LA-UR-08-07497

*Approved for public release;
distribution is unlimited.*

Title: Large-Scale Telephone Network Simulation: Discrete Event vs. Steady State

Author(s): Russell Bent, Stephan Eidenbenz, Sunil Thulasidasan

Intended for: 12th Communications and Networking Simulation Symposium (CNS 09)

Large-Scale Telephone Network Simulation: Discrete Event vs. Steady State

Russell Bent, Stephan Eidenbenz, Sunil Thulasidasan
Decision Applications Division and Computer, Computational, and Statistics Division
Los Alamos National Laboratory
{rbent, eidenben, sunil}@lanl.gov

Keywords: PSTN, discrete event simulation, steady-state simulation

Abstract

Motivated by the need of the emergency response planning community to obtain timely and accurate simulation-based predictions of the follow-up effects of telecommunications infrastructure outages, we study the trade-offs, advantages, and disadvantages of a high-fidelity time-dependent Public Switched Telephone Network (PSTN) simulator, called MIITS-P and its steady-state challenger, called IEISS. While IEISS computes results for individual simulation runs in a few seconds, MIITS-P call-by-call level simulation takes up to two magnitudes longer to complete runs, resulting in response times of tens of minutes. We study the quality of the results produced by both simulators for single-contingency cases, where we remove a wirecenter node from the network for an entire day and see what effect its removal has on the ability of the network to route its call traffic. Against our intuition that is based on a few theoretical examples suggesting that IEISS may not be able to even get base line traffic correct, we find that IEISS actually predicts relative rankings of wirecenter nodes extremely well on practical instances. This positive result holds under a specific transformation of the time-dependent simulation model into a steady-state equivalent; we look at a variety of such transformations and identify the best one. In an attempt to explain the good performance of IEISS, we take a closer graph-theoretic look at our network model, which reveals that most highly ranked wirecenters are connected in such a way that their removal cuts all possible paths between a large set of source and destination wirecenters. These cuts make time-dependencies and load variations over the course of a day insignificant and thus IEISS's steady state approach sufficient.

1. INTRODUCTION

The telecommunications sector is a key sector among the 17 critical infrastructure sectors defined by the US Department of Homeland Security[5], [6]. While communication networks are essential to the proper functioning of the national economy and government operations, they also satisfy social and emotional needs of individuals. The infrastructure of the Public Switched Telephone Network (PSTN) continues to be the backbone of the US communication fabric with its underlying technology now often in use for the Internet, other data networks, and the cellular phone network. The communication sector is interdependent with other sectors as it relies, for instance, on electric power (or the timely refueling of backup generators in case of emergencies), and in turn Supervisory Control And Data Acquisition (SCADA) systems for most infrastructure sectors rely on functioning communication networks; more complex interactions among sectors happen through demand shocks, where emergency events, such as an impending hurricane landfall or an unexpected bridge collapse, reliably lead to demand shocks or overloads in the telephone networks (particularly in the relatively little over-provisioned cellular networks).

The capability of predicting the effects of a telecommunications outage typically caused by an external event (such as a hurricane) is a crucial tool in contingency planning and real-time emergency mitigation, which is done at municipal, state and federal levels of government as part of an analysis routine, where such tools are used in decision support. The measures of interest in this context are the following: number of telecommunication assets that will be outaged, number of people without phone service, number of calls that can no longer be placed, and the economic impact of these call losses. The two metrics in the analysis cycle are timeliness and accuracy of the simulation prediction. With respect to response times, an analyst ideally "clicks a button" on a software tool that delivers the prediction instantly (e.g, up to 30 seconds); in some scenarios, response times up to one hour are permissible. Accuracy is obviously important as decisions based on the results of such tools have wide-ranging impact, for example in deciding a restoration order or pre-deployment of mobile replacement units in the case of cellular base stations. In this paper, we study the trade-offs between the two metrics of accuracy and timeliness on the concrete examples of a detailed CPU-intense call-by-call level, time-dependent PSTN simulator MIITS-P and its CPU-friendly, flow-based, steady-state challenger IEISS. We start by giving a few more details about the telephone network and how we model it in Section 2.

MIITS-P (Multi-Scale Integrated Information and Telecommunications System - Phone Sector) is a full-fledged PSTN simulator whose main simulator module takes as input a set of calls to be made over a 24 hour period as well as a network topology with link capacities and routes the calls over the network following an abstracted version of the SONET protocol stack. IEISS (Interdependent Energy Infrastructure Simulation System) is a steady-state simulator that takes as input a link-capacitated network topology and pair-wise flows between network nodes that have no notion of time. Both tools have been developed at Los Alamos National Laboratory. We present MIITS-P in more detail in Section 3 and IEISS in Section 4.

