

High Level Internet Modeling

Chris White¹, Alan Weiss¹, Gerard O'Reilly¹, Rene LeClaire²

¹Bell Laboratories, Alcatel-Lucent, Summit, N.J.

²Los Alamos National Laboratory, Los Alamos, NM

E-mail: goreilly@alcatel-lucent.com, whitec@alcatel-lucent.com, apdoo@alcatel-lucent.com,

Abstract

This paper presents a hierarchical internet model suitable for simulating the impact of a localized physical disaster on the Internet within a typical metropolitan area. This first order model reports internet availability as the ratio of offered to carried traffic. We compute the carried traffic by examining the fraction of offered network traffic which flows through potentially impaired network entities. The impact of a physical disruption on a network entity reduces the carried traffic by the fraction of traffic which passes through that entity. The model is fully parameterized with respect to application properties (locality, latency, bandwidth, etc.) and connections to other infrastructure models (electric, phone, etc.). We also present a simplified system dynamics representation of this model using Vensim.

1. Problem Description

Our studies are part of a collaboration among Sandia National Laboratories, Los Alamos National Laboratories, and Bell Laboratories, to understand the cascading of impacts across infrastructures when natural or man-made disruptions occur [18, 19, 20, 21]. In recent years, the Internet has become a crucial component of our nation's communication network. The reliance of other critical infrastructures (power, transportation, communications, emergency services, etc.) on the Internet for monitoring, control, and scheduling elevates the Internet to critical infrastructure status. In addition, the inherent failure tolerance of internet routing allows the Internet to survive disruptions which might disable these other infrastructures making it an important tool for disaster recovery and impairment mitigation. For these reasons, a suitable model of the Internet

is required for assessing the impact of and recovery from a regional or nationwide disruption.

We present the details of a model suitable for simulating the impact of various events on the functioning of the Internet. These events could arise from a national disaster or from the failure of another critical infrastructure like the electrical power grid. The utility of such a model will be to quantify the impact of and recovery from a nationwide or regional disruption of the internet infrastructure. Because nationwide critical infrastructures are often intimately linked, the simulation of these other critical infrastructures requires an internet model that reflects the interdependencies.

We develop a very high level regional model which does not require detailed knowledge of internet topology or statistics.

1.1 Critical Infrastructure Modeling

The CIPDSS (Critical Infrastructure Protection Decision Support System) project at Argonne, Los Alamos, and Sandia National Laboratories has developed a risk-informed decision support system that provides insights for making critical infrastructure protection decisions by considering all critical infrastructures and key resources, and their primary interdependencies. Initiated as a proof-of-concept in August 2003, the CIPDSS project has demonstrated how it will assist decision makers in making informed choices by a) functionally representing all critical infrastructures and key resources with their interdependencies; b) computing human health and safety, economic, public confidence, national security, and environmental impacts; and c) synthesizing a

methodology for decision making that is technically sound, defensible, and extendable.

System dynamics consequence models representing the key infrastructures were built using Vensim [6]. The consequence models simulate the dynamics of individual infrastructures and couple separate infrastructures to each other according to their interdependencies. Dynamic processes like these are represented in the CIPDSS infrastructure sector simulations by differential equations, discrete events, and codified rules of operation.

The initial CIPDSS prototype used nearly 5000 variables to simulate the dynamics of the critical infrastructures and key resources at the national and metropolitan scales: many of these variables are output metrics estimating the human health (e.g., deaths from an event), economic (monetary damage), or environmental effects (e.g., air contamination) resulting from disturbances to the infrastructures.

In addition to contributing a valuable stand-alone representation of Internet structure, function and vulnerability, the Internet model adds an important component to the CIPDSS suite of coupled infrastructure models. Together these models can examine the potentially complex dynamics and interdependencies between the internet and other critical infrastructures.

1.2 Applicability of the Internet Model

Constructing a globally applicable model that simulates the detailed protocols, topology, and scope of the Internet is a difficult challenge. To create a tractable simulation model, several limitations need to be applied to the scope of the problem. We describe a geographically and connectivity based model suitable for simulating the impact of a localized physical disaster. This includes such events as a hurricane, earthquake or terrorist attack causing damages within a typical metropolitan area. The model is not suitable for the simulation of “Cyber Attacks” or network level impairments which require a detailed

knowledge of Internet protocols or a network level view of network traffic.

The range of treated applications and the coupling to other infrastructure models is very general. It allows a broad range of network applications through a comprehensive parameterization of application properties. Future applications can be incorporated by adjusting the values of the parameters. In a similar fashion, we incorporate a broad range of infrastructure impairments through connections to other infrastructure models.

2. Network Model

The model is a hierarchical structure describing a typical point to point connection in the Internet. This represents the use of a single application connecting a single user with a second endpoint in the network. We use the natural aggregation points of the Internet to subdivide the path for this connection (Figure 1).




Figure 1 – Network Model: the arrows indicate parametric coupling between network entities.

Conveniently, these aggregation points for a typical metropolitan area coincide roughly with geographically relevant subdivisions. In the spirit of defining a hierarchical model with the potential for various levels of complexity, we propose each subdivision will have its own model for performance and potential impairments.

