominal rate of repair when the total level of damage in the network is considered to be at normal levels of routine damage. As the amount of damage begins to exceed routine, workers will begin to extend their normal shift with overtime. Under severe conditions, workers may double their regular shift up to a sixteen hour shift to maintain network health; longer shifts are prohibited due to typical labor regulations in the United States.

As the worker shift length increases, and the duration of extended shifts increases, the workers become fatigued. For example, a worker can work a double shift occasionally and productivity will not suffer. However, if workers are continually working overtime, even for a few extra hours every day, their productivity will begin to suffer and repairs will take longer. Fatigue is a delayed degradation of worker effectiveness, and creates a positive feedback to the demand for additional dispatched workers. This formulation is similar to [Hines, 2005], with the fatigue function developed in Sterman [Sterman, 2000] and additional parameters from Oliva [Oliva, 1996].

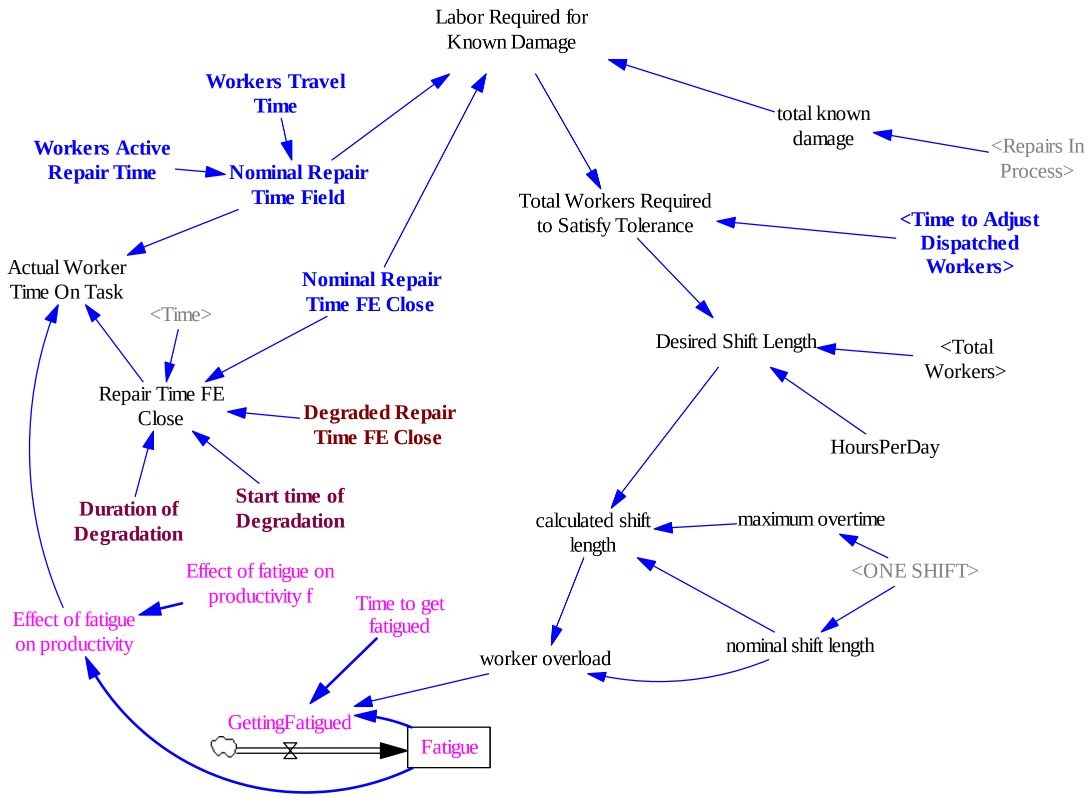

Figure 2: Fatigue, Overtime, and Productivity

Simulations: Scenarios and Results

Default Model Parameters

The steady state expressed in this model is intended to capture a mid-sized metro area under normal, non-disaster, conditions. Since the workforce operations portion of this model has been isolated to explore disaster response, the steady state has workers becoming available and never dispatching because there is nothing to repair.

To simulate a disaster scenario, 1 million lines are taken out of service over a 24 hour period in the first week of the simulation. This is done by damaging loop elements (the physical line from the subscriber to the central office) which are the elements that are typically damaged in heavy wind conditions as they or the poles they are attached to are blown down. This is about the same number of BellSouth's lines, the region's largest carrier, lost during Hurricane Katrina [BellSouth Press Release]. The key data values shown in Table 1 were used as the baseline.