Squeezing the inherently time-dependent call sessions into a steady-state equivalent presents a modeling and logical challenge that can be solved by selective sampling and capacity modification approaches or by introducing a coarse time-stepped version of executing different flows at a small number of distinct points in time. Intuitively, it seems that aggregating calls over a longer period of time results in misleading flow values as a typical call duration is on the order of a few minutes and network loads shift geographically over the course of a day from residential to business areas and back. We present two small example topologies and call sets that would cause any steady-state approach to produce drastically wrong results (Section 5).

For the purpose of comparison, we choose as a reference scenario the single-contingency case for a 24 hour period and as our measure we look at the number of lost calls. A single-contingency case is a scenario, where one network node gets removed from the network. Such a removal will always result in the loss of calls that originate or terminate at an end device that has its last-mile connection to the removed node, it may also result in additional call losses for calls that were originally routed through the removed node and cannot find capacity on replacement paths. Somewhat surprisingly, we find through a large set of comparison runs that IEISS matches the MIITS-P results extremely well. A closer examination of the graph-theoretic structure of our networks explains this welcome phenomenon: it turns out that, at least for highly ranked wirecenters, their removal usually results in cutting the remaining graph into disconnected components that simply disallow calls to be routed between many pairs of nodes at any time of day. Thus, the time-varying and shifting call loads do not matter much. As further evidence for the correctness of this finding, we show that in the few cases, where calls get dropped due to lack of available capacities, MIITS-P's results are quite different from IEISS as MIITS-P can use the residual (reduced but still existing) capacity on paths to reroute. We present these results in Section 6.

Related work. The architecture and modeling paradigms of MIITS-P as well as aggregate results on national-level asset ranking have been described in [9]. A detailed introduction to IEISS can be found in [1]. Scaling trade-offs for PSTN simulations have been studied by O'Reilly et al. [8] in the context of comparing discrete-event simulation with a systems dynamics approach; their main finding is that the coarse systems dynamics approach matches well with aggregate results from the more detailed simulator. Unlike the system dynamic approach, IEISS is an asset-level based modeling approach. Thus, our finding is that IEISS and MIITS-P often agree on the aggregate call losses on an individual node level, whereas [8] focuses on the system-wide agreement. [3] studies scaling tradeoffs in the TCP networks by comparing discrete-event simulations with flow based simulations. Their flow based approach is competitive with discrete-event simulations. Similar to a systems dynamics model the flow based approach provides aggregate level information on call losses between nodes in the network. It is important to note that there is a large body of work addressing discrete-event simulations for a wide variety of communications applications. A small sample of this work includes [7], [2], [10]

2. MODELING THE NATIONAL TELEPHONE NETWORK (PSTN)

In the United States, the telephone network is divided into geographical zones called Local Access Transport Areas (LATAs). Calls within a LATA are generally treated as local calls and are within the service region of a local phone company (often referred to as Regional Bell Operating Companies, such as BellSouth, Pacific Bell etc.), while inter-LATA calls are assumed to be long distance calls, where the service is provided by long-distance carriers. A LATA is further subdivided into calling areas called wire-center regions. Each such region is typically serviced by a single wire-center facility (commonly referred to as Exchanges or Central Offices). End-subscribers (landline) located in a wire-center region generally connect to the same exchange. Even subscribers belonging to different providers are usually serviced from the same exchange, since more than one provider can be present at a particular exchange. Often, smaller carriers, i.e those distinct from the regional Bell companies [1], lease physical circuits from the larger operators—thus, physically calls are almost always routed through the principal wire center facility. Calls originating and terminating within a LATA, but in different wire-centers are usually routed through a specialized switch known as a local tandem. Long distance calls, (inter-LATA calls) are routed through an access tandem—these are tandem switches that are connected to various long distance provider networks; the circuit switching at this point depends on the subscriber's choice of long distance providers.

Thus, the telephone network, at the LATA level, is more or less hierarchical, with the end-subscriber devices, wire centers, local tandems and access tandems representing increasing levels within the hierarchy.

