Telecommunications Operations Resiliency: Labor Shortages and the Voice Network

Andjelka Kelic[1]
E. P. Mathews[2]
Walt Beyeler[1]
[1]Critical Infrastructure Modeling & Simulation, Sandia National Laboratories,
PO Box 5800, MS 1137, Albuquerque, NM 87185-1137
[2] Bell Laboratories, Alcatel-Lucent, 600 Mountain Ave, Murray Hill, NJ 07974
Tel: (505) 844-5266 / Fax: (505) 284-3850
akelic@sandia.gov, epmatthews@alcatel-lucent.com, webeyel@sandia.gov

Abstract

Models of the voice telecommunications infrastructure have focused on the availability of the network during a disruption without accounting for the workforce necessary to provide repair and recovery functions for that network. This paper describes a system dynamics model of the maintenance operations of the voice telecommunications infrastructure and explores the effects of large and prolonged worker absence on the ability to keep the infrastructure operating. Analysis shows that the voice telecommunications infrastructure is highly resilient to the loss of a large portion of its workforce.

Keywords: telecommunications, labor, operations, PSTN

Introduction

The National Infrastructure Simulation and Analysis Center (NISAC)[1] is intended to provide DHS with an effective understanding of the performance of fourteen critical infrastructures, especially under extraordinary circumstances. NISAC draws on industry experience and knowledge regarding individual infrastructures, developing new models to represent the processes and cross-infrastructure interactions that may precipitate or control disruptions.

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]). These models have not included the repair and maintenance functions that keep the network operating and respond in the event of a large scale failure. The rate of restoring the

---

[1] The National Infrastructure Simulation and Analysis Center is a joint program at Sandia National Laboratories and Los Alamos National Laboratory, funded and managed by the Department of Homeland Security's (DHS) Preparedness Directorate.

system is controlled by the rate of moving, installing, and testing equipment. In order to estimate outage duration those processes must be represented.

In September 2005 Sandia National Labs and Bell Laboratories, working in partnership, began developing a model of telephony operations that included human resources, warehouse replenishment, and the driving forces of exceptional network damage and routine maintenance and repair. By the end of the year, the Telephony Operations model had been designed and the initial version built. Simulations were run, including a baseline and several disaster scenarios, to validate the model and to understand its weaknesses. This paper documents the model, baseline data, model structure, assumptions and the effect that absenteeism has on repair and recovery.

Overview of Model Structure

The Telecom Operations model consists of three interconnected systems: the network infrastructure state, workers, and the warehouse models. Figure 1 and Figure 2 show the governing loop diagrams represented in the model. One of the key elements of interest in the model is the amount of functioning equipment in the network; this is shown in Figure 1. This figure shows how equipment failure affects the perceived damage in the network, and how material is used in the repair process. Figure 2 shows the causal loops that govern the response of the workforce to damage. 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.

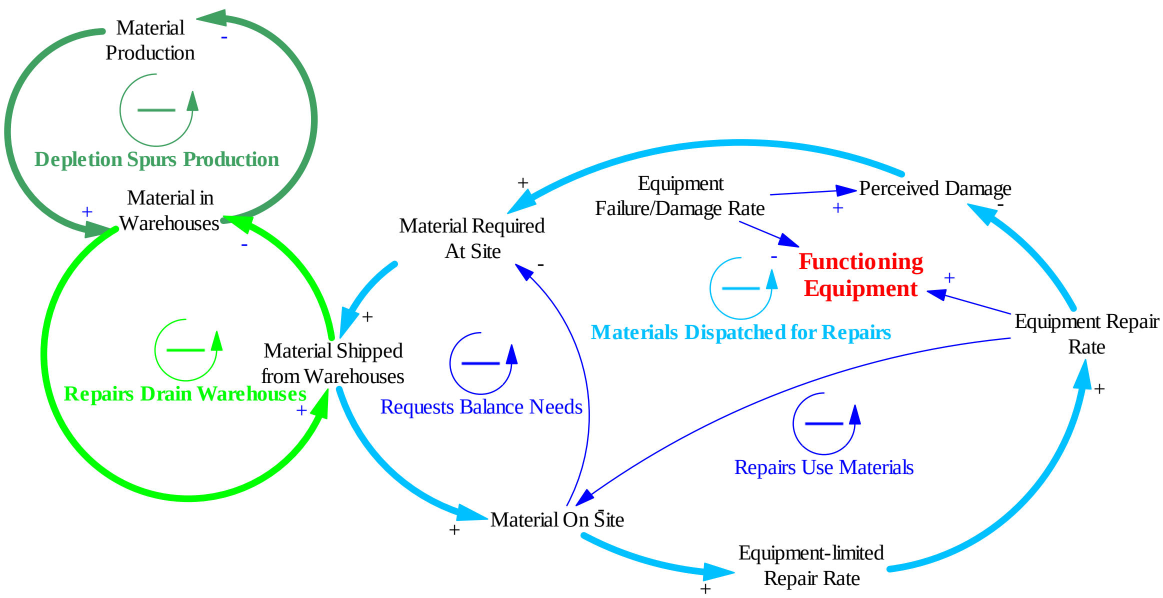

Figure 1: Network Failure Rates and Material Supply Overview

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. The similar structure of the material and worker flows allows us to use arrays to manage these distinctions.

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.