Table 1: Default Model Parameters

PARAMETER	VALUE
Size of metropolitan area	6,000,000 subscriber lines, including: Business, Residence, Redundancy, Overbuild
Size of Central Office equipment	6000 one thousand port cards
Average repair time	▪ Switch: 0.5 hour per line unit ▪ Frame: 0.5 hour per line unit ▪ Transport: 4 hours per cable break ▪ Loop: 0.75 hour per residential repair
Average travel time between field sites	▪ 2 hours: transport sites ▪ .75 hour: loop (residential) sites
Average Front end close time	20 minutes for switch and frame components 15 minutes for transport components 9 minutes for loop components
Average Front end close rate	25% for switch and frame components 10 % for transport components 40% for loop components
Number of repair workers per damage report	1 per switch and loop 2 per frame 4 per transport report (cable break)

In a normal disaster scenario, workers are brought in from surrounding areas to help augment the workforce in place. This additional workforce is not enough to repair the damage in a timely fashion using an eight hour work day, so they are also subjected to overtime and fatigue. Thus in this model we neglect the ability to bring in additional workers and simply focus on workforce management strategies to most efficiently repair the damage with the pre-existing workforce.

Baseline Scenario Results

The overall result for damage levels in the network is shown in Figure 3 when the model is run with the baseline parameters listed above under the disaster scenario. As can be seen in the figure, the damage occurs in the first week of the model run, then the workforce begins slowly working off that damage. The network is returned to normal in week 42 of the simulation.

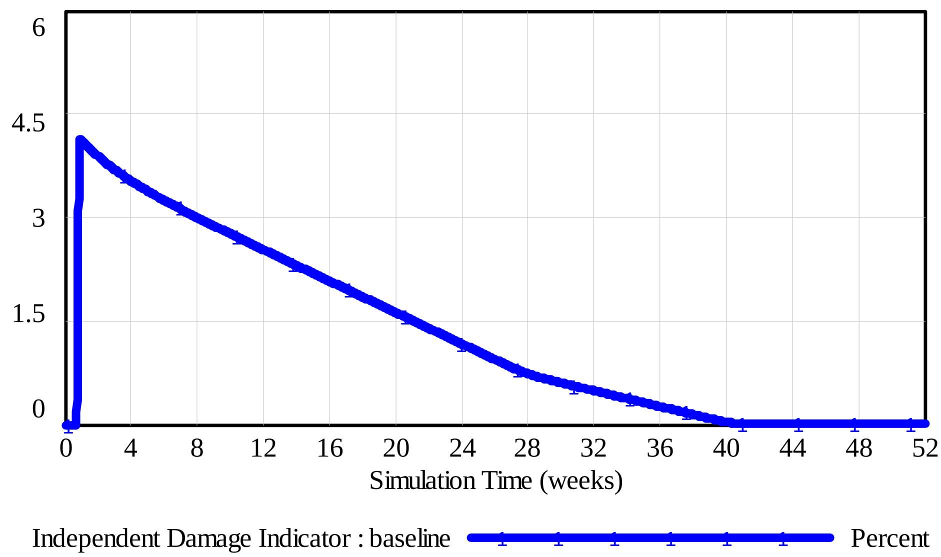

Figure 3: Results of Baseline Scenario Model – Percent Network Damaged

Workers are dispatched as show in Figure 4 to deal with the damage. The number of field workers well exceeds the number of operations center workers required, as it does in reality. The number of workers dispatched to perform repairs reaches equilibrium quickly and continues until the network is returned to an operational state.