The elements of a cellular network are the base-stations (cell towers), and Mobile Switching Centers (MSC). An end-

[1]These companies are referred to as CLECS—Competitive Local Exchange Carriers.

device, i.e a cell phone is associated with a single cell tower; a cell tower often provides service to hundreds of end-devices, and is in turn connected to an MSC. An MSC can be associated with a number of base stations, and is also connected to a tandem in the local exchange (wire center). This represents the point of interface between the cellular and landline networks. If the end-devices are both mobile, then calls within a LATA are usually routed entirely through the tower and MSC network. If one of the devices is landline, then the calls are routed through one or more local tandems. Long distance calls originating from a cell phone are either routed through the long distance network of the cellular operator or handed off to access tandems like wireline calls.

Thus, each LATA has several central office buildings that house at least one switch. The switches may be either an end-office switch, a tandem switch, an access tandem or an MSC, i.e different kinds of network switches are usually co-located in the same building.

3. THE DISCRETE-EVENT INCUMBENT: MIITS-P

MIITS-P stands for Multi-Scale Integrated Information and Telecommunications System - Phone sector. MIITS-P models and simulates the wireline and cellular telephone networks of the entire continental United States on a call-by-call resolution. It is the one of the sectors of the MIITS-P suite, other sectors include the Internet and mesh networks. MIITS-P relies on scalable distributed event-driven simulation and on an end-to-end socio-technical modeling paradigm. See [11] for an introduction to the software engineering principles used in MIITS-P. We give a brief overview of MIITS-P, for more details we refer to earlier work [9].

MIITS-P's conceptual building blocks are the network generation module, the session generation module, and the actual simulator module. The *network generation module* models each of the roughly 31,000 network switches in the roughly 25,000 PSTN wire centers that are spread across the continental United States. Each switch is assigned to a wire center location that can house multiple switches. We also keep track of the functionality that the switch offers and we model the connections between switches and wire centers. The network generation module relies on industry data and engineering principles. The *session generation module* generates individual calls between end-user telephones aggregated up to wire center switches. It relies on US census data [4] and surveys of end-user behavior. Each call is assigned to a source and a destination wire center switch based on the attractiveness (according to the day- and night population of that switch for work-based or home-based calls). A wire center only gets assigned calls that originate in its service area polygon. Call generation is based on a number of parameters that effect the calling behavior such as time of day and the activity of the caller (work or home). Using these parameters and other social heuristics we could model different scenarios — even those that deviate significantly from the norm such as calling patterns during a hurricane or a terrorist attack. For our base case, we generate on the order of one billion calls. The *simulator module* routes the calls over the network imitating real-life routing algorithms as closely as possible through demand-based over provisioning and load-balancing. The simulator module is deployed on a distributed cluster and runs as a discrete event-driven simulation system, which enables us to scale to a national level within reasonable simulation time.

MIITS-P has been used in a variety of PSTN studies for emergency scenarios, in particular for modeling effects of hurricanes on the PSTN availability. The main MIITS-P result (see [9] for details) is a national scale ranking of each wire center that is based on the number of call-seconds routed over each wire-center over the course of a 24 hour work day.

4. THE STEADY STATE CHALLENGER: IEISS

Though MIITS-P provides a high fidelity representation of telecommunication systems, it suffers from significant computational requirements when performing simulations on a national scale. In emergency situations, such as hurricanes, wildfires, etc., when quick results are needed, the computational requirements are often too heavy to provide results before events render them meaningless. Such situations demonstrate that there exists a need to approximate the results of discrete event simulations such as MIITS-P and understand the quality loss when utilizing the approximation both in worst case and in practice.

We considered a number of different approximation techniques, (some are discussed in Section 6), however a steady-state based approach yielded the best results. Our steady-state simulation tool was developed and tested within LANL's interdependency simulator, IEISS (Interdependent Energy Infrastructure Simulation System) [1]. Intuitively, our steady-state model of telecommunications is equivalent to simulating a single time unit in a discrete event model. Indeed, the IEISS telecommunications models are derived directly from the MIITS-P simulator modules. IEISS's *session generation module* uses the *session generation module* of MIITS-P to generate a set of time indexed sessions. IEISS then selects a sample of those sessions without time to use as its set of sessions. Later in Section 6 we will compare some techniques for selecting the sample. Clearly, using sessions from a single time unit and simulating that time unit directly in MIITS-P model is likely to yield poor results as a single time unit's set of sessions represents a very small fraction of all the traffic routed by the model. Thus, the challenge is how to select a sufficient numbers of sessions that characterizes the traffic in the MIITS-P model. The *network generation module* is exactly the same between MIITS-P and IEISS for constructing the network topology. The main difference in the two modules is determining capacities on the links in the network as IEISS must simultaneously route sessions from disparate time periods that are never routed together in the MIITS-P. IEISS's capacities are chosen by first calculating the largest number of sessions that co-exist in the discrete event model (the peak load on the network) and second calculating the number of sessions generated by IEISS's *session generation module*. The ratio between these two calculations is multiplied by the capacities in the MIITS-P network to determine the capacities of the IEISS network. The *simulator module* represents the largest departure between the MIITS-P and IEISS approaches. IEISS simulates telecommunications by iteratively (at random) choosing unrouted sessions until every session is either routed or dropped. For each session, IEISS chooses a route between the session's source and destination using the routing preference algorithms developed for MIITS-P. The most preferred route with available capacity is chosen to handle the session. If all possible routes contain a link that has reached saturation, then the session is dropped and considered lost.