The rapid growth of new applications and the ongoing evolution of existing applications require the model incorporate a simple mechanism for creating new applications and modifying the properties of pre-existing applications. To accomplish this, we parameterize internet applications based upon a limited set of properties and requirements for the applications.

To couple this internet model with the other infrastructure models, we include a similar parameterization with respect to quantities computed by the other infrastructure models. For example, the loss of electrical power will limit the use of certain access mechanisms or change the details of packet transport in the network.

Figure 1 shows the hierarchical point to point model of the Internet. Each network entity of this model will include an independent model for operation. The individual models are parametrically coupled allowing the sophistication of treatment to vary with each network entity. The following subsections describe each entity in more detail.

2.1 User

The user represents the starting point for all network traffic in the model. The model currently assumes a single typical user which greatly simplifies the computation of internet traffic profiles. Obviously, future enhancements could further subdivide this user element to create different types or classes of users (home, business, affluent, impoverished, young, old, etc.) Each of these classes might have different access mechanisms, application usage profiles, as well as expectations on internet quality of service (QoS). Although it is possible to incorporate this level of detail, it is not clear that without the use of sophisticated models which describe the actions of each type of user within the context of a physical disaster that the additional detail would yield significantly better overall internet traffic profiles. It is for this reason, that we have based the first order model on a single typical user.

2.2 Access

The user has many possible access mechanisms. These range from low-speed, inexpensive, widely available access mechanisms like dialup to broadband access mechanisms like cable modems and Digital Subscriber Line (DSL) access (Figure 2). Each access mechanism has different interactions with other infrastructures as well as different types of disruptions arising from a physical

disaster. The model must incorporate how a single disruption impacts each access mechanism.

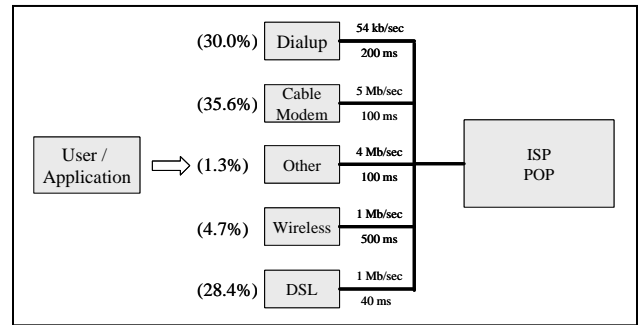


Figure 2 – Access Methods

The model focuses on the fraction of users selecting each access mechanism. These fractions will change not only with the different access methods selected, but also with the type of application. For example, users will rarely use a low bandwidth connection to access an internet application which requires large amounts of data transfer (i.e. streaming video). These changes in access fraction need to be adequately adjusted for the model to provide a realistic view of the losses in internet traffic arising from a disruption.

The fraction of users selecting specific access mechanisms can be determined through an exploration of publicly available data. Figure 3 shows a table and plot of the fraction of broadband users in the United States subdivided by access mechanism [1]. This information with information describing the total number of internet access points (roughly 70 M) [2] can be used to estimate the access fractions for our typical user. Similar statistics can be obtained at the international level [3].

Table 1
High-Speed Lines¹
(Over 200 kbps in at least one direction)

Technology ²	1999	2000	2001	2002	2003	2004	2005	
	Dec	Dec	Dec	Dec	Dec	Dec	Jun	Dec
ADSL	369,792	1,977,101	3,947,808	6,471,716	9,509,442	13,817,280	16,316,309	19,514,318
SDSL and Traditional Wireline	-	1,021,291	1,078,597	1,216,208	1,305,070	1,468,566	898,468	876,286
SDSL	-	-	-	-	-	-	411,731	366,376
Traditional Wireline	609,909	-	-	-	-	-	486,737	509,910
Cable Modem	1,411,977	3,582,874	7,059,598	11,369,087	16,446,322	21,357,400	23,936,536	25,583,233
Fiber ³	312,204	376,203	494,199	548,471	602,197	697,779	315,651	448,196
Satellite and Wireless	50,404	112,405	212,610	276,067	367,118	549,621	965,068	3,809,247
Satellite	-	-	-	-	-	-	376,837	426,928
Fixed Wireless	-	-	-	-	-	-	208,695	256,538
Mobile Wireless	-	-	-	-	-	-	379,536	3,125,781
Power Line and Other	-	-	-	-	-	-	4,872	5,859
Total Lines	2,754,286	7,069,874	12,792,812	19,881,549	28,230,149	37,890,646	42,436,904	50,237,139




Figure 3 – Access Method Statistics

One aspect not addressed by the publicly available statistics concerns the number of individuals utilizing multiple access mechanisms. This information is difficult to determine; however, the impact of multiple access mechanisms is considered second order when compared to the terms included in this model. Future efforts to understand how a user or the network responds to a network impairment will increase the importance of multiple access mechanisms. For example, if a user can not use their primary access mechanism, how rapidly and to which secondary access mechanism they transition will be important. This becomes exceptionally difficult to anticipate because the number of secondary internet access mechanisms would be very large for a typical user (access from work, school, public library, internet café, cell phone, etc.)

