Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

DEVELOPMENT OF AN INTERNET BACKBONE TOPOLOGY FOR LARGE-SCALE NETWORK SIMULATIONS

Michael Liljenstam
Jason Liu
David M. Nicol

Institute for Security Technology Studies
Dartmouth College
45 Lyme Rd., Suite 300
Hanover, NH 03755, U.S.A.

ABSTRACT

A number of network simulators are now capable of simulating systems with millions of devices, at the IP packet level. With this ability comes a need for realistic network descriptions of commensurate size. This paper describes our effort to build a detailed model of the U.S. Internet backbone based on measurements taken from a variety of mapping sources and tools. We identify key attributes of a network design that are needed to use the model in a simulation, describe which components are available and which must be modeled, and discuss the pros and cons of this approach as compared to synthetic generation. As for attributes that we have to model, we also briefly discuss some measurement efforts that can potentially provide the missing pieces, and thus improve the fidelity of the model. Finally, we describe the resulting network model of the U.S. Internet backbone, which is being made publicly available.

1 INTRODUCTION

Simulation is a key technology for the investigation of communication networks. Not only is it useful for preliminary study of protocols and network applications, it can reveal unexpected system dynamics. The importance of the Internet makes it an attractive object of simulation-based study. Measurement data has revealed interesting behaviors, including the well known long range dependency aspect of traffic reports, topological characteristics (e.g., power-law connectivity distributions of Autonomous Systems, Faloutsos, Faloutsos, and Faloutsos 1999), and instability of convergence using the Border Gateway Protocol (BGP) (Labovitz et al. 2001). There is a strong suspicion that network behaviors are heavily influenced by network topology. Research projects for mapping the Internet have produced diverse tools and techniques for observing network topologies, while at the same time other projects look at ways of generating synthetic networks that have some of the characteristics observed in real networks. Nevertheless, "what is a representative network topology to use for simulations?" appears to be a recurring question.

In this paper we try to capitalize on recent developments in Internet mapping tools, such as *skitter* from CAIDA (McRobb and Klaffy 1998), *Mercator* by the SCAN project at ISI (Govindan and Tangmunarunkit 2000), and *RocketFuel* from Univ. of Washington, WA (Spring, Mahajan, and Wetherall 2002), to generate realistic network models from collected maps and some supplementary data sets.

1.1 Related Work

Several research efforts have tackled the problem of generating synthetic representative topologies, and some of them, like GT-ITM (Calvert, Doar, and Zegura 1997) and BRITE (Medina et al. 2001), also include functionality for exporting these topologies in formats that can be read by certain network simulators.

Less appears to have been done towards applying discovered topology data sets from Internet mapping projects directly for simulation. Perhaps, because at first blush it appears to be more a question of implementation than research. One notable exception to this seeming lack of attention is the Boston University Representative Internet Topology Generator (BRITE) (Medina et al. 2001). The current version of BRITE provides facilities for importing data from CAIDA's skitter (McRobb and Klaffy 1998) IP-level topologies and router-level maps from the SCAN project (Govindan and Tangmunarunkit 2000). It also provides facilities for exporting the topology to configuration file formats used by the `ns` (ns 2003) simulator, SSFNet (Cowie, Nicol, and Ogielski 1999), and JavaSim (DRCL JavaSim 2003). It has

also been stated that other import formats will be supported in the future. BRITE strives to become a unifying tool that can be used to incorporate data and algorithms from different sources and readily provide output for a variety of uses.

We find this effort laudible as a unified approach and a method for data interchange is very useful. Nevertheless, it is worth noting that whether or not contents of new import or export formats can be handled by this method inherently depends on how rich the internal representation in the tool is. That is, in order to make use of any possible data type in the input streams to generate output for a simulation tool, the internal format would have to be the union of all data types that can occur in the inputs. Hence, although it is not necessarily prohibitively difficult, having full interchangability of data from inputs to outputs (without loss of information) does not come for free as it requires an internal representation that encompasses all possible inputs and strategies for dealing with lack of data in specific input and output formats.

More importantly though, the focus in this tool is on the topological structure of the graph. While these features are important, many other features such as link bandwidths, various latencies, queue sizes and protocol parameters are also important and no tools currently exist that can generate representative graphs equipped with all these "labels". For instance, the simulation configuration files generated by BRITE for SSFNet have several limitations in terms of the generation of these "labels".

Our approach is different in that we exclusively use (parts of) the real Internet topology, rather than synthesized topologies, and we focus on how we can use a combination of different collected data sets about the Internet to add more "labels" to the graph and thus build a more complete simulation model. But we use only the input data sources we find necessary for our purposes and we limit the available output formats to two options: *i)* **DML** configuration files for the the SSFNet simulator. *ii)* A 'distilled' network topology in an **XML**-based format. It should be straight-forward to convert to other formats for other uses. Thus, our approach is not meant to be 'universal', but we would not exclude the possibility of adding the methods we describe into tools such as BRITE in the future.