5. THE CHALLENGE

The quality of the IEISS approximation when compared to MIITS-P is dependent on the metric used for comparison. In this paper, we use metrics that calculate the relative importance of wire centers in telecommunications networks based on the amount of traffic that is lost if the wire center is out-of-service (often referred to as contingency-case analysis). This metric was chosen due to the types of analysis the authors are engaged in for telecommunications (i.e. [9]). We illustrate the challenge of going from an event-driven simulation with an explicit notion of time to a steady-state model with two small examples that point out the basic difficulty in making time disappear for this metric and illustrates possible worst cases for our approximations.

Consider the small network example and the associated small session set of Figure 1. In a base case scenario, MIITS-P will route the ten 11 am calls from vertex W to destination vertex D by routing 5 calls each through vertices a and b and then taking one of the two possible paths to D. Similarly, it will obviously route the ten 7 pm calls from vertex H to D over one of the two possible paths without violating any capacity constraints. In a contingency case scenario, where vertex a is removed from the network, MIITS-P would only be able to route five of the 10 calls originating from vertex W, resulting in 5 non-routed or lost calls. In a contingency case, where vertex c is removed, MIITS-P could still route all calls as the ten calls from W and H happen at different times of the day leaving enough capacity for all traffic to be routed

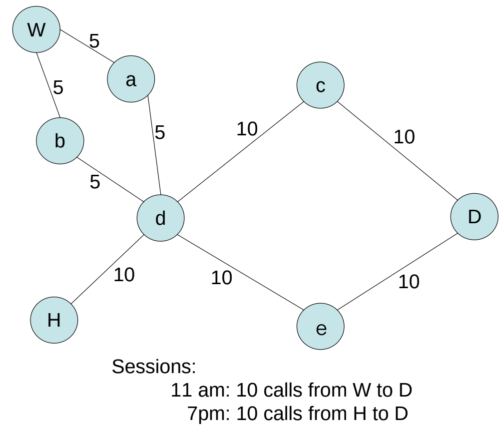

Figure 1. Pathological instance leading to IEISS mis-estimation of call losses. Numbers represent link capacities

through vertex e.

IEISS, on the other hand, samples the set of sessions and sets capacities as described in Section 4. Assume first that IEISS samples randomly from all 20 sessions and for the sake of argument, let us assume without loss of generality that it samples all 20 sessions. The capacity re-normalization step outlined in Section 4 results in the exact same capacities as in MIITS-P. In a base case run, IEISS would thus correctly predict that no calls will be lost. In the contingency case with node a removed, IEISS also correctly predicts that 5 calls will be lost. However, in the contingency case with node c removed, IEISS wrongly predicts that 10 calls will be lost as it tries to route 20 calls in one synthetic point in time through the only path through vertex e. It is easy to see that different capacity re-normalization procedures (as long as they multiply the original capacity with a constant) will not solve this issue but merely switch the wrong predictions between the two contingency cases. To be more precise: if IEISS re-normalizes capacities by a factor 2, the contingency case with vertex a removed would result in a wrong prediction of zero call losses, while the contingency case with vertex c removed results in a correct prediction of zero call losses. Similar reasoning shows that IEISS can be forced to make arbitrarily wrong predictions even if it samples in a biased manner.