The fractions presented here and used within the model should be regarded as a current best estimate typical of the entire United States. As time passes, or if the model is applied to a specific regional area, it is important that these values be

updated to express the best current information. Due to the nature of the fractions, they will change very rapidly in time as new access methods are introduced, or as old access methods become more or less commonplace. They will also change with region based upon implementation and deployment plans of the regional access providers.

For example a metropolitan area with a large population density will have a higher fraction of users using broadband access mechanisms than a sparsely populated rural area.

2.3 Internet Service Provider (ISP)

The internet service provider (ISP) provides the connectivity from the various access mechanisms to the broader Internet. In addition, the ISP is typically responsible for billing the end user, providing customer service, and providing several necessary infrastructure services (Domain Name Service (DNS), email, Usenet (NNTP), Web hosting etc.). Although many ISPs are nationwide in scope, we are primarily interested in the portion

(network hardware and servers) of a nationwide ISP which provides network access and services to a specific region. Figure 4 shows a list of the largest US ISPs with rough subscriber numbers [7]

ISP	Subscribership
America Online	22,200,000*
NetZero	8,600,000
Comcast	8,142,000
Microsoft Network	8,000,000
Spinway	6,700,000
United Online	6,600,000***
SBC Communication	6,496,000****
EarthLink	5,400,000
RoadRunner	4,557,000
Verizon	4,531,000****
Prodigy	3,500,000
1stUp.com	3,500,000
Freei.Net	3,200,000
AltaVista (via 1stUp)	3,000,000
Cox Communications	2,975,000
Bell South	2,678,000****
Charter Communications	2,120,000
Compuserve (AoL)	2,000,000*
Adelphia	1,656,700
Cablevision	1,600,000
AT&T Worldnet	1,500,000

Figure 4 – Nationwide ISP

2.3.1 ISP Point of Presence (POP)

The ISP point of presence (POP) is the primary connection from an access medium to the ISP's backbone network. A POP might consist of banks of modems providing dialup access or a head-end switch of a local cable modem or DSL provider. Generally, the POP provides the location and mechanism for aggregating the traffic from a large number of end users to the ISP backbone network.




Figure 5 – ISP Entry to Network

Depending on the nature of the traffic, the ISP backbone may represent the final destination of the traffic (i.e., access to email or another ISP provided service) or as a conduit for traffic destined for the wider Internet. The traffic not local to the ISP will travel either to the Internet core or to a regional internet exchange point (IXP) (see Figure 5). The transport to the internet core occurs via a direct connection with a Tier 1 Internet provider. The location of this connection could occur at physical location of the POP, or within a centralized aggregation point located on the ISP's backbone network. These details depend on the size and traffic served by the ISP as well as the extent of the ISP's backbone network. Since the ISP will pay for any traffic carried by another provider, they have substantial financial incentives to carry the traffic on their own backbone when possible. A second manner of saving transport costs arises when the traffic is meant for another ISP located within the same geographic region. In this case, the traffic will travel to a local internet exchange point (IXP), saving the transport costs paid to the Tier 1 provider.

2.3.2 ISP Infrastructure

In addition to the transport of traffic from an access method to the broader Internet, the ISP also provides several critical services. The ISP infrastructure consists of the services provided by a typical ISP beyond the direct transport of packets (Figure 6). In some cases, the removal of these services effectively causes the loss of connectivity even if the primary packet transport mechanisms remain functional. For purposes of this task, we assume the disruptions of ISP services globally reduce the aggregate traffic passing through the ISP. Future models will subdivide these services in a manner which can impact the traffic fractions in a much more detailed manner. For example, the disruption of local services (mail, web hosting, etc.) will have little or no impact on packets destined for the internet core. In contrast, the loss of the Domain Name Service (DNS) or the ISP network operations center

(NOC) might drastically limit the flow of any packets beyond the local ISP.

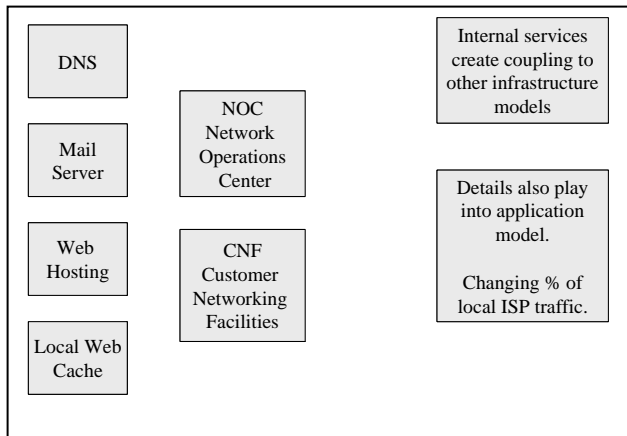


Figure 6 – Typical ISP Services

2.4 Internet Exchange Point (IXP)

An Internet exchange point (IXP) is a location which hosts network connections within a geographic region where ISPs within that region can exchange network traffic. This traffic will remain local to that region and will not incur transport costs for traveling to the internet core [11]. IXPs provide a physical location where ISPs can place network equipment and connect to the IXPs switching fabric to exchange traffic.