1.2 Synthetic Topologies vs. Collected Maps

In the simulation literature, see for instance Law and Kelton (2000), Banks et al. (2000), there are well known trade-offs when choosing between using collected data sets to directly drive a simulation and using some theoretical distributions derived from the data as simulation input. Most of these are also relevant for network topologies. Generally speaking, directly using a collected data set is excellent in terms of providing a realistic instance of input (modulo the quality of the data set). On the other hand, it is limited since it provides only one input. Hence, using statistical models based on the data is preferable in many cases as it "exercises" the model in more ways than a single input set can provide. However, ultimately what is preferable depends on the purpose of the study and for synthesized models the identification of representative generalizations is an issue in itself.

Finding representative generalizations of network topologies is an active research area, and there appears to be an ongoing debate as to what the most relevant features to capture are (Tangmunarunkit et al. 2002). Moreover, the relative importance of features of the topology depends on the system being studied and the effects of topology on protocols is not well known in many cases.

When contemplating the choice of synthesizing topologies or directly using measurements, a useful analogy can be made to models of radio channels used in the design and deployment of radio systems (Ahlin and Zander 1997, page 116). One may classify channel models into two categories: *design models* and *verification models*. Design models are typically more abstract statistical generalizations of many situations used to subject designs to many different conditions. Thus, the goal at this stage is to find a design that will work well under all realistically varying conditions, but in so doing will tend to abstract away many details. Validation models, on the other hand, are used to provide a very detailed model, as close to reality as possible, to test a proposed design and thus it may in some cases even replace actual field tests. Because of the high level of detail, these models are frequently limited to a specific situation. For instance, in channel modeling it may involve a ray-tracing model of radio propagation based on the construction plans and topographical maps of a specific site where a system is to be deployed. Consequently, the goals differ for these categories of models, as well as the techniques used.

We believe synthetic network topologies offer suitable design model generalizations, while models based on collected network maps and measurements can serve as validation models. Consequently, we envision many cases where it is useful to be able to directly base a network simulation model on a collected network topology, and this provides the motivation for our focus here. Of particular interest to us, are studies of the impact of attacks or disaster events on the current Internet infrastructure. Clearly, this is one instance where a detailed model of the actual infrastructure is desirable.

1.3 Contributions

We focus on directly using available router level maps and other data sets to create realistic models of parts of the Internet. Since we are aware of no single data set that contains all the information we need, such as router adjacencies, geographical mappings, link bandwidths and

delays, and so forth, we combine multiple data sets. In this process we have to deal with many issues related to imperfections in the data sets, lack of data, and how to combine data sets into a coherent whole.

The main contributions of this paper lie in recounting our experiences in building a large-scale model from Internet topology maps and other data sets. We also identify areas where data is lacking, in the hope of spurring measurement efforts on in directions beneficial to this kind of model building. Finally, we describe our initial resulting model of the U.S. Internet backbone, and we are making this model available to the research community for experimentation.

The remainder of this paper is organized as follows: Section 2 describes the steps we go through to create a simulation model from collected network maps of the Internet backbone and some supplementary data. Section 3 describes the end result in quantitative and qualitative terms. As we went through this process we identified several missing pieces of information and in Section 4 we discuss efforts in network research that could potentially provide the missing pieces of data. Finally, Section 5 concludes.

2 TOWARDS A METHOD FOR GENERATING REALISTIC NETWORK MODELS

The first question to ask is what data is necessary to collect for the model. Any list of data types 'necessary' to build a simulation model will, by necessity, be biased depending on what the simulator in question models and can represent, and the purpose of the study. Our list is based on creating a model in the SSFNet simulator (Cowie, Nicol, and Ogielski 1999, SSFNet 2003). SSFNet is a packet level network simulator written in Java that includes standard TCP/IP protoocols (IP, ICMP, TCP, UPD, HTTP, ...) and detailed implementations of routing protocols such as BGP and OSPFv2. A distinguishing feature of SSFNet is that it was built "from the ground up" to support parallel and distributed execution of network models for large-scale simulation.

A non-exhaustive list of data types needed to build an SSFNet model follows:

Router connectivity Router adjacencies as typically given by router-level maps.

AS mapping Routers should be annotated with the Autonomous System (AS) they belong to. Thus, it can be inferred which part of a network forms an AS (or similarly belongs to a particular ISP).

Exchange points and inter-domain routing policies Network exchange points, i.e. where different ISPs (ASes) peer and exchange traffic need to be configured. If we have router-level maps and know the mapping to ASes the locations of peering points can be deduced from the router maps. However, the policies governing the exchange of traffic (customer-provider, peer-peer, and sibling-sibling relationships) have to be inferred from other sources or based on heuristics.