The network in Figure 2 provides an example that results in call losses in the base case. In the base case with a complete sample set, IEISS and MIITS-P are able to achieve the same result using the call loss metric described earlier (i.e. 10 calls lost). However, this example is demonstrative of how the challenge of approximating the session-loss metric is not specific to our metric choice. For a metric based on calculating call origination loss for each vertex in the network, this network is difficult for IEISS to accurately approximate. In

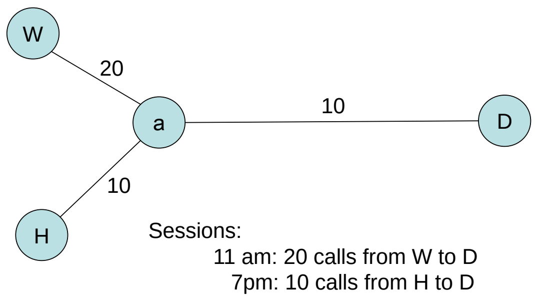

Figure 2. Pathological instance leading to IEISS errors in base case. Numbers are link capacities

the MIITS-P simulation, 50% of the W-originated calls are lost and 0% of the H-originated calls are lost. However, in the IEISS simulation, due to the random simulteanous session routing algorithm, we would expect 33% of the W-originated calls to be lost and 33% of the H-originated calls to be lost, a result that is a considerable distance from the MIITS-P result. Indeed, there is no way to scale the capacities of this network (using a uniform constant scalar) and achieve the MIITS-P result. Clearly, a more sophisticated (non-constant) capacity callibration technique is needed to address the difficulties of this example.

These two examples suggest that IEISS could face an uphill battle in correctly approximately call losses. However, it turns out that our small examples apparently do not occur very often in reality. We next study IEISS performance on real LATA networks and realistic session sets in the next section.

6. RESULTS

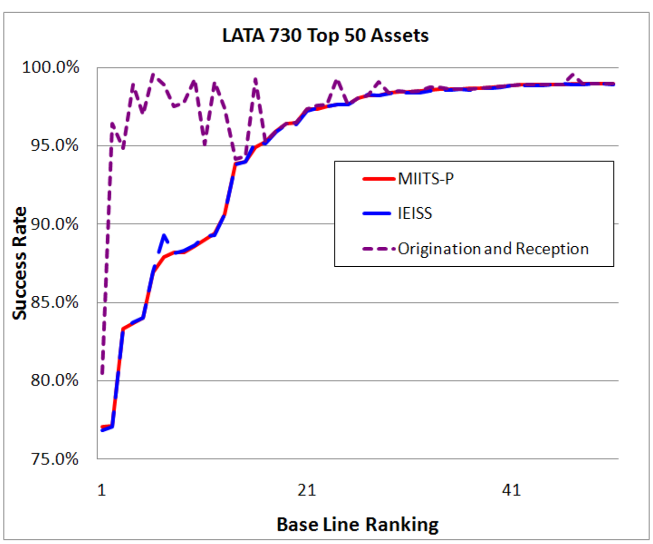

Figure 3. Success Rate Comparison

In order to compare the performance of MIITS-P and IEISS, we considered a number of different models based on publically available telecommunications data. Ultimately, models based on LATA 730 (Southern California) and LATA 236 (Washington, DC) yielded model characteristics that produced the most interesting results in comparing MIITS-P and IEISS. LATA 730 was interesting due to the very high call volumes in Southern California and LATA 236 was interesting due to the existence of many alternate paths between wire centers that made IEISS's approximation task more difficult.

Figures 3 and 4 compare IEISS with MIITS-P based on telecommunications data for LATA 730 and LATA 236. For each wire center in these models, a call completion success rate is calculated by simulating the model with the center out-of-service using MIITS-P (i.e. a contingency analysis). The wire centers are plotted from lowest success rate to highest success rate in a solid red line in Figures 3 and 4. The blue dashed line plots the same wire centers according to the success rate calculated using IEISS. IEISS's model contains 1% of the sessions in the MIITS-P model, drawn uniformly at random. A good result is a plot of IEISS that is monotonically increasing, indicating that the steady state simulator is a good approximation of the relative rank calculated by MIITS. Interestingly, from these results not only is IEISS able to achieve a monotonically increasing approximation of MIITS-P (relative success rate), it also does a remarkable job of approximating the actual success rate except in a handful of instances. Also included in Figures 3 and 4 is a success rate calculation for each wire center based on the number of sessions originating or destined for the wire center (plotted in a purple dotted line). The jaggedness of this approach's plot indicate that such a technique is not as effective at approximating the relative importance of a wire center as it does not account for traffic that passes through the wire center.