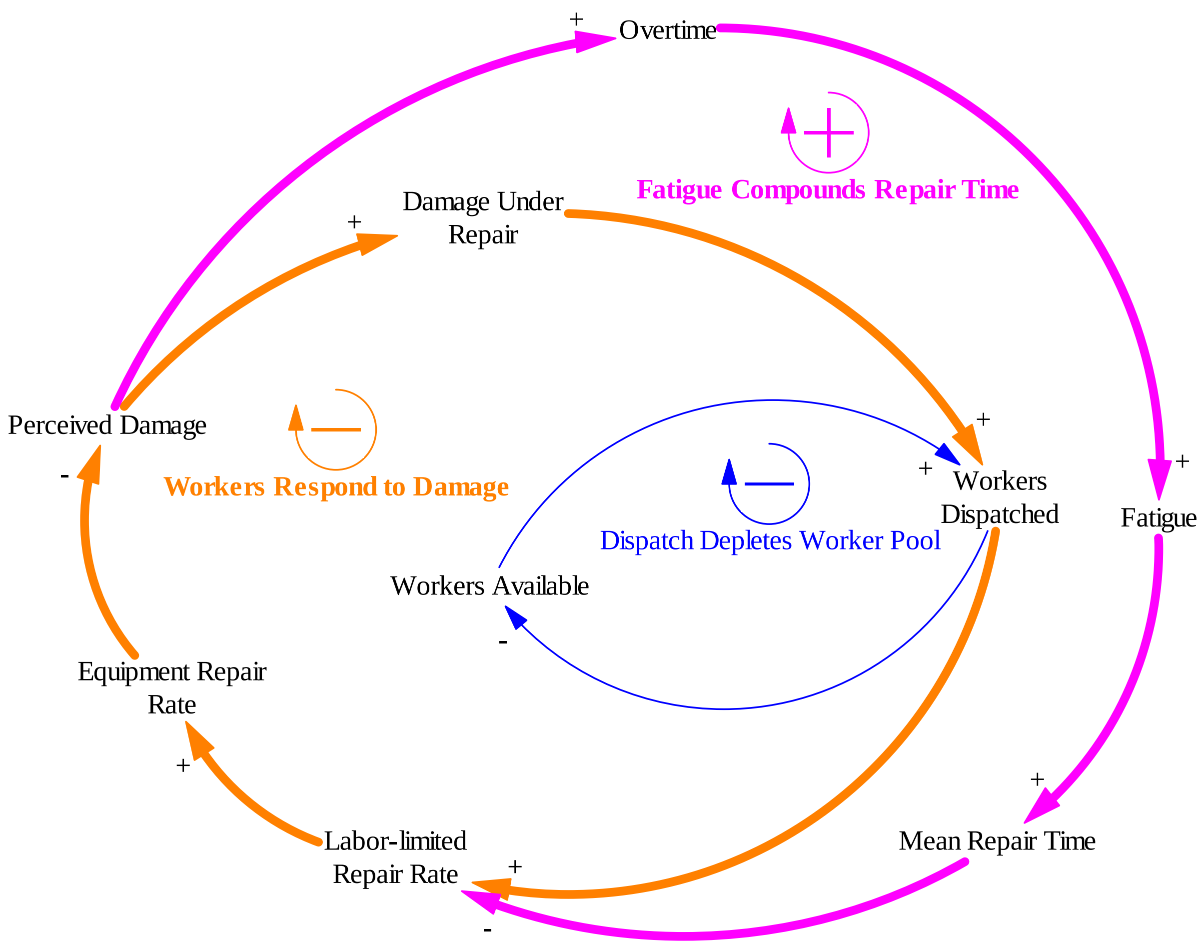

Figure 2: Worker Dispatch and Repair Overview

Tracking the operating state of the network is the driving goal of the model. Failures occur in the network at varying rates, and are repaired at varying rates. These rates depend on the type of network element, the status of the worker resources and the status of the materials to be delivered from a warehouse. In a disaster scenario, extraordinary failures occur in addition to ongoing "business as usual" failures. Repairs can be performed either remotely or not, depending on the nature of the failure, and certain repairs are performed in two stages ("initial patch" and "repair"). 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. In a disaster scenario, worker presence may apply different assumptions, overriding "business as usual" process and worker availability may be severely altered (as during a snowstorm or epidemic). 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.

The activity of a human worker making a repair on a network element will require the delivery of a spare or repair component from the warehouse. The warehouse, as it delivers spare parts, must replenish its supply from its factory sources without over-ordering or running short. In the event of extraordinary demand, as during a disaster, the warehouse delivery processes will throttle the rate at which repairs can be made, and may cause repairs to be effected in non-optimal order, or cause worker resources to sit idle, waiting for parts.

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 and is shown in Figure 3. In a normally operating and maintained network, critical components fail at a particular rate, governed by the “component damage rate” variable. 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.

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 and material necessary for the repair. The rate of repair of failures is dependent on the type of component and on the state of worker and replacement component resources. 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.

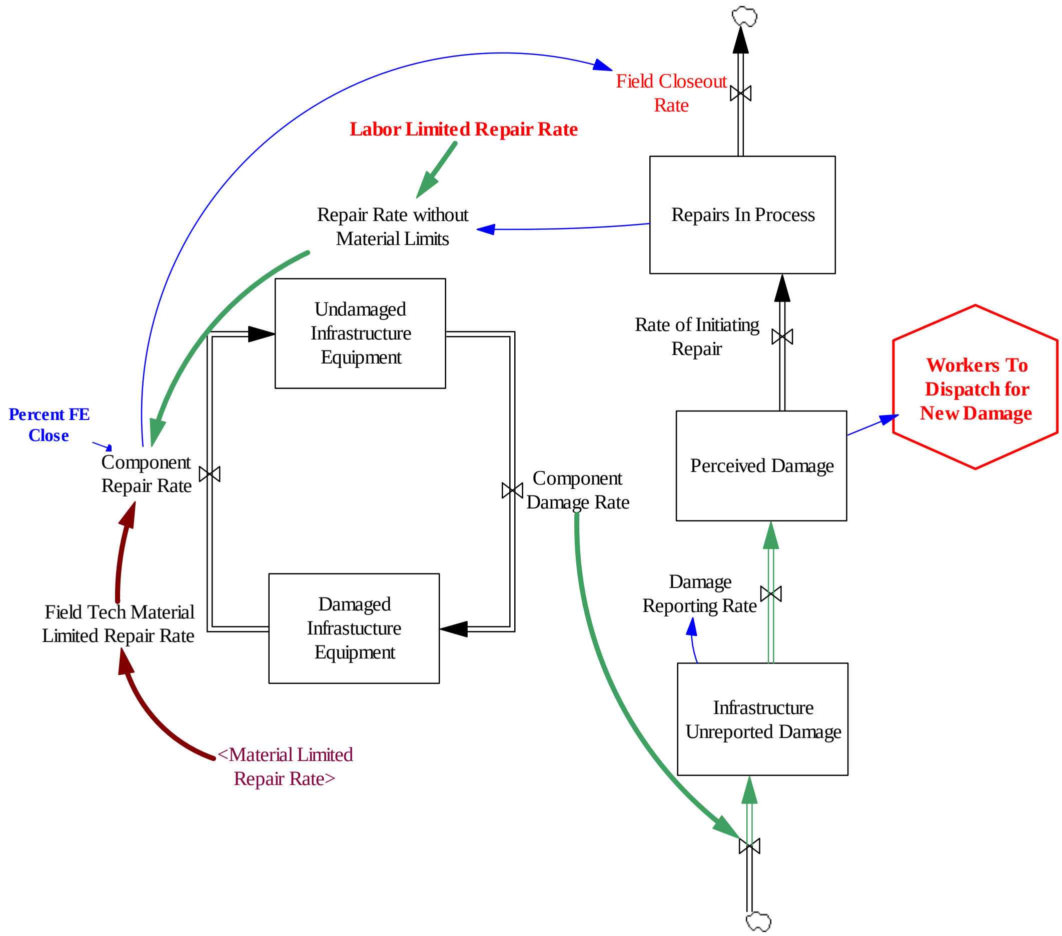

Figure 3: Network Infrastructure Model Component

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 4 and Figure 5. 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. This portion of the model allows for temporary removal of personnel from the workforce due to absence - either in small volume due to ordinary illness or large volume due to epidemic or emergency conditions. The parts of the model governing the removal will be discussed in the section on worker absence.

The portion of the model shown in Figure 4 calculates the number of workers that need to be dispatched to complete repairs in the desired amount of time. 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 can wait before being repaired (in addition to the time it takes to repair that piece of equipment). Field technicians require materials to conduct repairs and will not be dispatched if the material is not available. 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.

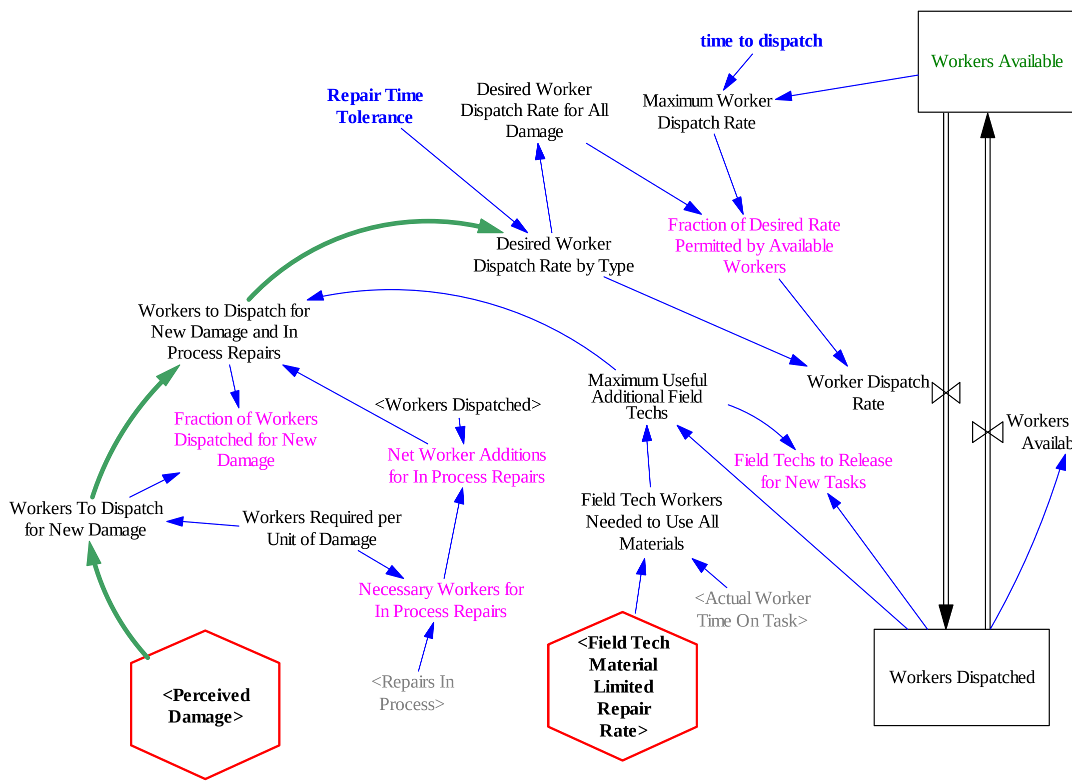

Figure 4: Worker Dispatch Calculations

Workers move through four states in the model as shown in Figure 5. 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 extraordinary event portions of the model are designed to remove workers from the pool of off-shift workers due to absenteeism. The model does not currently represent hiring additional workers.