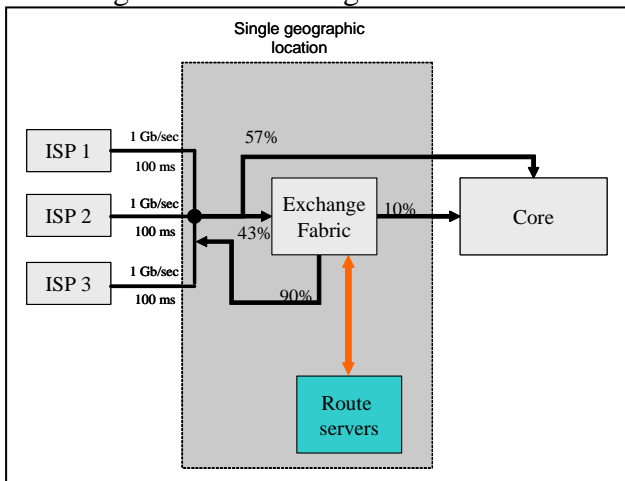


Figure 7 – IXP Functions

An IXP can be a purely commercial entity run by a provider who in addition to providing peering

between ISPs may also provide packet transport at a national level. See as an example the site location map for PAIX [8] a nationwide IXP. They provide multiple peering points across the US. In addition, they will provide transport to the other peering points within their network.



Figure 8 – PAIX IXP

Several non-commercial IXPs exist which provide access to regional ISPs purely based upon the cost savings to the local ISPs. These IXPs typically provide peering only. No packets destined for the wider Internet will be exchanged.

- **Major geographic peering points**
 - San Francisco Bay, Washington DC, Chicago Region, Los Angeles, Dallas, New York Region
- **Commercial**
 - Verizon
 - MAE-East, MAE-West, MAE Central, MAE Miami
 - Equinix
 - Chicago, Dallas, Washington DC, Honolulu, Los Angeles, New York City, Silicon Valley
 - Switch and Data (PAIX):
 - Atlanta, Dallas, New York City, Palo Alto, Philadelphia, San Jose, Seattle, Northern Virginia
- **Other smaller regional peering points**
 - Pittsburgh, Blacksburg, Oregon, San Diego, Seattle, New Mexico, Indianapolis, Columbus, Utah, Dayton

Figure 9 – Nationwide IXP

Traffic details of the commercial IXPs [8,9, 10] are difficult to obtain. However, non-commercial IXPs [11,12,13,14,15,16] often provide extensive

details on the switching fabric as well as traffic carried. These numbers can be used to determine a typical packet traffic fraction. However, depending on the density of Internet users within a region, the numbers for a specific region may vary from the norm.

2.5 Core

Within the context of this model, we assume once a packet reaches the internet core it will eventually reach its final destination. This assumption is predicated on the fact that there are multiple Tier 1 providers which provide the national level transport. See, for example, Figure 10 showing the network map from UUNET [map from 17].



Figure 10 – Example Tier 1 network map (UUNET)

Disruptions of the core would require detailed knowledge of the backbone networks of the Tier 1 providers. Although some public maps exist, comprehensive knowledge of the networks of Tier 1 providers is considered highly proprietary information. Since we focus primarily on developing a regional model, we will not model the impact of a disruption on the core. Future models may address this deficiency by linking together multiple IXPs with links having the typical properties of the internet core. However, even within such a model, the existence of multiple Tier 1 routes between any two IXPs will lead to a “core” link which is very robust to physical disasters.

3. Internet Availability / Model Parameterization

For simplicity, we define internet availability as the ratio of carried to offered traffic in the network. In this model with no impairments, all offered traffic will be carried traffic. This assumes the network is designed and provisioned with the capacity to operate at this level during normal operating conditions. Obviously, this is not true in general since the growth in network traffic continues to fuel network improvements. However we are primarily interested in the variations of the network properties away from the normal operating conditions. Thus we assume the network is adequately provisioned.

As we add impairments to the network some fraction of network traffic will not be carried which in turn will decrease the internet availability. The decrease depends in detail on the nature of the application as well as the geographic location of the impairment within the network. In the future we hope to extend and expand this definition to include quality of transport metrics rather than a simple, binary carried / not carried decision. Additional work on computing carried traffic on the network will further refine the definition of availability, expanding the number of computed output quantities in the model.

3.1 Offered to Carried Traffic Conversion

The current model defines carried traffic as the fraction of offered traffic which has an unimpaired path through the network. To compute the carried traffic, we must determine the fraction of traffic which travels through each network entity. This fraction will change based upon the nature of the application. For example a local application will typically have most of the network packet streams exit at the ISP without ever traveling deeper in the network. Figure 11 shows two different paths through the network for a single application. In this case, 23% of the packet streams exit at the ISP and 34% of the packet streams proceed to the core. These per-

centages represent the fraction of application sessions which travel various paths of the network, not the fraction of packets within a given user session.