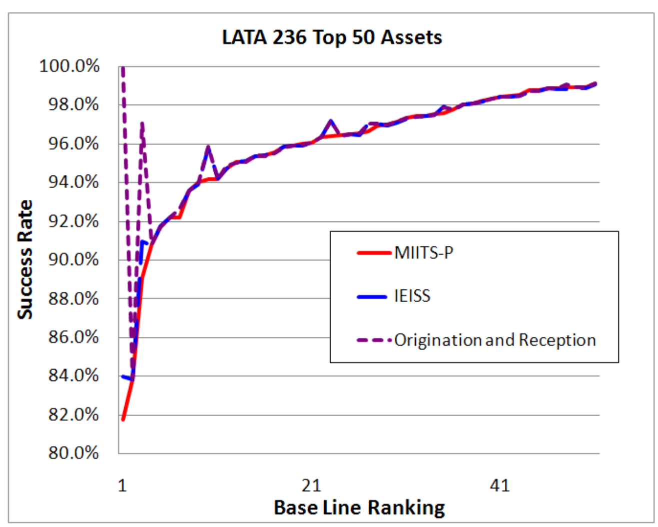

Figure 4. Success Rate Comparison

Often, the authors are required to rank wire centers across multiple disparate telecommunications models. Thus it becomes important to determine if minor errors in individual models are magnified when results are combined. Figure 5 considers a case of combining results from models built from separate LATAs (120, 236, 490, and 730). The figure shows a log plot of wire centers by number of calls dropped when the wire center is out-of-service. The order of the wire centers is determined by the number of calls dropped in the MIITS-P simulation. Once again, the generally monotonically decreasing plot of the IEISS results indicates that it is a good approximation of MIITS-P's relative call drop rates. The poor shape of the plot approximating call drops based on origination and destination suggests that this approach's errors are greatly magnified when combing results from disparate models.

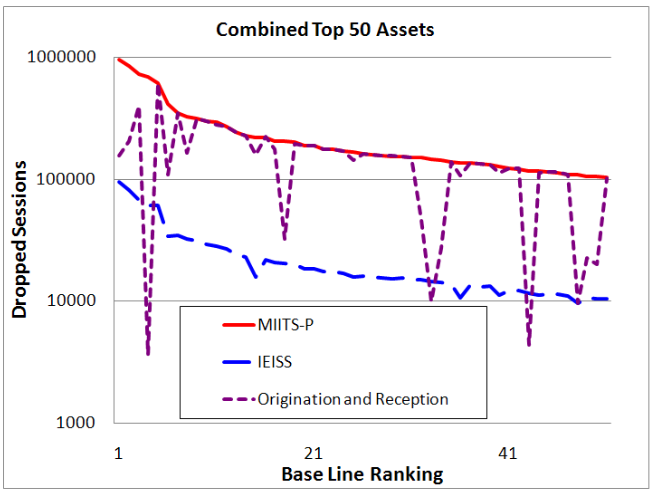

Figure 5. Success Rate Comparison

The key benefit of the steady state simulation technique arises from its computation time requirements. As seen in Table 1, the computational requirements are two orders of magnitude smaller than MIITS-P, while achieving results that are very comparable, as seen in the results discussed earlier.

CPU Seconds	Min	Max	Average
MIITS-P	218	1826	591
IEISS	5	17	9

Table 1. CPU Time comparison

6.1. Sample Sensitivity

It is natural to consider whether or not the sampling choice has an effect on the quality of the IEISS simulations. Figure 6 compares uniform sampling with a sampling approach based on choosing calls from load peaks in the model. The intuition behind such approach is that in many scenarios it is during the peaks that most negative effects from a wire center being out-of-service are observed. As is seen in the figure, the results remain reasonable; however sampling uniformly clearly provides better results overall.