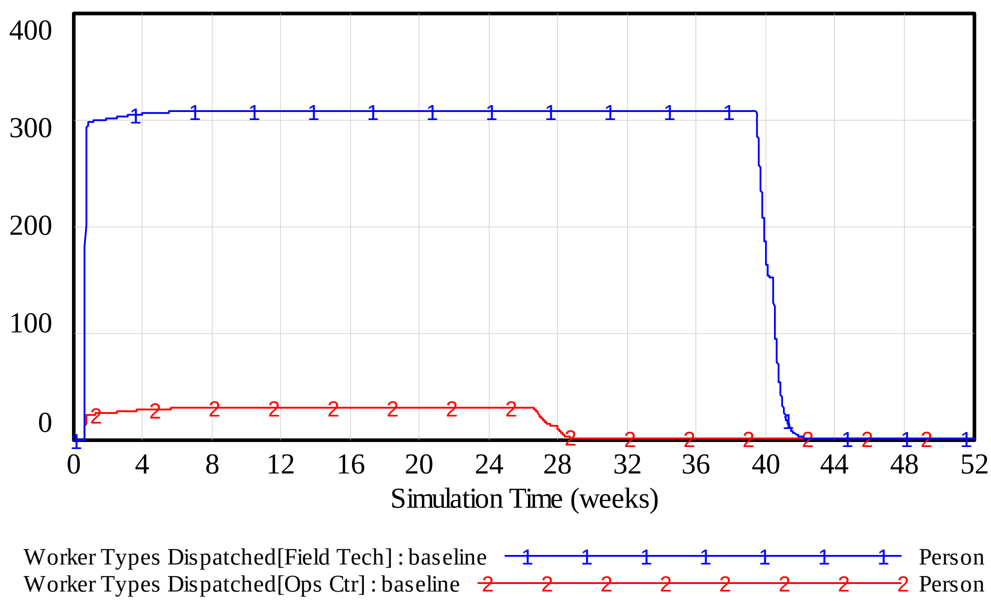

Figure 4: Results of Baseline Scenario Model – Workers Dispatched

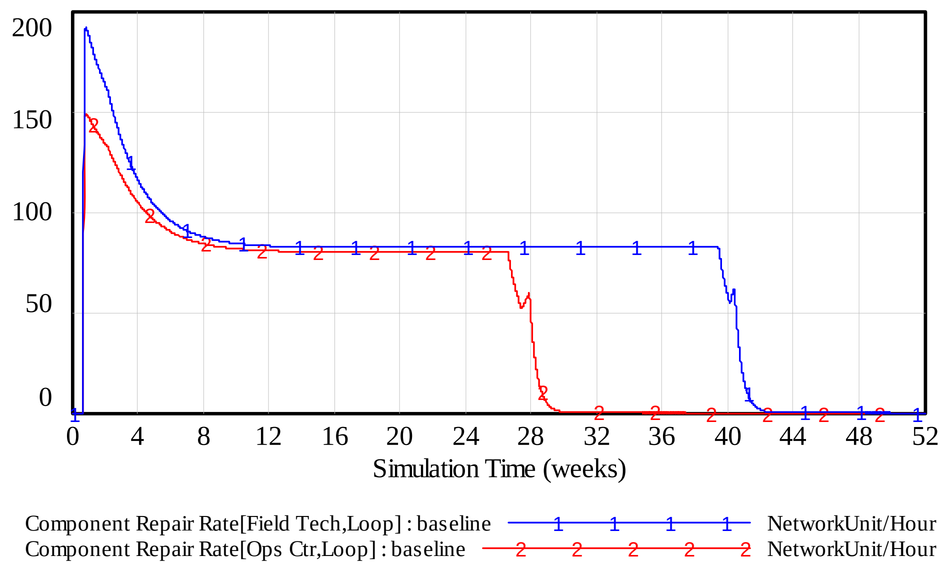

Figure 5: Results of Baseline Scenario Model – Component Repair Rate

The component repair rate show in Figure 5 governs the rate at which damage in the network decreases. The figure shows that initially repair rates are high and then fall by more than a half over a short period of time. This drop off in the repair rate is due to the onset of fatigue which is shown in Figure 6. The onset of fatigue causes the workforce to be only 40% as productive as they would otherwise be, significantly delaying the restoration of the network, despite the longer hours (a double shift, 16 hours) being worked by the workforce.

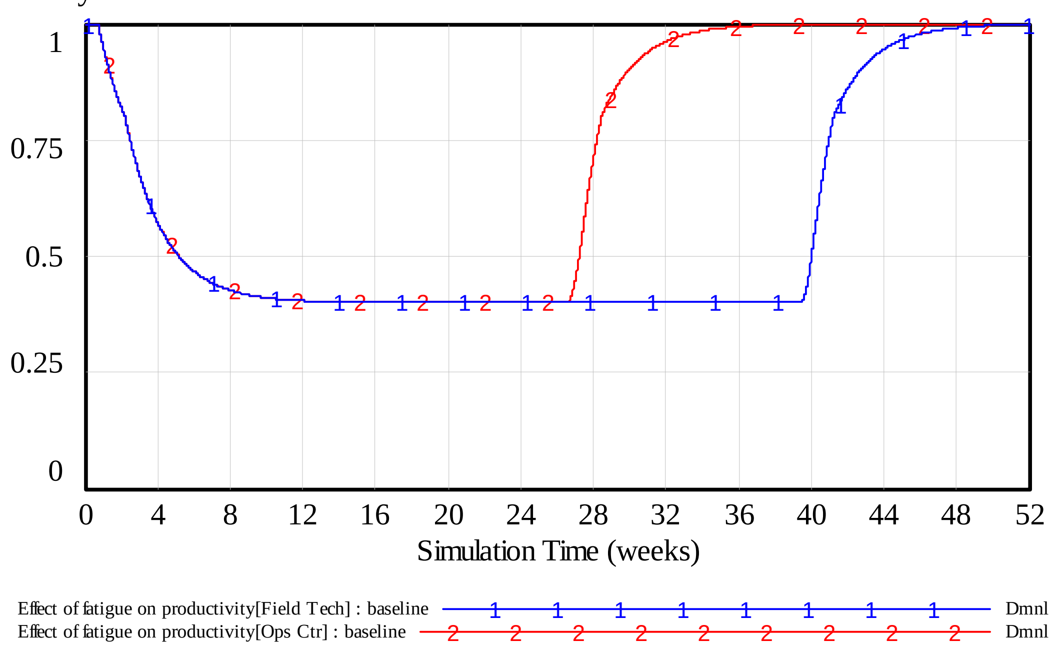

Figure 6: Results of Baseline Scenario Model – Fatigue Effect on Productivity

The total cost of this repair is just under 2 million "cost units", where a cost unit is one for each person-hour worked in a normal shift, and 1.5 for each person-hour worked in overtime. The model allows for different charging rates to be used for field technicians and operations centers employees. Detailed research into the cost of these personnel was not performed for this analysis, thus the model assumes that their costs are the same and normalizes to cost units

.