Figure 11 – Fractions for two paths through the network

Figure 12 shows a full set of fractions for all traffic paths in the model for a typical internet application. The percentages reported within each blue circle show the fraction of packets which exit from this point in the network. In this example 30% (9+7.5+3.0+7.5+3.0) of the packet streams are terminated at the ISP. The sum of the values in all blue circles is 100%.




Figure 12 – All routes through the network for one application

3.2 Base Routes

The base route (Figure 13) expresses how a typical packet would travel through the network. Initial parameters were obtained by looking at typical traffic flows and network design parameters –

assuming the network design capacity fractions roughly follow network utilization fractions.

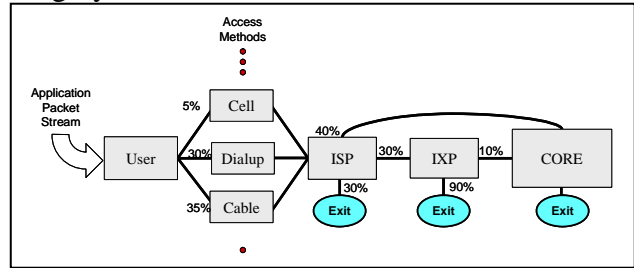


Figure 13 – Basic Routes

3.3 Application attributes

Application attributes parameterize an application with respect to typical application requirements. Based upon these quantities, traffic route fractions are modified in a manner consistent with methods of achieving the attribute. By defining a comprehensive set of attributes, we enable the model to work for a wide variety of applications as well as for subtle changes to the standard set of applications. In general each application attribution has a numerical value on the interval [0,1]. Typically, a 0 would indicate the attribute is not important at all for the given application and a 1 would indicate the attribute is critically important. A value of 0.5 would indicate an importance consistent with the importance of the attribute to normal web browsing. Thus, using this scale, importance is quantified in direct comparison with web browsing.

For the initial model, these numbers enter as a simple linear relationship with the fraction of traffic carried. As an example, when the locality parameter is set equal to 1 (maximally local) the fraction of traffic traveling to the internet core will be zero. When the locality attribute is set to 0.5 (and thus matches typical web browsing) the core fraction will be equal to the base fraction. So the linear relationship is determined using two endpoints, the base fraction at 0.5 and a 0 or 1.0 fraction at one of the extremes (0 or 1) of the attribute value. The selection of the extremum point (0 or 1) and traffic fraction depends on the nature of the attribute.

For the locality parameter, the equation representing the fraction of core traffic will thus be given by:

$$\text{CoreFraction} = \text{CoreBaseFraction} (1 + 2(0.5 - \text{locality})) + \text{Other Terms}$$

The *Other Terms* in this case would be terms arising from the other application attributes. The locality coefficients for all of the fractions (Core, ISP, IXP, etc) are selected such that the sum of all of the traffic fractions will sum to 1.0. Unfortunately, with independent linear relationships, the end result of adjusting more than one of the attributes for the same application will lead to sums greater or less than 1.0. In this case, the individual fractions should be scaled such that the sum again equals 1.0. Since we have assumed independence between the attributions, it is also possible to obtain negative traffic fractions due to the contributions of the other terms arising from other attributions. In this case, we restrict the range of fraction values to [0,1]. Again this will require scaling the individual fractions to achieve a sum equal to 1.0.

The use of linear relationships has many deficiencies which will be addressed in future enhancements to this model. Obviously in the future, the dependence on and interdependence between these various attributes can be included in the model. For the present, for applications similar to web browsing, or applications which deviate from web browsing in a single, well defined attribute, the current model adequately represents the changes to the traffic fractions with a degree of precision adequate for understanding the overall changes to internet availability. Future models may include a non-linear (although probably still polynomial) model for determining the traffic fractions as well as including the direct couplings between attributes. This would allow a cleaner treatment of the boundary cases alleviating the necessity of the scaling described above.

Table 1 shows an example set of typical application profile weights.

	Local-ity	Mobil-ity	Band-width	La-tency	Secu-rity	Cost	ISPI-nf
Web	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Email	0.7	0.1	0.1	0.1	0.5	0.1	0.8
VOIP	0.8	0.1	0.3	0.8	0.7	0.6	0.8
File Shar-ing	0.4	0.1	0.9	0.1	0.1	0.1	0.1
Stream-ing Media	0.4	0.4	0.8	0.7	0.1	0.1	0.5
Data Backup	0.7	0.0	0.8	0.1	1.0	0.6	0.6

Table 1 – Application Profile Weights

In the next sections, we describe the definitions of the application attributes included in the model.

3.3.1 Locality

Locality refers to the relationship between the application user and the end point of the communication. Specifically it refers to how far within the network traffic packets must travel. For example, a local application would have a large amount of traffic remain within the local ISP. A non-local application would have most traffic travel to the core. As an example of locality within the current Public Switched Telephone Network (PSTN), the fraction of all calls that are long distance is approximately 15%.

3.3.2 Mobility

Mobility refers to the need of the user to be mobile. This would have a tendency to favor mobile access mechanisms in comparison to fixed line access. An example would be an application for which most users would complete using wireless or cell phone access such as a location based directory service, or a GPS based map service.