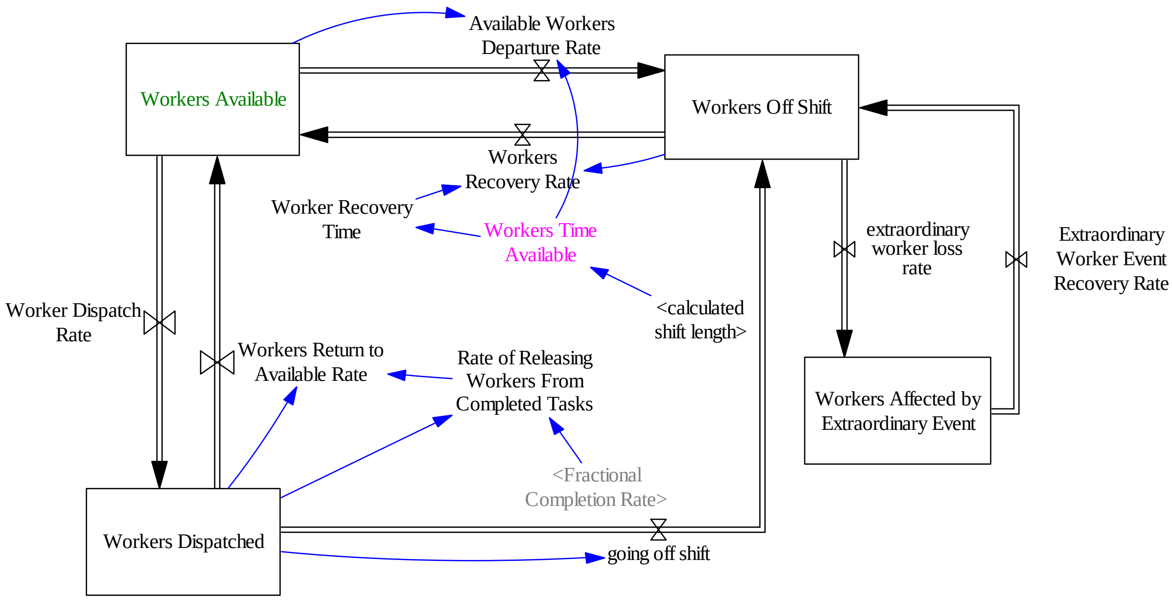

Figure 5: Worker Dispatch Model Component

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 6. 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], however the specific function for calculating worker fatigue based on overtime needs to be further researched and replaced with one from industry data or what is consider to be standard for the type of repair work performed by the field and operations center workers.

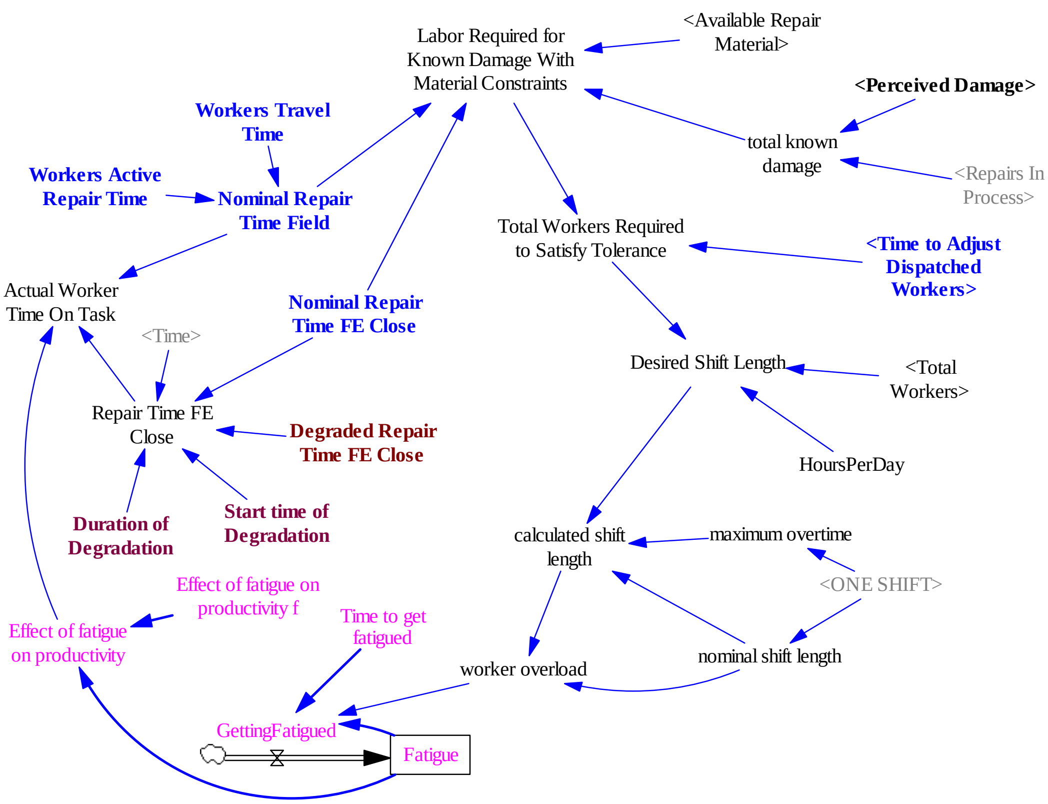

Figure 6: Fatigue, Overtime, and Productivity

Warehouse

The warehouse portion of the model tracks the availability of repair components from the supplier as shown in Figure 7. Both the supplier and the warehouse have desired levels of inventory, and produce or order supplies to maintain that desired level while being able to cover incoming requests. The formulation follows Sterman[2000].

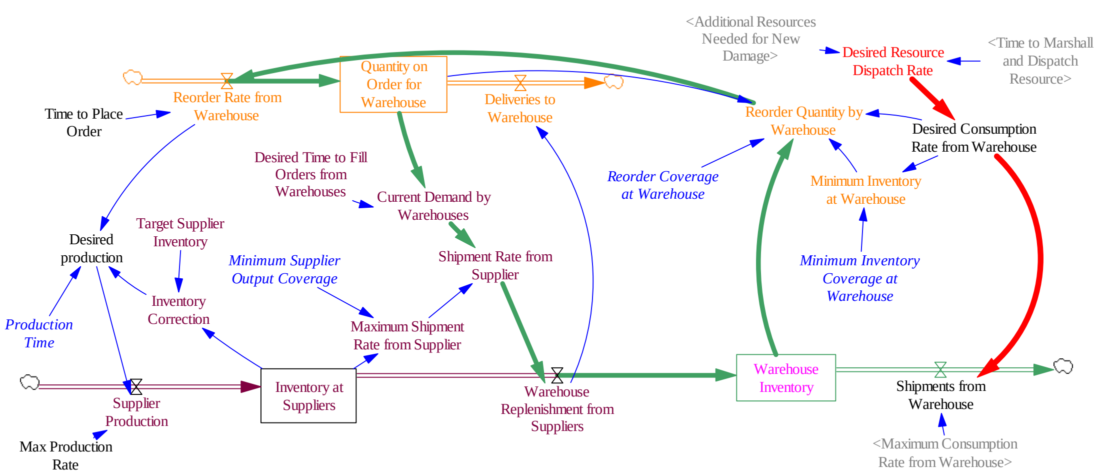

Figure 7: Warehouse Ordering and Inventory

Figure 8 shows the piece of the model that ensures that materials are on site for repair. If the materials are not available, workers will not be dispatched to perform the repair. Thus a disruption in the supply chain or an inventory shortage could cause degradation in the state of the network.