Intra-domain routing configurations Depending on the nature of the study, the intra-domain routing protocol used (such as OSPFv2) and its configuration (e.g. link weights).

Link Bandwidths Link bandwidths are needed for many studies.

Link Delays Link delays need to be set.

Router queues For studies related to things such as congestion and Quality of Service, the sizes and configuration of queues in the network is important.

2.1 A Starting Point

As our starting point we selected the RocketFuel data set (Spring, Mahajan, and Wetherall 2002). RocketFuel is a router-level mapping tool where one of the underlying ideas is to focus on one specific ISP network at the time and try to map it as completely as possible. This is in contrast to other projects that typically try to map as much as possible of the whole Internet with no particular focus.

In building our model we decided to focus on Internet backbone covering the U.S. and thus select the 6 ISPs in the RocketFuel data set that operate in the U.S. Table 1 lists these 6 ISPs. For comparison, the top ten ISPs in the year 2000, according to Haynal (2000), are listed in Table 2.

Table 1: The 6 ISPs Used from the RocketFuel Data Set

AS	ISP Name	Routers bb	total
1239	Sprintlink	547	8,355
2914	Verio	1,018	7,336
3356	Level3	624	3,446
3967	Exodus	338	900
6461	Abovenet	367	2,259
7018	AT&T	733	10,214

Table 2: Top 10 ISPs in 2000 According to "Russ Haynal's ISP page"

By Market Share		By Connectivity
Share	ISP Name	ISP Name
27.9%	UUNet/WorldCom	UUNet/WorldCom
10.0%	AT&T	Sprint
6.5%	Sprint	Cable & Wireless
6.3%	Genuity	Genuity
4.1%	PSINet	AboveNet
3.5%	Cable & Wireless	AT&T
2.8%	XO Communications	Qwest
2.6%	Verio	Verio
1.5%	Qwest	Global Crossing
1.3%	Global Crossing	TeleGlobe

Since the RocketFuel data lacks some of the largest tier-1 ISPs, and a "backbone topology" would not be very representative without the largest networks, we decided to attempt to also include map data from the SCAN project (Govindan and Tangmunarunkit 2000) collected using their tool Mercator. The 6 ISP maps from the RocketFuel data set contain a total of 32,510 routers, out of which 3,627 are believed to belong the ISPs' backbone networks. The basic data provided is a list of router adjacencies (the physical connectivity of the network) and a list of IP aliases for each router. In addition, there is some information attempting to classify the routers and provide physical locations. Heuristics are applied to the router mapping process to try to tag them by their distance from the ISPs backbone. For instance, routers tagged by distance zero ('r0') are believed to be backbone or gateway routers judging by their names. Routers at distance one (marked 'r1') connect to 'r0' routers and could be first level access routers. Routers marked 'r2' connect to 'r1' routers, and so on. Also, the data set contains, in most cases, information about the geographical location of the routers, also derived based on heuristics. For the 6 maps we use, location information is given in all but one case for the backbone routers, and is missing for 845 out of the total 32,510 routers. Both of these heuristics are based on known patterns in the router names, as given by a reverse DNS lookup of the IP address.

The SCAN project map contains more than 200,000 routers. The publicly available data set contains router adjacencies for anonymized routers (SCAN 1999). In addition to this we obtained IP alias lists from the authors. When comparing the SCAN project map with the RocketFuel data we note that the SCAN project study was conducted in 1999, which is much earlier than the RocketFuel study in 2001. Thus, a comment is in place which relates to both the SCAN data set and to the list of ISPs in Table 2: with the turmoil in the industry over the last few years, some of these ISPs are now defunct or have been absorbed into other ISPs. As a first order approximation we assume that networks have changed hands rather than been dismantled and we retain the network ownerships as they were in 1999, assuming that although the owner may now be part of another business entity, the basic network structure remains.

Other than that, one should note the following differences: *i)* SCAN data lacks annotations about geographical locations and relation to backbone, *ii)* SCAN data does not include information about which AS (ISP) the routers belong to, and *iii)* RocketFuel study includes an attempt to estimate how large a fraction of each ISPs network was successfully covered by the tool (not the case for the SCAN data).

We decided to add annotations to selected parts of the SCAN project data set and try to turn it into the same format as the RocketFuel data.

2.2 Adding Annotations to SCAN Project Data

The SCAN project map we use was collected on Aug 8, 1999 and contains 228,263 routers. We are interested in certain ISPs in this map and pick a subset of it as follows. We use a BGP (Border Gateway Protocol) routing table dump, from the same date, from the RouteViews project (Meyer 2003). The routing table provides a mapping from IP prefixes to origin ASes. This simple method is not perfect, however: in 417 cases the mapping result is ambiguous and in 1219 cases a mapping cannot be found. Even so it manages to resolve an AS for more than 99% of the routers.