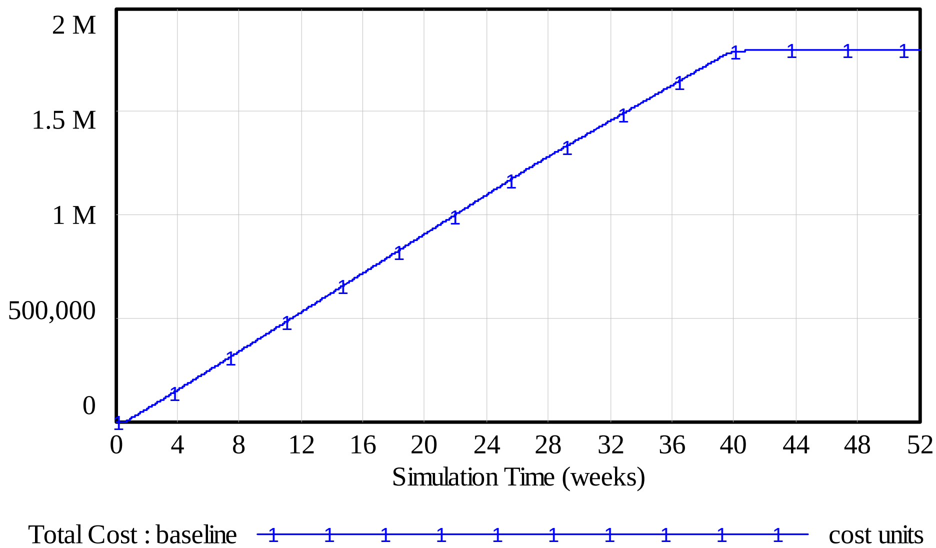

Figure 7: Results of Baseline Scenario Model – Total Cost of Repair in Cost Units

Workforce Strategies and Optimization

The telecommunication operations model was simulated with 10 different maximum shift values ranging between 8.8 and 16.0 hours. The results from these simulations are shown in the following two figures. Figure 8 shows the time period required to repair the damage caused by the disruption plotted against the shift length. This curve gives a minimum value of approximately 145 days for the repairs if shift length was limited to 12 hours. This minimum point corresponds to the middle point of the optimal productivity range identified by Sterman (Sterman, 2000) of 60 hours worked per week.