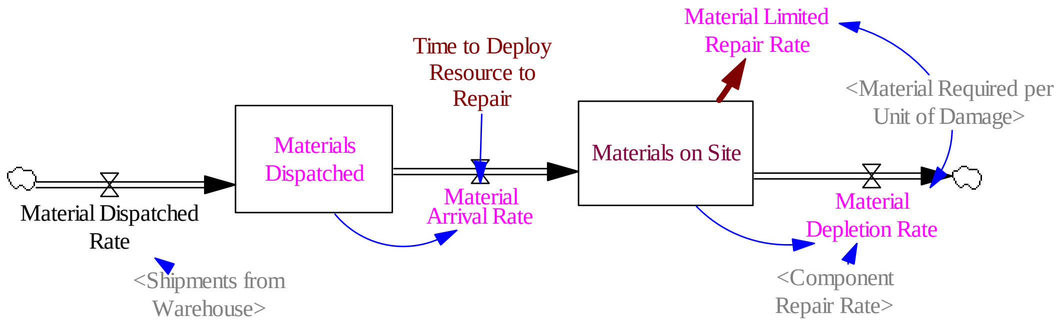

Figure 8: Warehouse and Material Dispatch Model Component

Simulations: Scenarios and Results

Overview of Approach for Baseline Scenario

We have exercised the integrated model to test its behavior under ordinary and disrupted conditions. Below we describe the results for a configuration that approximates a generic mid-sized metro area. Ordinary operations form a baseline on which two kinds of workforce disruption are imposed: absenteeism and illness. The results demonstrate the wide range in behavior that the model dynamics can produce, and the way it can be used to identify constraints on restoration time and effectiveness of mitigating measures under diverse disruptions. Several of the model parameters are notional, and the model has not yet undergone extensive testing. The results give a concrete illustration of the model scope and capabilities, but should not be interpreted as an analysis of a real system.

Baseline Scenario Parameters

The steady state expressed in this model is intended to capture a mid-sized metro area under normal, non-disaster, conditions. This is the scenario used as the baseline against which to isolate and analyze perturbations created by other scenario conditions. We analyzed and included historical data to the greatest degree possible, as allowed by the level of detail of the model. Under these conditions, with accurate initial values for all key variables, the system quickly achieves equilibrium in all areas of interest. The following key data values shown in Table 1 were used as the baseline view:

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)

Table 1: Default Model Parameters

Baseline Scenario Results

When the model was run using the above parameters, which were chosen to reflect the information available from actual telecom service providers, the resulting failure rates corresponded closely with actual failure rates observed in the network. As shown in Figure 9, the network is approximately 0.25% damaged at any given time in the course of normal functioning. This means the total number of network elements in the model that are out at any one time will be close to 15,000; actual telecom operations information shows that this is reasonable for a city with six million lines. Given the lag time between damage occurring in the network and the actual reporting of damage so that it can be repaired, the perceived damage in the network is slightly smaller than the actual damage levels. Figure 10 shows that the repair rate is steady and that loop repair makes up the largest portion of the repair rate, and of the damage.