At this point we can select, one at a time, the ASes that we are interested in by picking the corresponding subset of the routers and their interconnections. In our case we are interested in supplementing our data with maps for some major ISPs like WorldCom (UUNet) and Cable&Wireless. Thus, we select one AS at the time and create a router map for that AS. Some large ISPs, such as WorldCom, break their network down into multiple ASes. WorldCom/UUNet uses AS numbers 701, 702, and 703. However, as far as we could determine, AS 702 covers Europe, and AS 703 covers Asia and the Pacific region. Since we focus on the U.S. part of the backbone, we only use 701 for our model. The networks we extract from the SCAN project map, and their sizes, are shown in Table 3.

Table 3: AS Maps Extracted from the SCAN Project Map, and Numbers of Routers Removed when Processing Them

ASN	Total Routers	Without Internal Links	Dis-connected	Remaining
701	9138	230	56	8852
3561	6016	256	148	5612
Total	15154	486	204	14464

It is also necessary to locate connections to neighboring ASes, so adjacencies to routers in other ASes are specially marked. In a manner similar to RocketFuel, we create a map per AS with external links to IP addresses indicating connection points to other networks (ASes). Since we do not have geographical location information at this point, we mark each router with the special marker used by RocketFuel to denote that the location could not be resolved by its heuristics. These markers will be replaced in a later step since we will need to find other means to fill in the missing location information in RocketFuel maps anyway.

In processing the data (per AS) we had to deal with a certain amount of "noise": *i)* Links are mostly asymmetric; we added links to form a symmetric graph. *ii)* Some routers lack links internally within the AS. This is most probably only a reflection of the incompleteness of the graph or missing data in the mapping from routers to ASes. We removed such routers. *iii)* Only routers that have "outgoing" links are listed the others are only "mentioned" as they are linked

to. We need to add the mentioned routers. *iv)* The graph (per AS) may not be connected. When this happens, we select the largest connnected component, and remove other routers. Table 3 shows how many routers we started out with, and how many of them that were removed in each step. The end result of our processing of the SCAN project map is a set of AS maps in the same format as the RocketFuel maps.

2.3 Merging RocketFuel and SCAN Maps

In a second step we need to merge together the multiple maps from RocketFuel and the SCAN project, hooking up the external links. Also, we fill in geographic data if it is missing (a few cases in the RocketFuel data and all cases for SCAN data).

To map the IP addresses to geographical locations, we use an Internet geographical database called NetGeo from CAIDA (CAIDA 2002b). The back-end of NetGeo is a database, that is used to cache the geographical information of IP addresses. The geographical information is derived heuristically by parsing the information retrieved from the *whois* server. NetGeo can be accessed either directly via the web or through a standard programming interface (both in Java and Perl). We found that, although most IP addresses can be resolved in this way, the data from NetGeo contains a fair amount of noise: there are IP addresses that cannot be resolved, and more problematically, some resolved locations are clearly incorrect. In those cases we compare the data with information obtained from the IP address location lookup service provided by Geobytes (Geobytes 2003). Data from Geobytes was primarily used in cases where *i)* none was obtained from NetGeo, or *ii)* NetGeo data (implausibly) fell back to company headquarters. There are still a number of IP addresses, less than 1.2%, that cannot be located using both methods. We simply discarded these routers from the network topology.

To merge the data from SCAN and RocketFuel, one also needs to find out the peering relationships between the ISPs. Fortunately, both SCAN and RocketFuel data sets contain IP addresses for routers belonging to adjacent ISPs. We have been successful in matching these foreign IP addresses with those in both data sets to identify the peering points.

2.4 Determining Link Bandwidths and Delays

The router-level maps of the Internet obtained from RocketFuel and the SCAN project does not contain information about links, such as bandwidths and delays. The lack of link characteristics motivated us to seek relevant information at each ISP that we intend to include in the study. Fortunately, all these ISPs we encounter have published network maps that contain, at least, information about connection types between major Points-of-Presence (PoPs), which can be then translated into link bandwidths. Some of this information has been collected by the Mapnet project at CAIDA and is publically available (CAIDA 2002a). Since our data contains geographical information for each router, we map the routers to known PoPs and then assign bandwidths from those given for PoP-to-PoP connections. However, this is imperfect at best, since we cannot distinguish between multiple trunk lines between two PoP sites. In these cases we simply pick bandwidths in a round-robin fashion, assigning the largest one first to ensure that most of the capacity is included. For links whose router locations are unknown and for those routers that are collocated, we apply some random distributions to assign bandwidths. We assign link delays as a function of link lengths based on speed of light. In the future, we would prefer to derive latencies from traceroute results.