3.3.3 Bandwidth

The bandwidth parameter refers to the bandwidth needs of the application. Applications requiring large amounts of bandwidth will shift the traffic to broadband access mechanisms.

3.3.4 Latency

The latency parameter refers to the effect of delay, latency and lag on the application. Applications such as streaming media or VoIP which require low latency will shift the traffic to lower latency access mechanisms. For example, real time internet gaming is not possible over a high latency wireless link.

3.3.5 Resiliency

The resiliency parameter refers to the ability of the application to survive dropped packets.

3.3.6 Security

The security parameter refers to the need of the application packets to remain secure. Although VPNs exist for ultra secure packet transport, typically users would favor wireline access mechanisms over wireless when transmitting something of a secure nature.

3.3.7 Cost

The cost parameter refers to the nominal valuation of the packet transport. Expensive packet streams would typically be transported over a more robust link. This could include secure banking transactions, the backup of corporate servers, or similar transfers which immediately impact the ability of a person or a corporation to make money.

3.4 Infrastructure Restrictions

Depending on the status of various infrastructure parameters, paths may be reduced or even eliminated. For this initial model, we've included coupling only to the power and telecom infrastructure models.

3.4.1 Power

The power infrastructure potentially has a significant impact on the status of the various network entities in our simple model. In particular, a power blackout means that some entities will go down after battery back-up facilities run out. See [22] for a description of a major blackout in 2003.

3.4.2 Telecom

The telecom infrastructure can be disrupted by natural disasters, such as hurricanes or flooding, or by acts of sabotage. This can limit the number of access mechanisms available to users. An example telecom disruption is shown in Section 5.

3.5 Infrastructure Failures

In addition to the infrastructure limitations created through the coupling to the other infrastructure models, parameters have been incorporated to simulate a direct failure of a network entity. In this model, we have enabled the direct failure of the ISP, IXP, and ISP infrastructures. Again these parameters enter as a value selected from the range of [0,1]. This is included as a direct multiplier for the associated traffic fraction. So a value of 0 will mean that the entity has completely failed, and the carried traffic through this entity will be 0. A value of 1 will allow all offered traffic to be carried.

4. Vensim Model

The model presented in the previous section was built in Vensim [6].

The model has 4 views as depicted in Figure 14, Figure 15, Figure 16, and Figure 17:

- Main model
- Application attributes
- Infrastructure restrictions
- Wireline and wireless infrastructure with/without power lifeline

These views can model each application (email, web browsing, and VoIP) separately, and end by giving the internet availability for each application. Note that we are not limited to just three ap-

applications. We can model with the same structures as many applications as the user wants. On the left side of Figure 14, we have our population of 5 million people, which then divides into households and businesses. Each of these has various rates of instances per day for the internet applications under consideration (web browsing, email, and VoIP).

Figure 15 shows the application attributes (just 4 for now, bandwidth, latency, mobility, and locality) and their influence on the access fractions (DSL, Dialup, Cable, Wireless, Other). The Base % for each is based on current values for these fractions. Each of these variables is an array depending on the application (just 3 for now, web browsing, email, and VoIP).




Figure 14 - Main Model

This is then broken out into the offered load by the various access methods (dialup, DSL, Cable, wireless, and other). Then the offered traffic by access method is divided up into ISP, IXP, and Core carried traffic. Finally, at the lower right, we calculate the internet availability for each application.




Figure 16 - Infrastructure Restrictions

Figure 16 shows various infrastructure restrictions resulting from a telecom disruption or an electric disruption (blackout). Each of these restrictions affects the various access modes differently. Additional work needs to be done on how the ISP, IXP, and Core Network infrastructures are affected by a blackout.

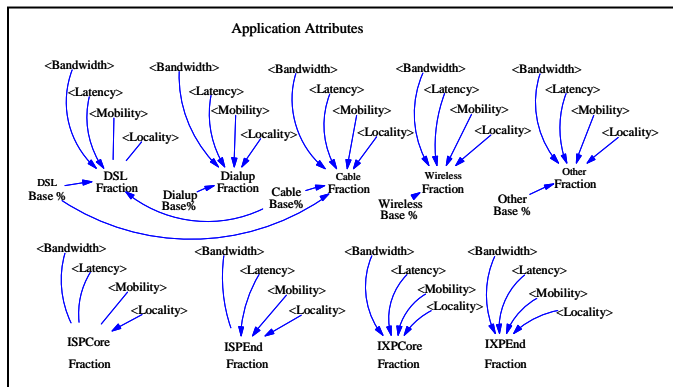


Figure 15 - Application Attributes

Figure 17 shows how the wireline and wireless infrastructures are subdivided into those sets of lines or mobile subscribers that have a power lifeline or not. Note that we have excluded business subscribers for simplification for now.




Figure 17 – Wireline and Wireless infrastructure with power impacts

Having a power lifeline means that the wireline or wireless service will or will not work in the event of a blackout. For example, since wireline central offices usually have diesel generators for emergency back-up power, wireline service still works in a blackout. The results from this view are fed back into the view on Infrastructure restrictions (Figure 16) to derive the impact on the telecom infrastructure of a power blackout.