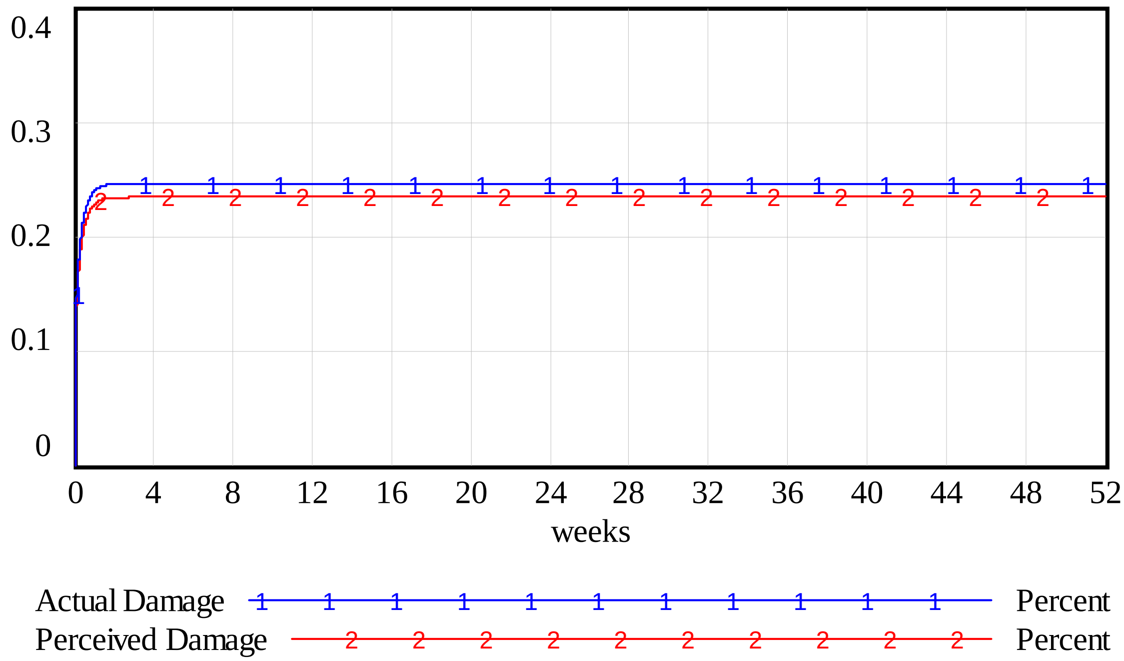

Figure 9: Results of Baseline Scenario Model – Actual and Perceived Damage

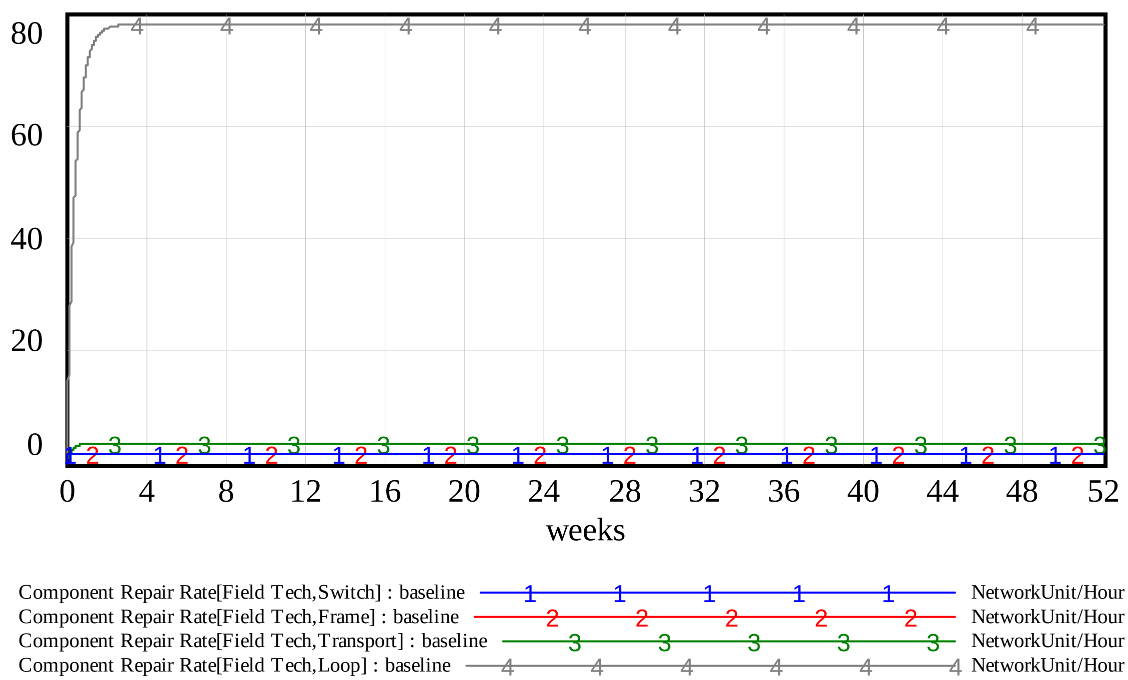

Figure 10: Results of Baseline Scenario Model – Component Rate of Repair

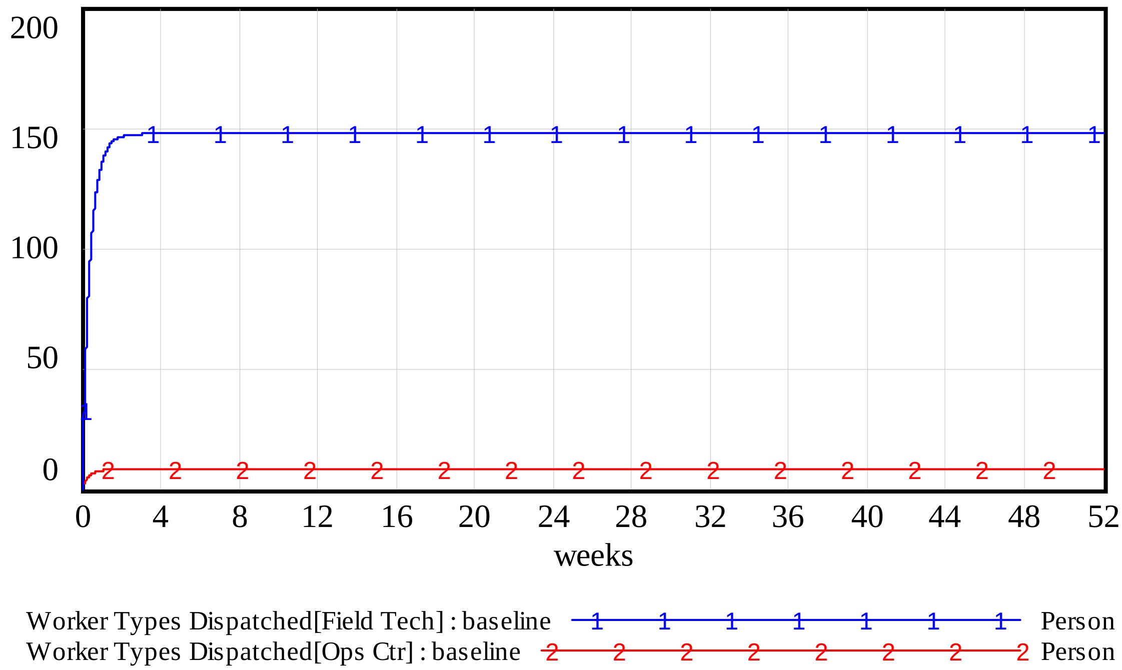

Figure 11: Results of Baseline Scenario Model – Workers Dispatched

The results from the baseline simulation correspond to telephony operations in the real world: The rate of repair reaches equilibrium very quickly and remains stable, with a small amount of ongoing network damage, as time is taken to notice, react to, and effect repairs to each damaged component.

In Figure 11, it can be clearly seen that the number of field workers well exceeds the number of NOC workers required, as it does in reality. The number of workers dispatched to perform repairs reaches equilibrium quickly, with gradual movement from initial zero, rather than sudden and drastic movement of the workforce.

Absenteeism Modeling Approach

Two general structures for modeling absenteeism are included in the model. These structures move workers into and out of a stock of unavailable workers. Flow rates between the stock and the pool of available workers can be stipulated, or can be derived through goal-seeking on an exogenous fraction of affected workers. Although, in reality, workers can leave work at any time, in the model, workers become absent only from the pool of off-shift workers and return only to that pool.

The variable “epi switch” controls whether or not the model generates its own worker loss numbers or if the numbers are externally provided.

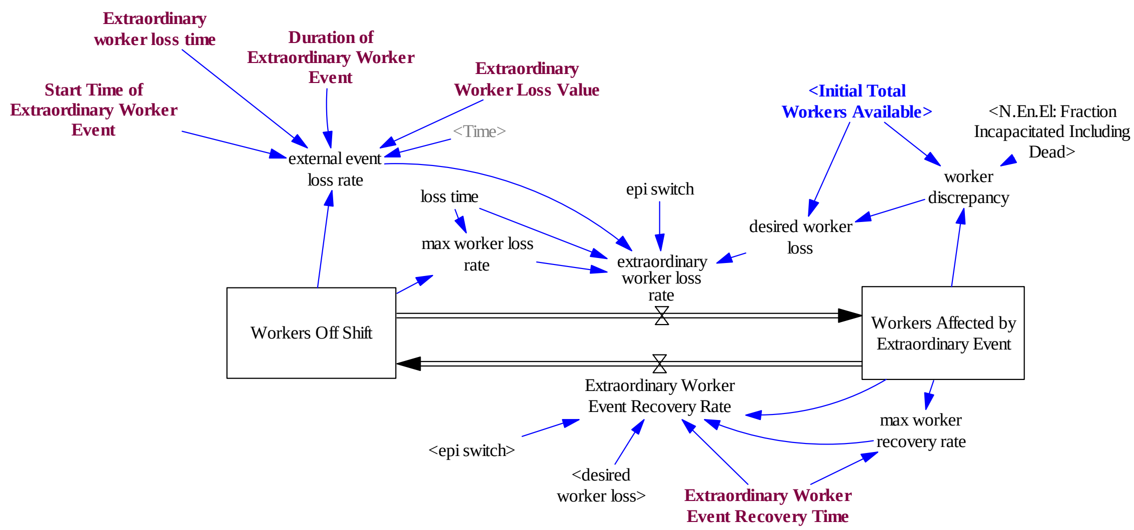

Figure 12: Overview of Worker Illness Models

The model shown in Figure 13 represents the effect of absenteeism on the work force when the variable “epi switch” is set to “0” (off). The variables governing the duration and rate of absenteeism are described in Table 2. In this configuration the model does not account for deaths due to illness, and all workers are at some point returned to the workforce.

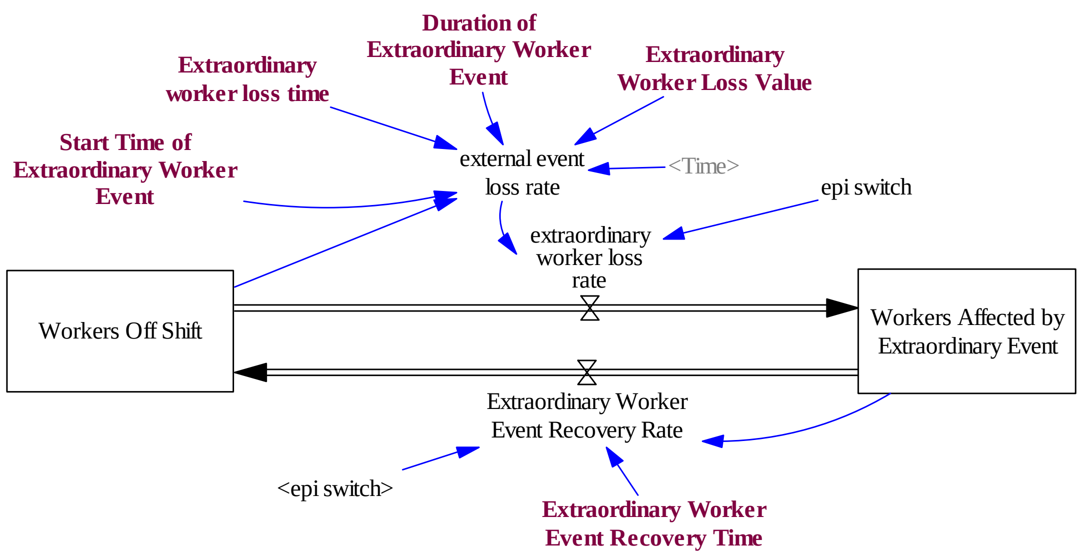

Figure 13: Internally Driven Absenteeism Model

PARAMETER	DESCRIPTION
Start Time of Extraordinary Worker Event	The two parameters control the timeframe of the absenteeism: at what hour of simulation it starts and how long it lasts.
Duration of Extraordinary Worker Event	
Extraordinary Worker Loss Time	The two parameters work in conjunction to drive the severity of the absenteeism: the workers become absent at the following rate: *Loss Value/Loss Time*
Extraordinary Worker Loss Value	
Extraordinary Worker Recovery Time	How many days before an absent worker can return to work

Table 2: Absenteeism Model Parameters

The model variables shown in Figure 14 allow for absenteeism to be input into the model from external sources when the variable “epi switch” is set to “1” (on). This allows a set number of workers to be removed from the active workforce at any given time. It can account for deaths due to illness by never returning workers to the work force (e.g. by having ten workers in the “workers affected by extraordinary event” stock at all time steps,).