Our method of assigning link bandwidths and delays is imprecise and still far from satisfactory. The quality of the assigned link characteristics depends very much on the quality of the published information we obtained from the individual ISP, which could be out-dated, and the transition from the macroscopic view to router-level Internet map contains a great deal of guess work and has not been tested. Thus, in the absense of other data sources (discussed in Section 4) this aspect of the model should be viewed with caution when used for simulation.

2.5 Configuring Routing Layers

A network model is more than just the topology of the physical connections. In fact, the traffic paths experienced by the user is determined not by the physical connectivity of the network, but by the logical connectivity determined by the pervasive dynamic routing protocols that run on top of this graph. For instance, BGP policies constrain the choice of inter-domain routes so that they frequently differ from the shortest path (Tangmunarunkit, Govindan, and Shenker 2001). Moreover, it will take some time for the routing protocol to respond to changes in the network so connectivity may temporarily be lost (Labovitz et al. 2001). Consequently, for a realistic model of connectivity it is necessary to have a realistic model of routing.

The SSFNet simulator contains detailed models of the Border Gateway Protocol (BGP) (Rekhter and Li 1995), the de-facto standard for inter-domain routing in the Internet, and OSPFv2 (Moy 1998), a commonly used intra-domain routing protocol. Thus, when building a simulation configuration description from the topology data, as in real life, it is also necessary to configure the routing protocols. (For simple models there is a simplified version of OSPF that does “static routing” which does not require configuration. Similarly, under certain conditions the BGP model can automatically configure itself to simplify model building. However, for

more complex models, such as the one we consider here and models that require dynamic intra-domain routing, it is necessary to manually configure certain aspects of the routing.) At this point we lack two important pieces of information: realistic BGP policies and realistic OSPF link weights. This will be discussed more in Section 4. Thus, as we build the simulation configuration script from our data sets we add keywords for "default" configurations of OSPFv2 and BGP. For OSPFv2 this means that the whole AS will be one OSPF area and that all the link weights are set to one. For BGP, we configure each of the routers that has a connection to some other AS to be a BGP speaker. Other routers run only OSPF. All BGP-speakers within the AS form a full mesh of internal peering relationships, and external peering sessions are configured with the neighboring BGP speaker in the other AS. The absense of any policy settings in BGP means that the operation essentially boils down to shortest AS path routing.

3 A STRAW-MAN MODEL OF THE U.S. INTERNET BACKBONE

The end result of this process is a model of the Internet backbone covering the United States, consisting of 8 national-level ISP networks. The total number of routers found in the U.S. part of the networks is 44,824 and there are 68,656 links. Also of importance for the simulation model building process is the number of BGP speakers in the model. In this model we found 1,185 routers that interconnect the ISPs, hence there are as many BGP speaking routers in the model. A break-down per AS is given in Table 4. Using the heuristic router markings of 'distance to backbone' we can pick a subset of the topology focusing on the backbone structure. The resulting network sizes are also shown in Table 4. This backbone subset is what we finally use to model the U.S. Internet backbone. Figure 1 shows the resulting network graph, where the geographical positions are mapped to the X-Y plane and each AS (ISP network) is given a different coordinate on the Z-axis. Thus, horizonal links (at the same "height") are internal links connecting routers within an AS, and vertical links connect different ASes. The network graph clearly shows certain exchange links (between ASes) spanning a great geographical distance. We find this puzzling. ISPs are generally said to connect at Network Access Points (NAPs) or Internet exchange points (IXs). For instance, one provider (MCI) operates large Internet exchange points, the "Metropolitan Area Exchange points": MAE West (San Jose, CA), MAE Central (Dallas, TX), MAE Chicago (Chicago, IL), and more. Lately, it appears that private exchange points have become more popular between the largest ISPs, but even so, we would expect to find that the links interconnecting different ISPs terminate in the same physical location. In fact closer examination of the data revealed that in the vast majority of cases, the termination points for exchange links were mapped to different geographic locations. This is a matter of ongoing investigations.