5. Some Example Simulation Results

Here we postulate several interruptions to the internet to demonstrate several key aspects of the model. The examples given are meant to be illustrative only – using this and future versions of the model with the CIPDSS models the user will be able to tailor similar examinations for a wide variety of interesting cases. For the examples shown the Internet model was merged with the CIPDSS energy and telecommunications models. The interruptions are set up in the CIPDSS models, passed to the Internet model and allowed to feed back into the CIPDSS models if appropriate.

5.1 Baseline – no telecom disruption and no blackout




Figure 18 – 24 hour profile of internet activity.

Figure 18 shows a typical traffic profile for telecommunications peaking during working hours.

Figure 18 shows the number of instances (in millions) of each of the applications (web browsing, email, and VoIP) across the day (1440 minutes) for the whole metropolitan area of 5 million people.

We assumed that each household had 5 web browsing sessions, 10 emails, and 15 VoIP calls per day. Each business had 2 web browsing, 25 emails, and 10 VoIP calls per day. These numbers were made up simply for illustrative purposes.




Figure 18 - Baseline demand for a day across applications

The internet availability is equal to 1.0 for all three applications in the baseline case since there are no disruptions.

5.2 Telecom Disruption

In this scenario, we assume that the wireline telecom network is disrupted. As an example, Figure 19 shows an illustrative telecom disruption created in the CIPDSS telecommunications model with the interruption starting at hour 2 and recovering over the remainder of the day.

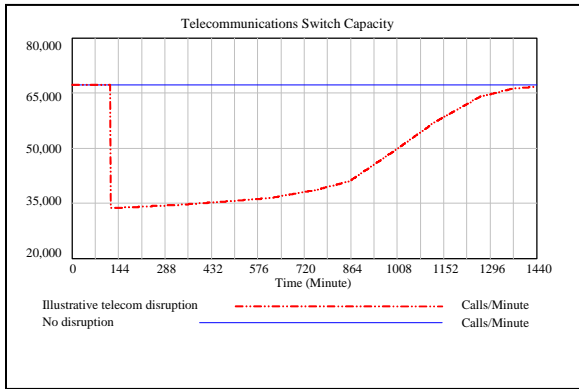


Figure 19 – Telecom Network Damage

Figure 20 shows that the telecom disruption affects on the web application with and without the disruption. The other applications for VoIP and email, look similar. Note that internet activity is not affected over as long a portion of the day as that of the telecommunications interruption of Figure 19 because the offered traffic in off-peak hours is so low.




Figure 20 - Carried traffic with telecom disruption

Figure 21 shows the resulting impact on internet availability from the telecom disruption for the web browsing and VoIP email application. The Email application exhibits similar behaviour.




Figure 21 - Internet Availability with telecom disruption

5.3 Telecom Disruption + Blackout

Figure 22 shows the impact of an assumed power blackout as created in the CIPDSS electricity model due to an interruption in electricity distribution capacity.




Figure 22 - An illustrative energy disruption.




Figure 23 – Carried traffic with energy disruption

Figure 23 shows a comparison of internet activity with and without the illustrative electricity disruption using the web channel as an example (the other channels have similar behavior). Note that the energy disruption scenario illustrated here has a stronger effect on internet activity.




Figure 24 – Energy and disruption

Figure 24 shows a comparison of internet activity with and without a combined telecommunications and electricity disruption using the web channel as an example (the other channels have similar behavior). The combined effects are most visible during peak telecommunications offered traffic.




Figure 25 – Internet availability with combined disrupts

Figure 25 shows a comparison of internet availability with this combined telecommunications and electricity disruption for the web and VOIP channels (the email channel has similar behavior to web).

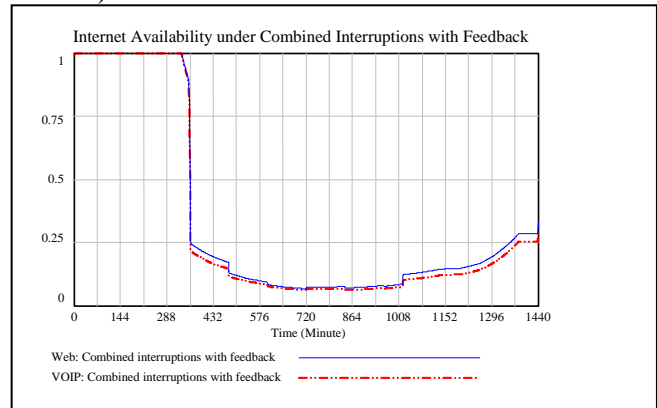


Figure 26 – Combined interruption with Feedback

Figure 26 shows a comparison of internet availability with and without a combined telecommunications and electricity disruption plus feedback for the web and VOIP channels (the email channel has similar behavior to web). The feedback is an illustrative connection between internet availability and SCADA (Supervisory, Control and Data Acquisition) operations – when internet is degraded, it is postulated to cause SCADA operation interruptions in the electricity sector which causes problems with electricity distribution. This electricity disruption then feeds back into the

telecommunications and internet models. Clearly under this postulated example the impact is the most severe case shown with the system unable to completely recover during the first day.