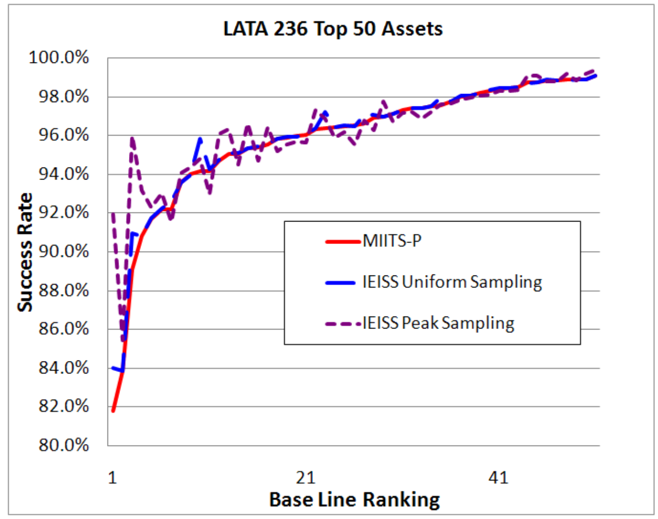

Figure 6. Sampling Comparison

6.2. Analysis

Finally, it is interesting to understand why the steady state simulation performs so well. On closer inspection of the results, in the majority of models, sessions are primarily dropped due to the lack of a path between the source and destination. Thus, as long as the session sample of IEISS is a good representation of the MIITS-P sessions, it is relatively easy to duplicate the results. This observation indicates that a simple path analysis that is computationally very efficient is sufficient to approximate the results of MIITS-P on real models. Figure 7 suggests otherwise as in a handful of wire center out-of-service scenarios, a path analysis is not as good as IEISS (though still reasonable). A closer inspection of those scenarios shows that they are cases where most calls are dropped due to congestion, suggesting that in models where calls are dropped due to congestion, a path analysis will perform poorly. From these results, it appears that the steady state simulation proposed here is able to handle congestion very well.

To further demonstrate that the steady simulation technique is adept at correctly handling congestion, we consider more closely a wire center out-of-service scenario where calls are dropped due to congestion. This is a scenario drawn from the LATA 236 model where more than half of the lost calls are due to congestion. In this scenario, MIITS-P predicts that 81.8% of the calls are able to be routed when the wire center is out-of-service and the steady state simulation indicates that 84.0% of the calls are able to be routed. These two results are

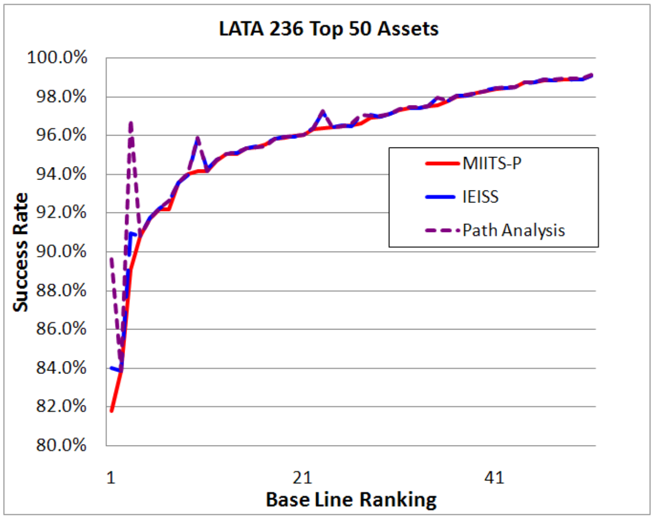

Figure 7. Path Analysis

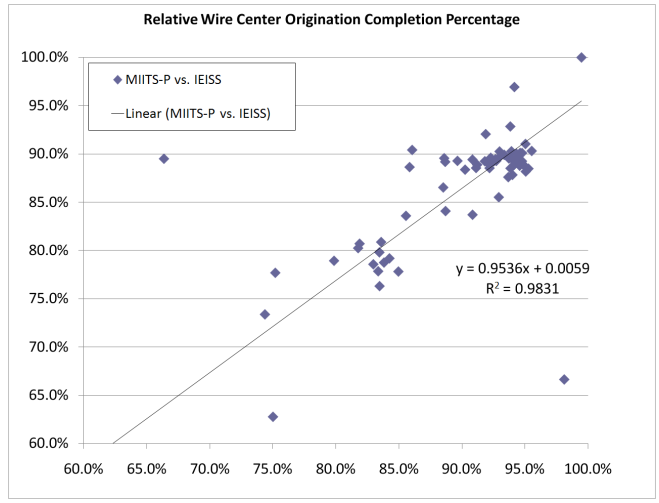

Figure 9. Wire Center Comparison

remarkably very close, indicating that IEISS is able to handle congestion correctly. Figure 8 considers another metric on this *hard* case which compares the load handled by each wire center in terms of a percentage of total possible load. The figure shows that the results from the two simulations are highly correlated with an R^2 value of .95 indicating that the routing profiles of the two simulation techniques are very similar. Furthermore, Figure 9 considers the completion percentage of calls originating at each wire center in this scenario for both simulation techniques and achieves an R^2 value of .98, indicating that local failure rates are also similar.