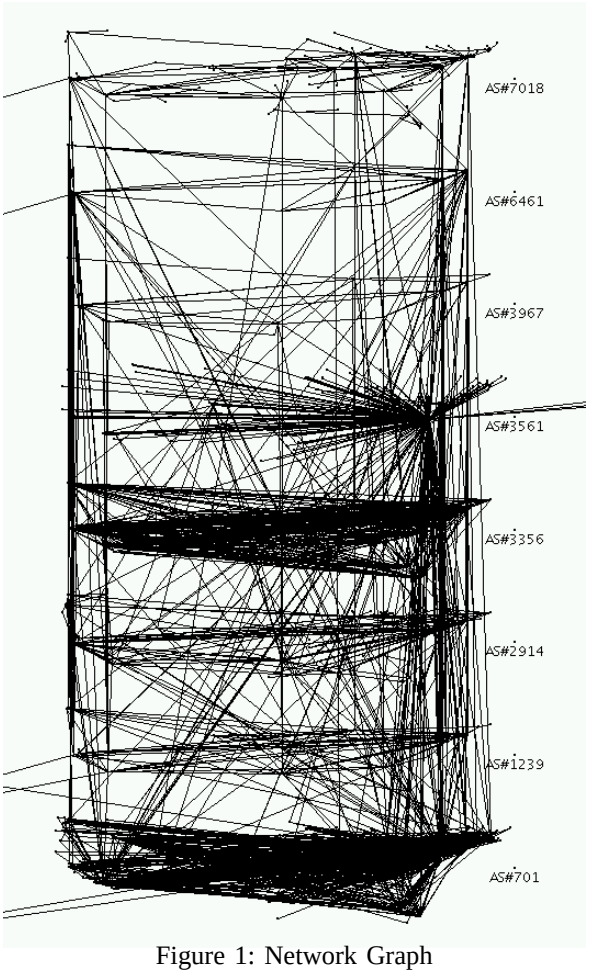

Figure 1: Network Graph

One question this raises is the quality of the geographical location information. In general, the heuristics used by RocketFuel for geographic locations appear to give reasonable results, although in a few cases the location was not provided. Indeed, the RocketFuel study verified the quality of their maps with a few ISPs and received positive comments. The results we got using other tools to add locations to other maps (Cable & Wireless and WorldCom/UUnet) and fill in the gaps were more problematic.

In constructing the model we discarded less than 1.2% of the total number of routers (not just backbone) because we were unable to map them to a geographical location. Thus, it seems that the geographical mapping facilities provide a mapping in most cases. However, the quality of the mapping is harder to determine except in some obvious cases. In a few cases it can easily be determined to be erroneous and for some networks it will obviously tend to fall back to the ISP's headquarter's address in many cases. In fact, AS 3561 (Cable & Wireless) appears to be exactly an example

Table 4: Total Network Sizes for Each AS in the Model

		Whole Network		"Backbone"	
ASN	ISP Name	Total Routers	BGP Speakers	Total Routers	BGP Speakers
7018	AT&T	11961	108	731	46
1239	Sprint	10531	171	497	56
701	WorldCom/UUNet	7983	263	4556	235
2914	Verio	6494	164	865	80
3561	Cable & Wireless	5429	261	2236	238
3356	Level3	1504	118	483	61
6461	AboveNet	490	59	247	39
3967	Exodus	432	41	213	32
Total		44824	1185	9828	787

of the latter problem. In Figure 1, AS 3561 appear appears to have a star topology which is clearly not correct judging by publicly available network maps (and would be highly implausible from an economic perspective). The reason is simply that more than 75% of the nodes are mapped to the same geographical location.

At this point it would be possible to discard AS 3561 from the model, in the absence of reliable geographic information. But we choose to keep it as the connectivity of the network is still correct. The effect of incorrect location information only affects the assignment of link attributes (delay in particular).

Another problematic point is the mapping of bandwidths to long-distance (inter-city) links. We found that we were only able to successfully map link end points to the network map containing bandwidth information for a minor fraction of the links. The majority of all links, in fact, turned out to be within PoPs where we had no information. And for the long-distance links the majority of all links could not be resolved to the map data we had. We believe the reason for this shortcoming is simply that the geographical granularity in the ISP maps (containing bandwidth information) is coarser than the naming in the router location information. Consequently, we are working on aggregating routers in larger metropolitan areas together to better match the ISP maps. But, even so, the ISP maps are largely unsatisfactory as a source of bandwith information and Section 4 discusses other possible sources.

Despite these shortcomings we have, in the end, a "straw-man model" of the U.S. Internet backbone. A model that we are making publicly available for download at `http://www.cs.dartmouth.edu/research/DaSSF/net-topology.html`, and as such has the distinction of being the only one available. Moreover, it serves as a starting point that we can keep improving upon as more data becomes available.

A "bare bones" simulation of this topology has been run using SSFNet 1.5.0, BGP running on border routers, and the simplified static OSPF on all routers. With only BGP convergence being simulated, i.e., no traffic, this model used approximately 4 GB of memory on a SUN Enterprise 6500. Naturally, adding LANs and application traffic will require additional memory.