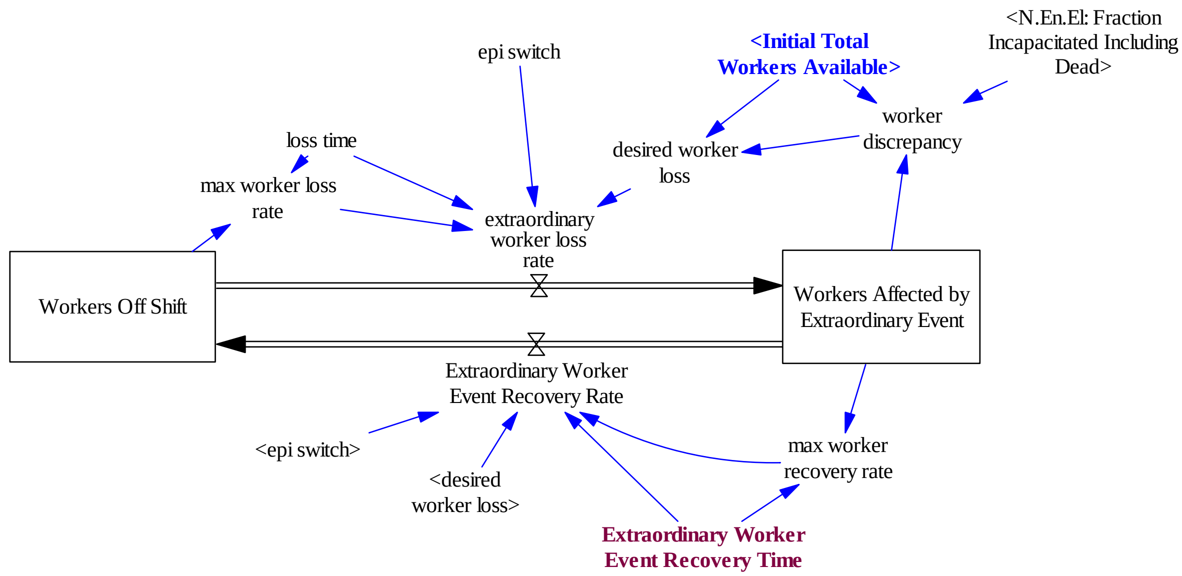

Figure 14: Externally Driven Absenteeism Model

Absenteeism Simulation Results

This section documents the baseline absence scenarios results. We find that even a small perturbation results in noticeable impact on overall levels of network failure. The two scenarios run were:

- Scenario A: Using internally generated absenteeism - 15% off shift-worker loss rate per day, with a 30 day absence period starting at day 10.
- Scenario B: Using externally generated goal-seeking absenteeism.

Scenario A

The worker loss rate in Scenario A is shown in Figure 15. This worker absenteeism removes 15% of off-shift workers per day for 30 days, placing them in the "workers affected by extraordinary event" pool for starting at day 30. Workers have on average a 10 day time period to be returned to the workforce. The resulting worker dispatch profile is shown in Figure 16.