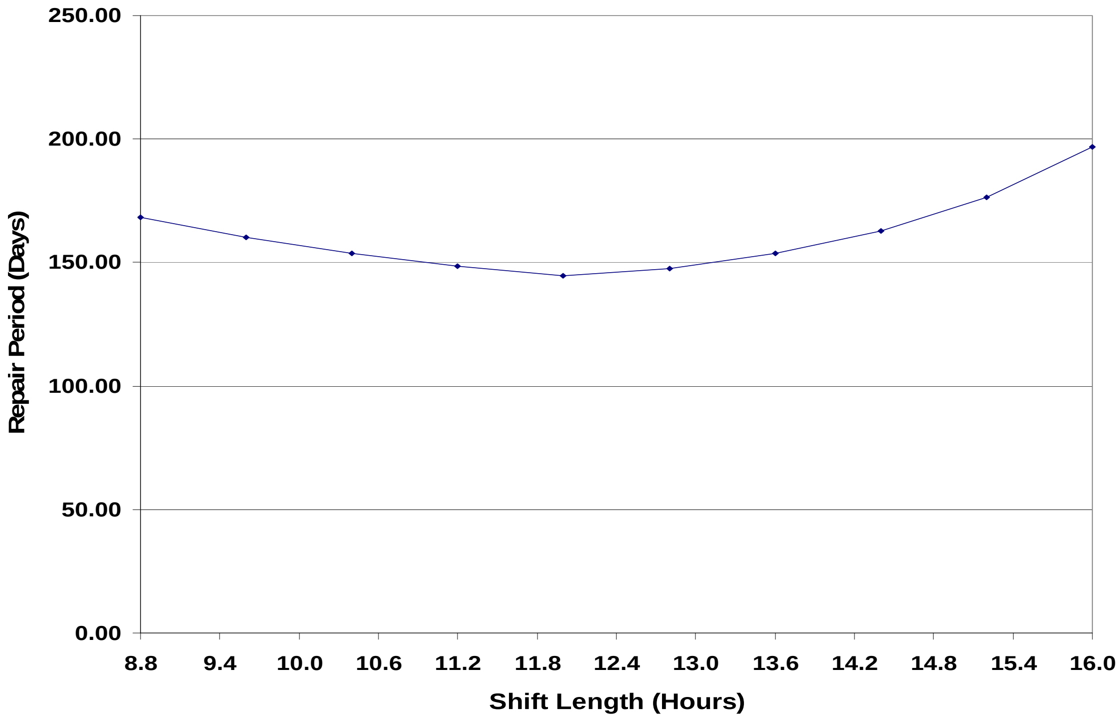

Figure 8: Repair Time Duration versus Shift Length

The model was modified to include a repair time optimization function in order to confirm the optimal shift length value of 12 hours. The maximum shift length was selected as the optimization variable or parameter. The results from the modified model with the repair time minimized are shown in Figure 9. This figure shows that the optimal shift length is not a constant 12 hours but one that assumes a high shift length spike of 16 hours at the beginning and end of the repair period. These spikes are caused by the model starting with the repair productivity at its maximum value which is achieved with a shift length of 16 hours. After a certain period of time, fatigue effects increase and repair productivity starts to decrease. The shift length ramps down to balance between fatigue and productivity for shortest repair. Towards the end of the repair period, the shift length increases again to 16 hours since dispatched workers will be able to rest before fatigue impacts repair productivity. Figure 9 also shows that a shift length less than 16 hours will shorten the repair period when compared with current industry practices.

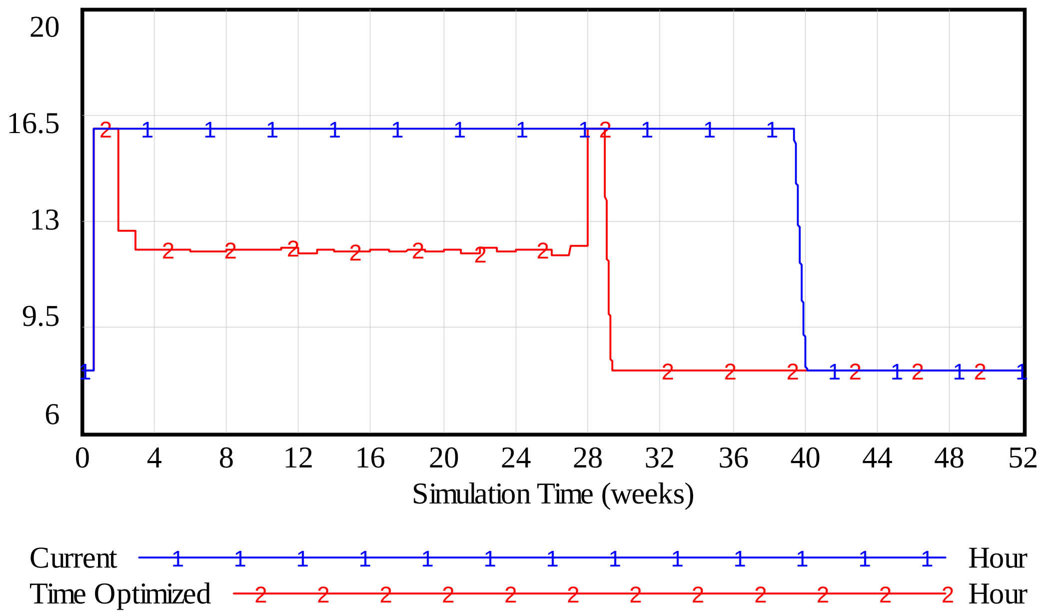

Figure 9: Repair Time Optimized Dispatch Work

Figure 10 shows the costs of workers during the repair period plotted against shift length. The curve shows that these costs grow exponentially with shift length and that the optimal value is a shift length of 8 hours.