4 THE NEED FOR MORE DATA

Even before undertaking this effort it was clear that there was a need for more data on the networks we attempted to model, but delving into it concretized this further. Here we briefly discuss some ongoing networking research efforts that could potentially be extended to provide missing pieces of information for our model. One aspect that needs improvement is some of the geographic location information in the model. But since this concerned only a subset of the model and commercial services exist that we have not explored yet we will not discuss it further here. Other aspects include:

4.1 Link Characteristics

Lack of link bandwidth information is one glaring issue with the model. Estimation of available bandwidth through active probing has been explored in studies such as Ratnasamy and McCanne (1999). Since link utilization is typically low on average (although bursty) and link bandwidth options are coarsely quantized, it appears reasonable to infer a link's bandwidth from a set of available bandwidth measurements. Improved link delay information is also needed but could be estimated from traceroutes. Finally, generating data for a complete network map is significantly more difficult than for a single link or path, but these methods could presumably be adapted similarly to other estimation techniques used for "Network tomography" (Bu et al. 2002) to obtain a network wide view.

4.2 Routing Configuration

For realistic traffic flows it is also necessary to model the effects of routing decisions, whether or not full routing dynamics are simulated. Studies such as Mahajan et al. (2002) attempt to infer OSPF (intra-domain) link weights as

an inverse problem from measured traceroutes. BGP (inter-domain) routing policies could, in principle, be obtained from Internet Routing Registries (IRR 2003) although this information may be incomplete or outdated since it is entered into the databases on a voluntary basis. A possible alternative approach would be to base policies on models of AS-to-AS relationships, such as the one proposed in Gao (2001).

5 CONCLUSIONS

Simulation is a key technology when experimentation with the real system is infeasible. In this spirit and in order to exploit advances in large-scale network simulation we desire to model, as accurately as possible, as a large a portion as possible of the Internet. Starting from router-level network maps generated by Internet mapping tools (here we used RocketFuel and Mercator generated data), we have attempted to build a model of the U.S. Internet backbone.

On the face of it, this may seem a straight-forward task, but in so doing we encounter numerous hurdles. We have to deal with a certain amount of "noise" in the data and many missing pieces of information for a realistic simulation model. Specific issues we encounter include: quality of sources of geographical location information, lack of information about link bandwidths and delays, and lack of routing configuration information (e.g., inter-domain routing policies and intra-domain link weights). We used simple PoP level network maps in an effort to assign reasonable bandwidths to long-distance links but find this unsatisfactory, largely due to problems with quality and inconsistent granularity of location information. We describe the resulting model, which has the distinction of being the only one available, and briefly review research efforts that we hope can fill in missing pieces of the puzzle.

ACKNOWLEDGMENTS

We thank Meiyuan Zhao for computing the AS overlay of the SCAN map. We also thank the RocketFuel project for making their data public, and Hongsuda Tangmunarunkit and Ramesh Govindan (SCAN project) for making their data available to us.

This research is supported in part by DARPA Contract N66001-96-C-8530, NSF Grant ANI-98 08964, NSF Grant EIA-98-02068, and Dept. of Justice contract 2000-CX-K001. Points of view in this document are those of the authors and do not necessarily represent the official position of the United States Department of Justice.

REFERENCES

Ahlin, L., and J. Zander. 1997. *Principles of Wireless Communications*. Lund, Sweden: Studentlitteratur.

Banks, J., J. Carson, B. Nelson, and D. Nicol. 2000. *Discrete-Event System Simulation, Third Edition*. Upper Saddle River, NJ: Prentice Hall.

Bu, T., N. Duffield, F. L. Presti, and D. Towsley. 2002, June. Network Tomography on General Topologies. *ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems* 30 (1): 21–30.

CAIDA. 2002a. Mapnet. `http://www.caida.org/tools/visualization/mapnet` [accessed July 14, 2003].

CAIDA. 2002b. Netgeo: Internet geographic database. `http://www.caida.org/tools/utilities/netgeo` [accessed July 14, 2003].

Calvert, K., M. Doar, and E. Zegura. 1997, Dec. Modeling Internet Topology. *IEEE Transactions on Communications*:160–163.

Cowie, J., D. Nicol, and A. Ogielski. 1999, Jan.–Feb. Modeling the Global Internet. *IEEE Computing in Science and Engineering* 1 (1): 42–50.

DRCL JavaSim 2003. `http://javasim.cs.uiuc.edu` [accessed July 14, 2003].

Faloutsos, M., P. Faloutsos, and C. Faloutsos. 1999, Aug. On Power-law Relationships of the Internet Topology. *ACM SIGCOMM Computer Communication Review, Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication* 29 (4): 251–262.

Gao, L. 2001, Dec. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking* 9 (6): 733–745.

Geobytes. 2003. `http://www.geobytes.com` [accessed July 14, 2003].

Govindan, R., and H. Tangmunarunkit. 2000. Heuristics for Internet Map Discovery. *Proceedings of the IEEE INFOCOM. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies.*, 1371–1380.