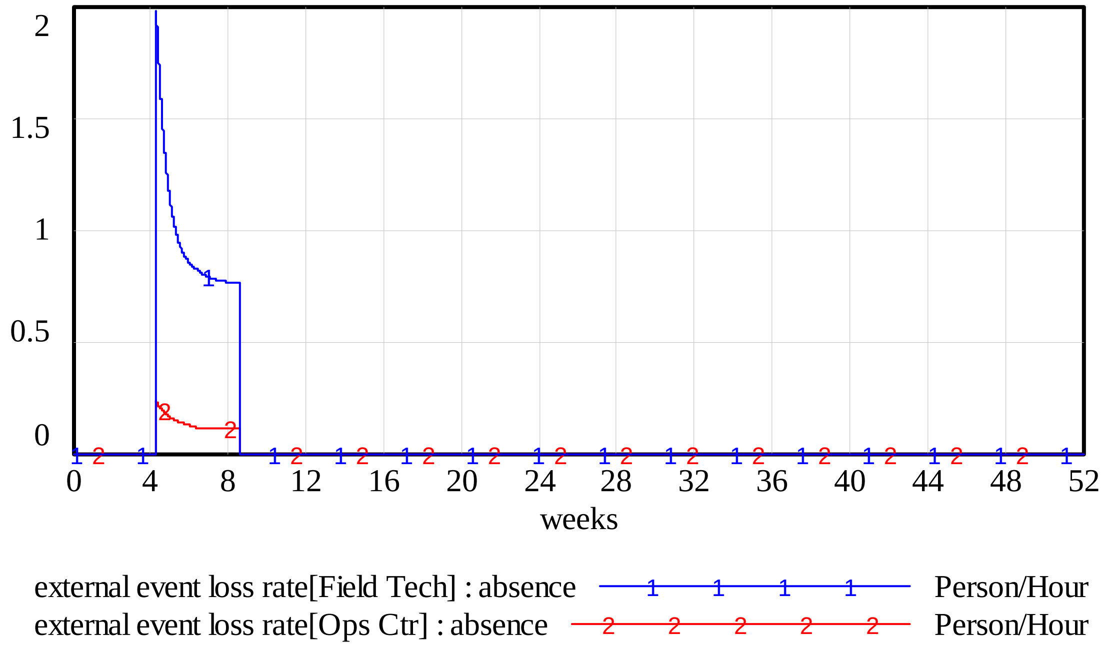

Figure 15: Absenteeism By Worker Type

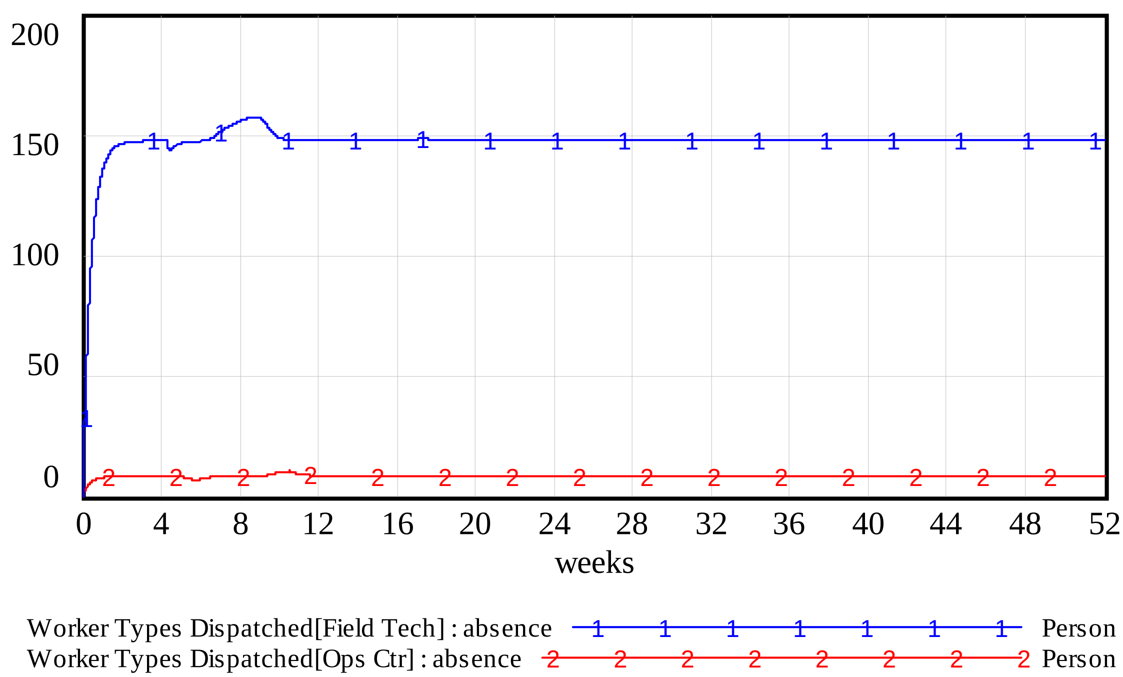

Figure 16: Dispatched Workers

As shown in Figure 17, as workers are removed from the pool of available workforce, the percent of the network that is damaged goes up. There is a slight increase to the percent network damaged after all of the workers are returned to the available pool due to an increase in perceived damage. When there are fewer workers in the workforce, damage is not noticed as rapidly. Once workers are returned to the workforce, unreported damage is again noticed at a normal rate and the amount of perceived damage and thus the percent of the network that is damaged goes up.

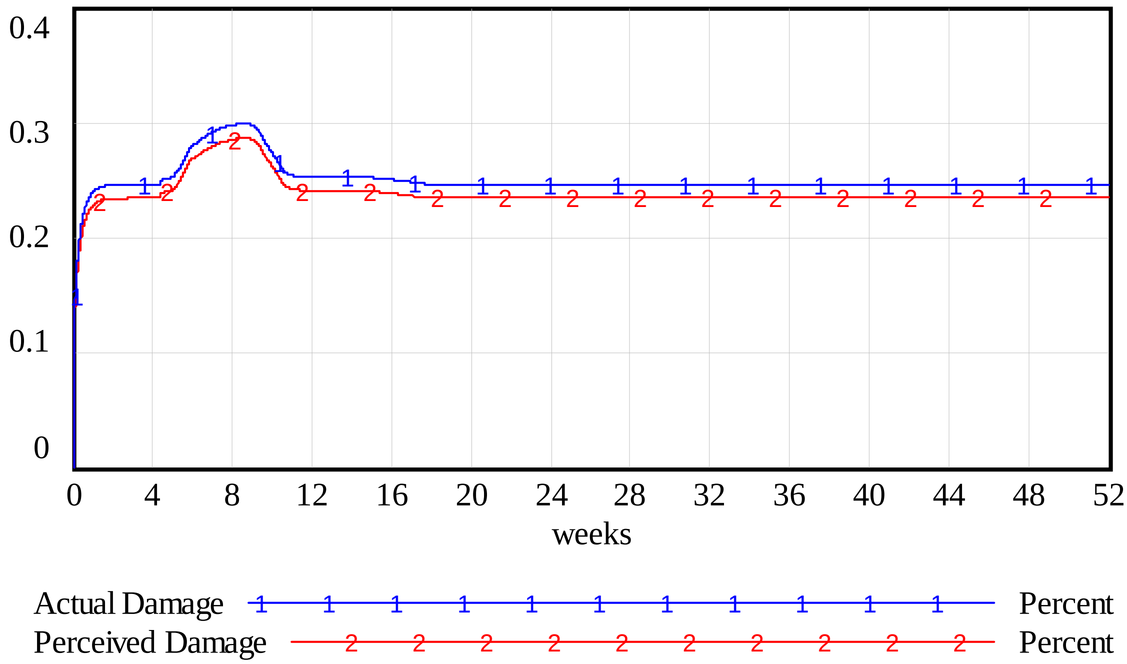

Figure 17: Perceived and Actual Damage

Figure 18 shows the necessary overtime for the workforce during the absenteeism period. When workers are initially removed from the workforce, the overtime required from the remaining workers goes up. It continues to rise as more workers are removed from the pool of available workers. Once workers are returned to the pool, necessary overtime begins to fall as the backlog of damage is worked off. Some amount of the backlog remains for the duration of the run due to repairs having a repair time tolerance, such that they don't need to be completed immediately.

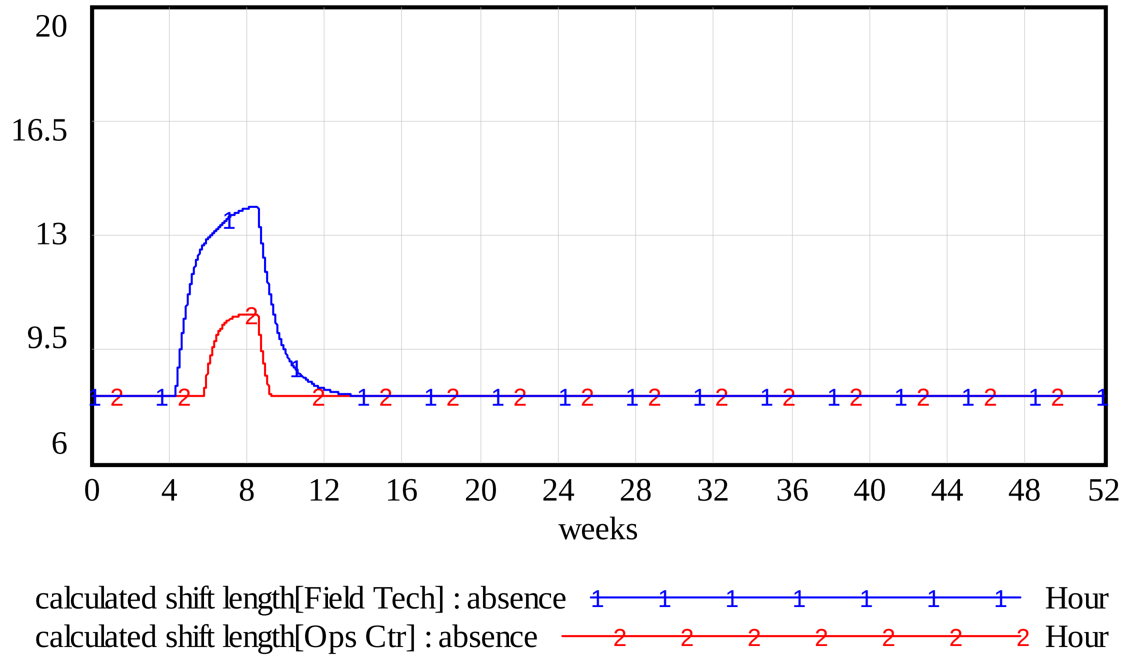

Figure 18: Necessary Shift Length by Worker Type

Scenario B

The worker loss rate in Scenario B is shown in Figure 19. The model goal seeks the number of workers absent to match the desired fraction of total workers absent resulting in the number of absent workers shown in the figure. The fraction of workers absent never returns to zero, simulating workers who never return to the workforce. The associated model results are shown in Figure 20.

Similar to scenario A, the percent network damaged peaks once the workforce has returned and the backlog of unnoticed damage is perceived. As shown in Figure 21, once the absenteeism begins, workers are required to work overtime. The model does not allow the workforce to go beyond a double shift, thus the necessary shift length never goes beyond sixteen hours. Overtime continues through the end of the model run as the backlog of damage slowly gets resolved.