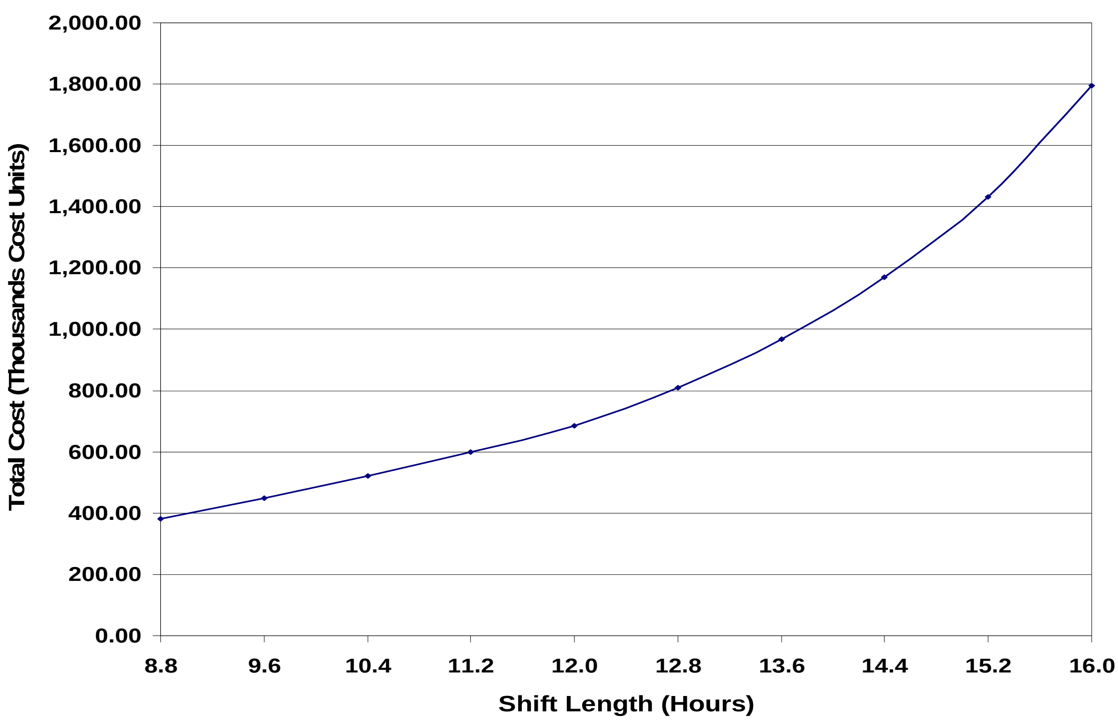

Figure 10: Total Worker Cost versus Shift Length

The model was also modified to include a cost optimization function in order to confirm the optimal shift length value of 8 hours. Similarly to the earlier repair time function, the maximum shift length was chosen as the optimization parameter. The results from the modified model with the cost optimization function are shown in Figure 11.

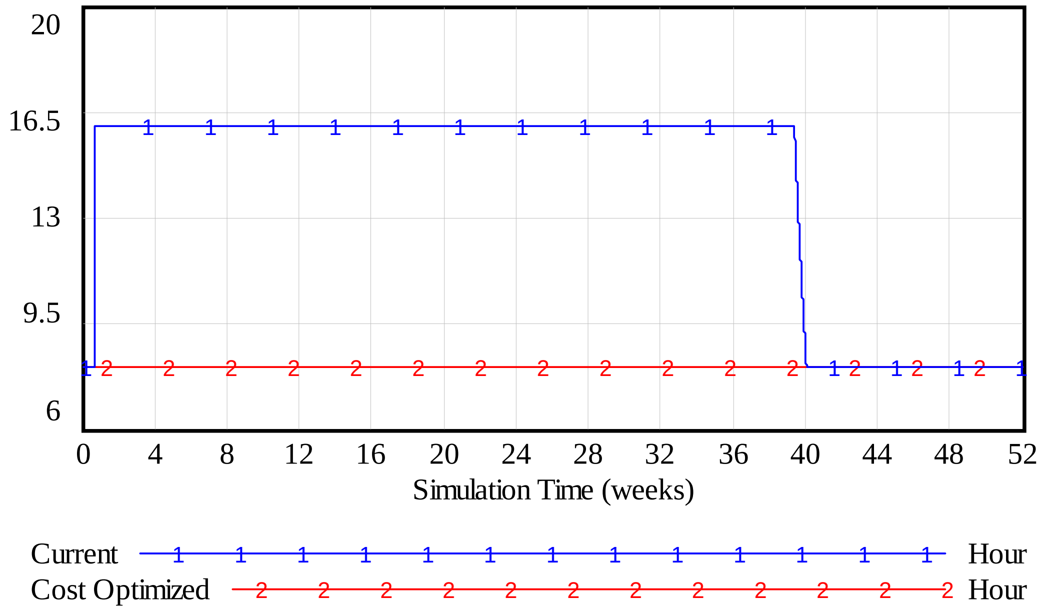

Figure 11: Cost Optimized Worker Shift Length

In order to minimize costs, shift time would be held constant at 8 hours, but that policy would cause the repair period to remain high. The shift length would have to be increased to decrease the repair period and would in turn raise the worker costs. These opposing goals hint that a balance may exist between repair period and worker cost. In order to investigate this balance, the model was further modified combining the cost and repair time functions together. Each function was weighted so that their current optimal values for shift time produced the same value from the combined function.

Figure 12 shows the results of the combined cost-repair time optimization along with current practice and the repair time optimization shift length profile. The combined cost-repair time optimization shift length profile shows the same characteristics as the repair time optimization. It has spikes at the beginning and end of the repair period, but at lower shift lengths and a longer repair period.