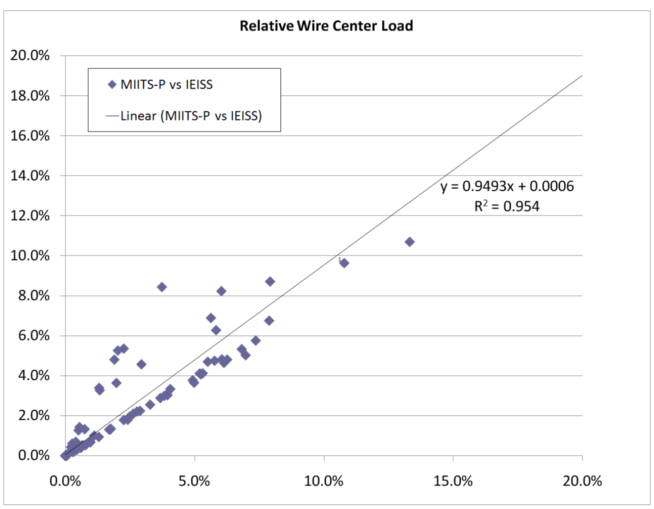

Figure 8. Wire Center Comparison

7. CONCLUSIONS

In recent years, considerable effort has been spent developing finely detailed discrete event simulation models for telecommunications. This work has advanced the ability of researchers to accurately model and simulate the behavior of real networks. The weakness of such approaches is in the required computational demands. In many applications such as the national ranking problem described in [9] the computational demands are not overly burdensome, however, in fast response situations such as during emergencies, it becomes necessary to consider more computationally efficient techniques. This paper explored a steady-state model for telecommunication that is two orders of magnitude faster than a discrete event model and whose results under metrics of interest compare very favorable to the discrete event model.

The results of this paper suggest a number of very interesting future research directions including the exploration of other comparison metrics to determine if the results on the metrics presented here hold under other conditions, determining techniques for inferring time-based information from steady-state results - such as the number of calls dropped between 6pm and 7pm -, and exploring ways of better approximating the pathological cases described in Section 5. One way to possibly address this last question is to attempt to statistically calibrate the capacity choices of the steady-state model using results from a handful of discrete event simulations or perhaps developing a hybrid discrete-event/steady state system. Finally, it will be interesting to consider how some of the steady-state approximations suggested here might be used in other physical networks such as the Internet, electric power, or natural gas.

Acknowledgments This work was supported in part by the U.S. Department of Homeland Security's NISAC (National Infrastructure Simulation and Analysis Center) program. We also thank the three anonymous referees for their helpful comments.

REFERENCES

[1] Brian Bush and et al. Interdependent energy infrastructure simulation system - ieiss technical reference manual. *LANL Technical Report LANL-D4-05-0027*, pages 1–254, 2005.

[2] R. Fujimoto, K. Perumalla, A. Park, H. Wu, M. H. Ammar, and G. F. Riley. Large-scale network simulation: How big? how fast? *Proceedings of the 11th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems. (MASCOTS 03)*, pages 116–123, 2003.

[3] Kayo Fujiwara and Henri Casanova. Speed and accuracy of network simulation in the simgrid framework. *Proceedings of the 2nd International Conference on Performance Evaluation Methodologies and Tools*, pages 1–10, 2007.

[4] T. McPherson, J. Rush, H. Khalsa, A. Ivey, and M. Brown. A day-night population exchange model for better exposure and consequence management assessments. *86th AMS (American Meteorological Society) Annual Meeting, Atlanta, GA.*, pages 1–6, 2006.

[5] U.S. Department of Homeland Security. Homeland security presidential directive 7 (hspd-7). 2003.

[6] U.S. Department of Homeland Security. National infrastructure protection plan. 2006.

[7] G. OReilly, A. Jrad, T. Brown, and S. Conrad. Critical infrastructure analysis of telecom for natural disasters. *Telecommunications Network Strategy and Planning Symposium*, pages 1–6, 2006.

[8] G.P O'Reilly, A. Jrad, A. Kelic, and R.R. LeClaire. Telecom critical infrastructure simulations: Discrete-event simulation vs. dynamic simulation how do they compare? *Proceedings of Global Telecommunications Conference (GLOBECOM)*, pages 2597 – 2601, 2007.

[9] V. Ramaswamy, S. Thulasidasan, P. Romero, S. Eidenbenz, and L. Cuellar. Simulating the national telephone network: A socio-technical approach to assessing infrastructure criticality. *Proceedings of 2007 Mililtary Communications Conference (MILCOM)*, pages 1–7, 2007.

[10] D. Rao and P. Wilsey. Simulation of ultralarge communication networks. *Proceedings of the 7th IEEE/ACM Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'99)*, pages 112–119, 1999.

[11] R. Waupotitsch, S. Eidenbenz, L. Kroc, and J.P. Smith. Multi-scale integrated information and telecommunications system (miits): First results from a large-scale end-to-end network simulator. *Proceedings of Winter Simulation Conference*, pages 2132–2139, 2006.