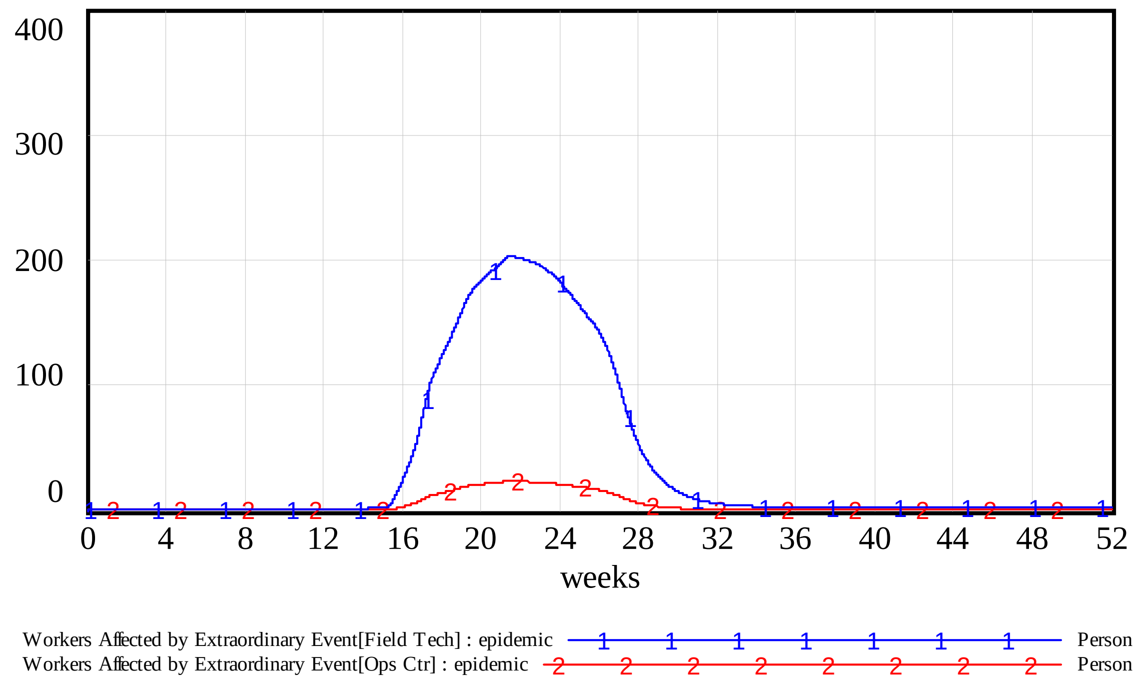

Figure 19: Absenteeism Scenario B

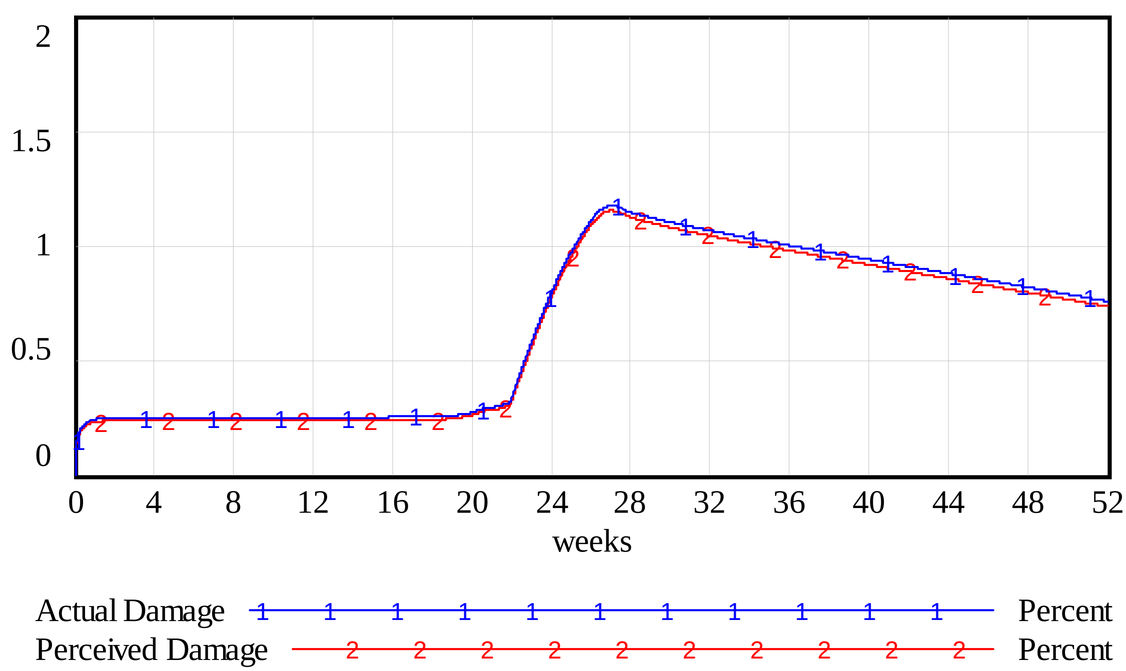

Figure 20: Absenteeism Scenario B results – Actual and Perceived Damage

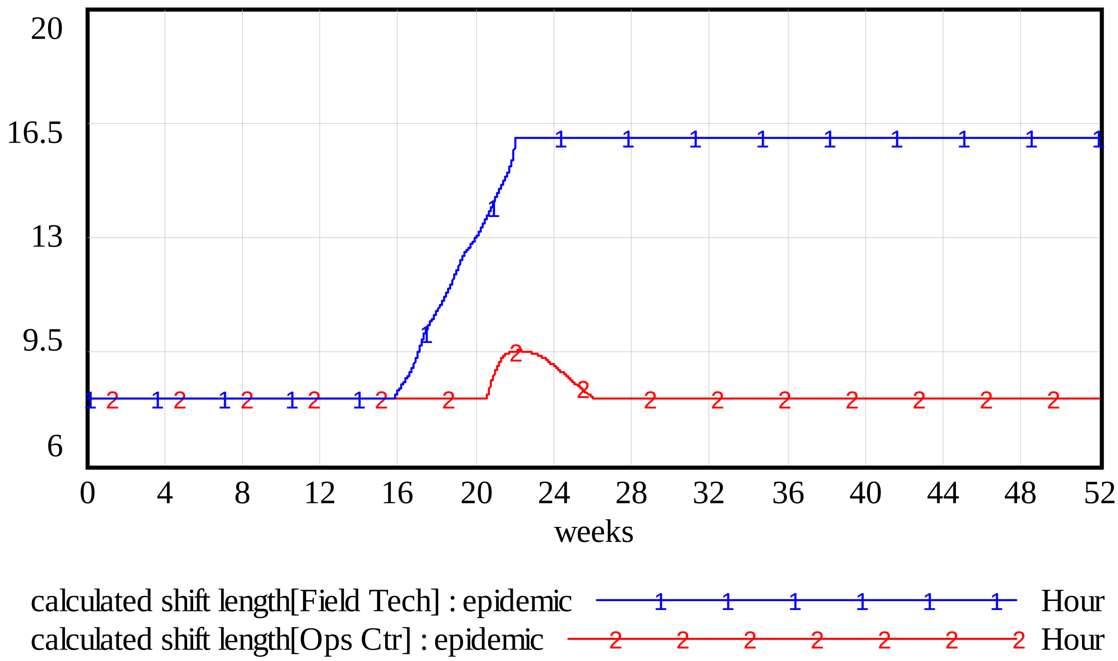

Figure 21: Necessary Shift Length by Worker Type

Conclusions and Future Work

The process of modeling telephony repair activities as well as analysis of the results shown here and additional scenarios not presented here, lead to several conclusions:

- The complexities of telephony repair are quite amenable to the modeling process. Simulation of historical events with known parameters showed close correlation between the resulting model repair times and the historically recorded repair times. The same is true for worker effort and staffing level values.
- The complexities of telephony repair are well modeled by SD flows, because although staff, inventory and failures occur in discrete units, the repair activities occur continuously over time and in fact, can be considered at any time to be percentage complete. The need to allocate staff to repair events lead quite cleanly to an interpretation of repair events as a stock of repairs-needed, calculated in terms of time, i.e. "staff-hours".
- In terms of insight derived from the results of running the model, it has been shown that repair and recovery functions in the voice telecommunications infrastructure are resilient to the loss of a significant portion of the work force. However, certain activities, such as sharing of warehouse inventory across regions and emergency staff augmentation during a crisis, is an increasingly effective response as the severity of the emergency increases.

Future work for the model includes testing of workforce management strategies in a crisis situation such as long-term worker loss or large scale outages.

References

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

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