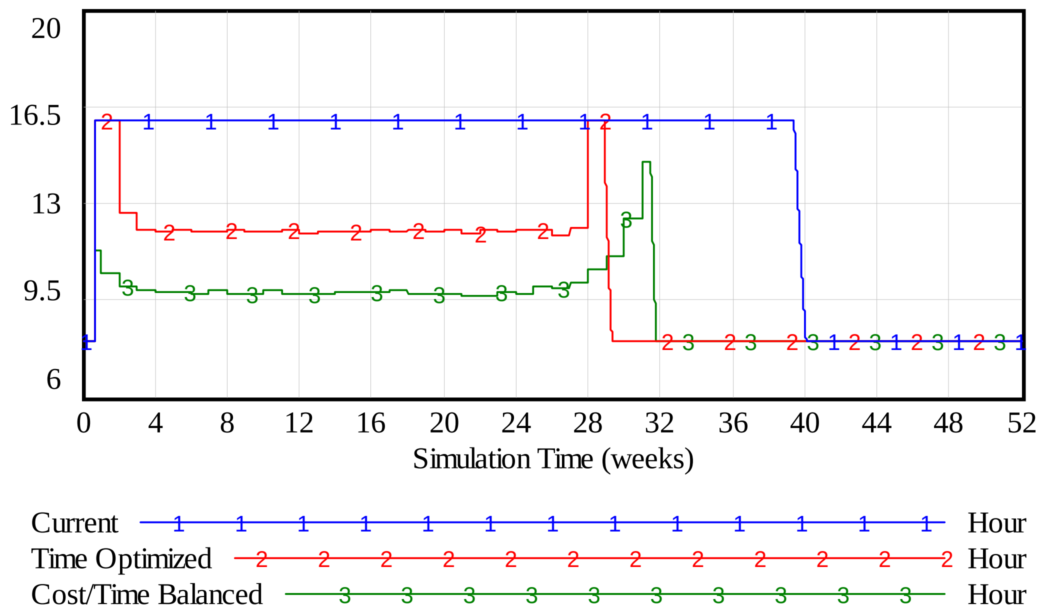

Figure 12: Balanced Cost and Repair Time Optimized Shift Length

The next two figures compare worker costs and repair times for the different optimization objectives and current practice, respectively. Figure 13 demonstrates that profiles with shift length less than 16 hours, which is the current industry practice, will result in a lower worker cost.

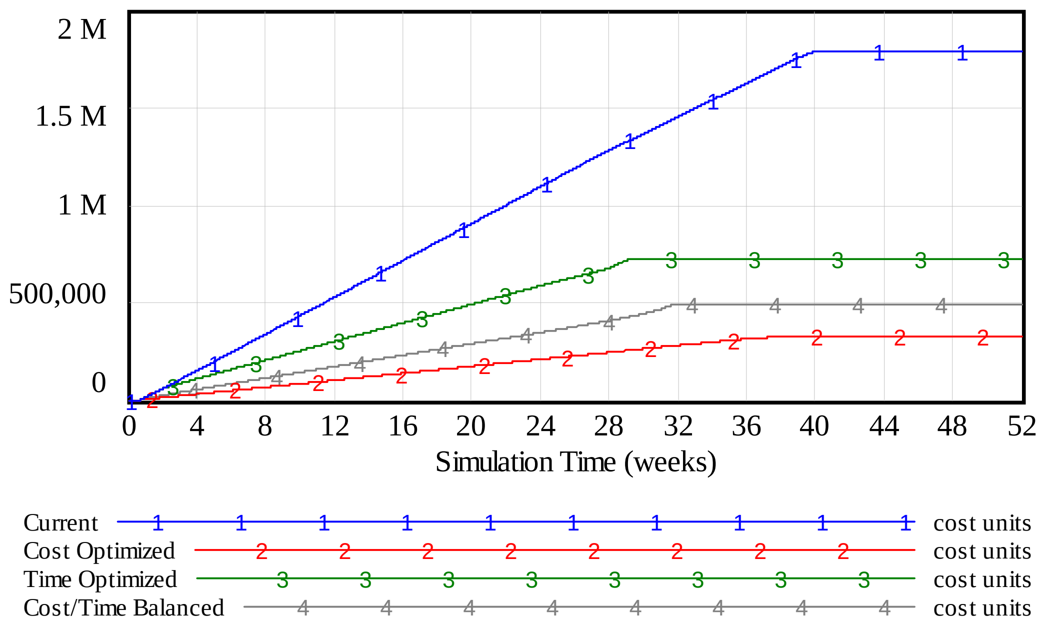

Figure 13: Worker Costs for All Optimization Objectives and Current Industry Practice

Similarly, Figure 14 demonstrates that profiles with shift length less than 16 hours, which is the current industry practice, will result in a shorter repair time or period.