Haynal, R. 2000. Russ Haynal's ISP Page. `http://navigators.com/isp.html` [accessed July 14, 2003].

IRR. 2003. Internet Routing Registry. `http://www.irr.net/` [accessed July 14, 2003].

Labovitz, C., A. Ahuja, A. Bose, and F. Jahanian. 2001, June. Delayed Internet Routing Convergence. *IEEE/ACM Transactions on Networking* 9 (3): 293–306.

Law, A., and W. Kelton. 2000. *Simulation Modeling and Analysis, Third edition*. New York, NY: McGraw-Hill.

Mahajan, R., N. Spring, D. Wetherall, and T. Anderson. 2002, Nov. Inferring Link Weights using End-to-End Measurements. *Proceedings of the Second ACM SIGCOMM Internet Measurement Workshop*, 231–236.

McRobb, D., and K. C. Klaffy. 1998. Skitter. `http://www.caida.org/tools/measurement/skitter/` [accessed July 14, 2003].

Medina, A., A. Lakhina, I. Matta, and J. Byers. 2001, Aug. BRITE: An Approach to Universal Topology Generation. *Proceedings of the Ninth International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS)*, 346–353.

Meyer, D. 2003. University of Oregon Route Views Project. `http://www.routeviews.org` [accessed July 14, 2003].

Moy, J. 1998, April. OSPF Version 2. RFC-2328.

ns. 2003. The Network Simulator: ns-2. `http://www.isi.edu/nsnam/ns` [accessed July 14, 2003].

Ratnasamy, S., and S. McCanne. 1999. Inference of Multicast Routing Trees and Bottleneck Bandwidths Using End-to-end Measurements. *Proceedings of IEEE INFOCOM*, 353–360.

Rekhter, Y., and T. Li. 1995, March. A Border Gateway Protocol 4 (BGP-4), RFC-1771.

SCAN. 1999. `http://www.isi.edu/scan/mercator/maps.html` [accessed July 14, 2003].

Spring, N., R. Mahajan, and D. Wetherall. 2002, Aug. Measuring ISP Topologies with Rocketfuel. *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications (ACM SIGCOMM)*, 133–145.

SSFNet. 2003. `http://www.ssfnet.org/` [accessed July 14, 2003].

Tangmunarunkit, H., R. Govindan, and S. Shenker. 2001, August. Internet Path Inflation Due to Policy Routing. *Proceedings of SPIE ITCom 2001*.

Tangmunarunkit, H., R. Govindan, S. Shenker, S. Jamin, and W. Willinger. 2002, Aug. Network Topology Generators: Degree-Based vs. Structural. *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications (ACM SIGCOMM)*, 147–159.

AUTHOR BIOGRAPHIES

MICHAEL LILJENSTAM is a Research Associate at the Institute for Security Technology Studies and Computer Science Department, Dartmouth College. His research interests include large-scale network simulation, security, routing, and modeling and simulation of wireless networks. He received his M.Sc. (1993) and Ph.D. (2000) from the Royal Institute of Technology, Stockholm, Sweden. His e-mail address is <michael.liljenstam@dartmouth.edu>, and his web page is <www.cs.dartmouth.edu/~mili>.

JASON LIU is a Research Associate with the Institute for Security Technology Studies at Dartmouth College. His research focuses on parallel discrete-event simulation, performance modeling and simulation of computer systems and communication networks, and large-scale simulation of wireless ad hoc networks. He received B.A. in Computer Science from Beijing Polytechnic University in China in 1993, M.S. in Computer Science from College of William and Mary in 2000, and Ph.D. in Computer Science from Dartmouth College in 2003. His e-mail address is <jasonliu@ists.dartmouth.edu>.

DAVID M. NICOL is Professor of Electrical and Computer Engineering at the University of Illinois, Urbana-Champaign, and member of the Coordinated Sciences Laboratory. He is co-author of the textbook *Discrete-Event Systems Simulation*, and served as Editor-in-Chief at ACM TOMACS from 1997-2003. He will serve as the General Chair of the Winter Simulation Conference in 2006. From 1996-2003 he was Professor of Computer Science at Dartmouth College, where he served as department chair, and at the Institute for Security Technology Studies served as Associate Director for Research and Development, and finally as Acting Director. From 1987-1996 he was on the faculty of the Computer Science department at the College of William and Mary; 1985-1987 he was a staff scientist at the Institute for Computer Applications in Science and Engineering. He has a B.A. in mathematics from Carleton College (1979), an M.S. (1983) and Ph.D. (1985) in computer science from the University of Virginia. His research interests are in high performance computing, performance analysis, simulation and modeling, and network security. He is a Fellow of the IEEE. His e-mail address is <nicol@crhc.uiuc.edu>.