6. Possible Future Directions

The current model includes several approximations which we expect to remedy in future work.

- Offered to carried traffic computation - Currently, carried traffic is computed based solely on the fraction of traffic traveling through the various network entities. Ideally, the carried traffic computation should incorporate the numerous constraints on network traffic flow. This includes limitations on network transport as well as the limitations caused by processing required within each network entity. We may also include multiple levels of carried traffic quality. For example, the traffic carried may be suitable for some applications but unsuitable for others.
- Parameter estimation and validation - The quality of the model depends strongly on the quality of the underlying values used as parameters in the model. To assess the validity of these parameters, we envision a comparison of computed quantities against detailed network level simulations. Such detailed network level simulations serve a dual purpose to determine the model parameters as well as to validate the nature of the parameterization. The design of the model allows the treatment of individual network entities at different levels of approximation. In this manner it is possible to increase the level of sophistication where it is most needed while allowing other entities to be treated with much less detail.
- Time dependence and network response - The current model simulates the static response of the network to the disruption of a network entity. A more realistic model must include some level of time dependence as well as some form of network response. For example in a real network, one would expect traffic to be routed around a potential impairment. Future work will attempt to include this response of the network to a disruption. To first order, this response can be computed through the addition a simple iteration of the current model. Beyond first order, the addition of network level simulations of the network entities would introduce time dependent congestion in the network. In addition to the network response, time dependence of the application usage patterns may play an important role in future computations. For example, when a user finds he cannot complete a VoIP call, he may instead attempt to send an email or and instant message. Similarly, a user which cannot reach the internet core through a cable access connection may attempt to use an alternative access mechanism. Such user level behavior may be difficult to quantify in a typical case, but understanding when they may be important represents an important area of exploration.
- Network topology - The current model reflects the minimal useful network topology. Future efforts will add more structure to the network model. This includes the addition of both details to the current network elements as well as extending the reach of the current network. For example, the point to point model could be extended to include multiple IXPs or multiple ISPs within a single region. The model could further be extended by defining a connection between two regional models. For additional detail, the ISP backbone or an IXP switching fabric could be simulated with more structure. This type of detailed topology would be coupled with more sophisticated models for computing carried traffic related to the constraints on traffic throughput.

- Applications - The current model presents a parametrization suitable for simple application profiles. Future efforts will expand the set of application attributes leading to a better expression of the application profiles. In addition, the more sophisticated models described as other future work will allow the addition of QoS metrics to the model. This will allow a more subtle description of application availability which could indicate the quality which the network can support a given application.

A second possible application enhancement to the model concerns the coupling of different applications in the network. Currently applications are assumed to interact independently with the network. This ignores the transition of one application to another based on network congestion. For example, the inability to make a VoIP call might lead to an increase in email or instant messaging traffic. Obviously, this coupling will be important as we allow the entities in the network to respond to a network disruption. This would also allow the increased traffic in one application to cause network congestion which impacts the service level of another application.

7. References

- [1] "High-Speed Services for Internet Access: Status as of December 31, 2005" – Industry Analysis and Technology Division, Wireline Competition Bureau, July 2006
- [2] www.cia.gov/cia/publications/factbook/geos/us.html
- [3] www.oecd.org/sti/ict/broadband
- [4] Automatic Reporting Management Information System (ARMIS), FCC website, <<http://www.fcc.gov/wcb/armis/>>
- [5] FCC Report 05-173A1, September 30, 2005, Annual Report and Analysis of Competitive Market Conditions With Respect to Commercial Mobile Services.
- [6] Ventana Systems, Inc. www.vensim.com
- [7] www.jetcafe.org/~npc/isp/large.html
- [8] www.switchanddata.com
- [9] www.mae.net
- [10] www.equinix.com
- [11] www.ep.net/ep-main.html
- [12] www.pitx.net
- [13] www.atlantaix.com
- [14] www.mass-ix.net
- [15] www.nyiix.net
- [16] www.oregon-ix.net
- [17] www.nthelp.com/maps.htm
- [18] Gerard O'Reilly, Huseyin Uzunalioglu, Stephen Conrad, Walt Beyeler, "Inter-Infrastructure Simulations across Telecom, Power, and Emergency Services, DRCN 2005.
- [19] S. Conrad, R. LeClaire, G. O'Reilly, H. Uzunalioglu, "Critical National Infrastructure Reliability Modeling and Analysis," Bell Labs Technical Journal, Volume 11, Number 3, pages 57-71, 2006.
- [20] Walt Beyeler, Stephen Conrad, Thomas Corbet, Gerard P. O'Reilly, David D. Picklesimer, "Inter- Infrastructure Modeling - Ports and Telecommunications," Bell Labs Technical Journal, Volume 9, Number 2, 2004, 91-105.
- [21] David J. Houck, Eunyoung Kim, Gerard P. O'Reilly, David D. Picklesimer, Huseyin Uzunalioglu, "A Network Survivability Model For Critical National Infrastructure," Bell Labs Technical Journal, Volume 8, Number 4, 2003.
- [22] U.S.- Canada Power System Outage Task Force, "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations, April, 2004.