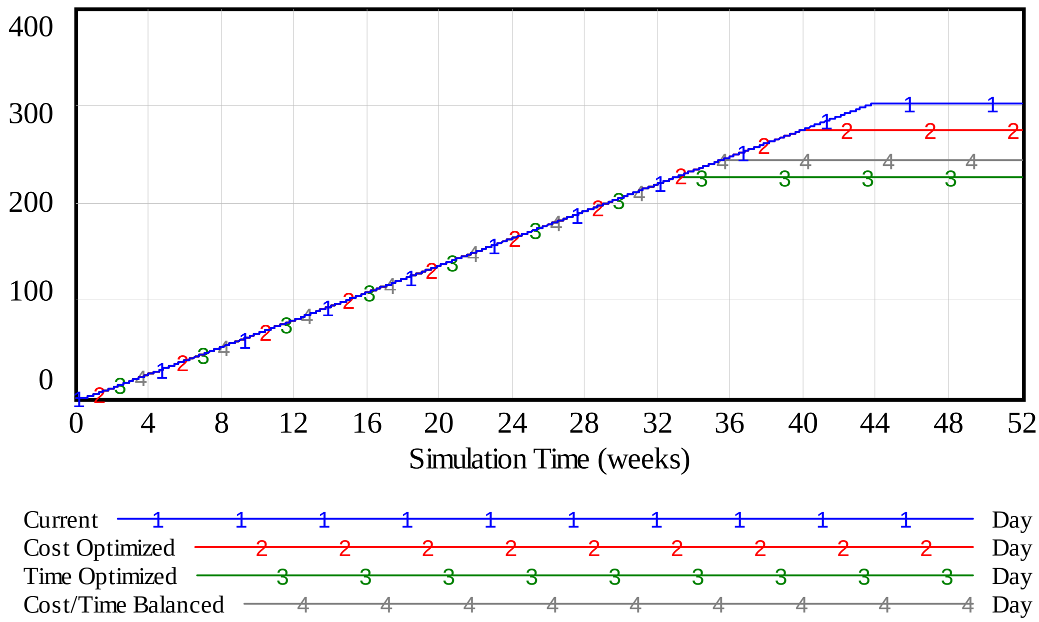

Figure 14: Repair Time for All Optimization Objectives and Current Industry Practice

These figures show that the minimizing costs and minimizing repair times are not entirely at odds with one another when compared to current industry practice.

Conclusions

These last two figures demonstrate that any of the optimization profiles achieve better worker costs and repair times than current industry operations. In addition, these figures show that the combined cost-repair time objective function does provide a good balance between these two goals. These different optimization profiles can be accomplished through additional mechanisms beyond limiting the worker shift lengths such as bringing fresh workers in staged groups or rotating groups as fatigue increases. Staged worker groups will delay the onset of fatigue while rotating groups will allow worker to rest reducing their fatigue.

In the end, a shift length shorter than 16 hours achieved by any mechanism resulted in better workforce management in terms of worker costs and repair times than current operations.

Future Work

During the course of this work, areas that needed additional focus and research were discovered. The current fatigue function does account for several conditions that would affect worker fatigue. First, workers will likely fatigue faster when the working environment is stressful and in many cases unknown, such as in a disaster recovery.

Extended period of long shift lengths do not increase fatigue levels and further deteriorate productivity as they would in a real long-term overtime scenario.

Beyond augmenting the model, the results and model assumptions need to be reviewed with companies and their field workers for the telecommunications industry and potentially other infrastructures to assess the validity and applicability of the model and its assumptions.

References

A. Jrad, H. Uzunalioglu, D. Houck, G. O'Reilly, S. Conrad, W. Beyeler. "Wireless and Wireline Network Interactions in Disaster Scenarios," MILCOM 2005.

BellSouth Press Release, "BellSouth crews work to assess damages, restore service as quickly as possible in Louisiana," 1 September 2005, http://bellsouth.mediaroom.com/index.php?o=100&s=press_releases&year=2005.

Gerard O'Reilly, Ahmad Jrad, Theresa Brown, Stephen Conrad, "Critical Infrastructure Analysis of Telecom for Natural Disasters," Networks 2006, November 2006.

Hines, Jim "Molecules of Structure: Building Blocks for System Dynamics Models," from www.vensim.com, July 2007.

Kelic, A, Matthews, E., and Beyeler, W. "Telecommunications Operations Resiliency: Labor Shortages and the Voice Network" The 2007 International Conference of the System Dynamics Society and 50th Anniversary Celebration, Boston, July 29-August 2, 2007.

Oliva Pue, Rogelio. "A dynamic theory of service delivery: implications for managing service quality," Massachusetts Institute of Technology Sloan School of Management, PhD, 1996.

Sterman, John. Business Dynamics. Boston: McGraw-Hill, 2000.