

# Routing Analysis and Peering Policy for Enhancing Internet Performance and Security

## Contents

<b>1</b>	<b>Project Description</b>	<b>1</b>
1.1	Introduction: the problem to be solved . . . . .	1
<b>2</b>	<b>Survey of current data, tools and methods</b>	<b>1</b>
2.1	Quality of topology data . . . . .	1
2.2	Quality of routing data . . . . .	2
2.3	Quality of analysis . . . . .	2
2.4	Quality of visualization . . . . .	3
<b>3</b>	<b>Research Topics and Methodology</b>	<b>3</b>
3.1	Research Topic 1: Describing and modeling Internet structure and change . . . . .	3
3.1.1	Data requirements . . . . .	3
3.1.2	Data collection . . . . .	3
3.1.3	Analysis . . . . .	4
3.1.4	Visualization . . . . .	4
3.1.5	Broader impact . . . . .	4
3.2	Research Topic 2: Characterization of Internet routing system expansion, refinement, and churn . . . . .	5
3.2.1	Data requirements . . . . .	5
3.2.2	Data collection . . . . .	5
3.2.3	Analysis . . . . .	5
3.2.4	Visualization . . . . .	7
3.2.5	Broader impact . . . . .	7
3.3	Research Topic 3: Classifying and assessing the effect of incongruity between topology and routing . . . . .	8
3.3.1	Data requirements . . . . .	8
3.3.2	Data collection . . . . .	8
3.3.3	Analysis . . . . .	8
3.3.4	Visualization . . . . .	9
3.3.5	Broader impacts . . . . .	9
3.4	Research Topic 4: Modeling peering policy dynamics . . . . .	9
3.4.1	Data requirements . . . . .	9
3.4.2	Data collection . . . . .	10
3.4.3	Analysis . . . . .	10
3.4.4	Visualization . . . . .	11
3.4.5	Broader impact . . . . .	11
<b>4</b>	<b>Common threads to overall mission</b>	<b>12</b>
4.1	Data collection requirements . . . . .	12
4.2	Analysis and visualization issues . . . . .	13
<b>5</b>	<b>Connection to R&amp;E community and industry</b>	<b>14</b>
5.1	Collaboration . . . . .	14
5.2	Impact on R&E and industrial community . . . . .	14
<b>6</b>	<b>Scientific impact on society</b>	<b>14</b>

<b>7 Results from Prior NSF Support</b>	<b>15</b>
<b>8 Milestones and Timeframes for funding and research completion</b>	<b>15</b>

## Project Summary

*There are strong parallels between the BGP routing space and the condition commonly referred to as the tragedy of the commons. The BGP routing space is simultaneously everyone's problem, as it impacts the stability and viability of the entire Internet, and no one's problem in that no single entity can be considered to manage this common resource. - Geoff Huston, IRTF mailing list, 15 Dec 2000*

Questions about Internet growth and continuing viability of the BGP4 routing system have become regular topics at operational as well as engineering and research meetings. Many discussions result in global credence given to presented opinions no matter the scope of the data analyzed. Inadequacies in current Internet routing analysis tools and methodology handicap efforts to plan future architectures. There are currently no reasonable inter-domain Internet routing models, and current methods to analyze routing dynamics are neither calibrated nor refined. This project is predicated on the hypothesis that scalable solutions for evolution of the Internet require understanding of current behavior and trends. CAIDA plans to develop techniques for correlating massive volumes of routing and topology data and tools to support modeling and prediction of routing behavior. The four proposed research tasks are of pressing interest to Internet research, operations, and public policy; each requires maintenance of instrumentation for data acquisition as well as management, analysis and visualization of resulting data. CAIDA's extensive measurement infrastructure and previous work in analysis of the Internet's core [1, 2, 3] provide groundwork for this research.

The first task is to develop a calculus to describe and model the structure and dynamics of global Internet topology. The second task is to characterize the evolution of the Internet routing system, especially with respect to expansion, refinement, and churn within core BGP tables. CAIDA's preliminary investigations [2] indicate that actual sources and trends in recent Internet growth and dynamics contradict prevailing wisdom of experts in the area. The third task is to evaluate the effect of incongruity between topology and articulated routing policy. Disparity between paths followed by packets and policy-based BGP paths have been observed, but the effects of these incongruities [4, 5] are unknown. CAIDA's topology probe monitors, many with co-located routing tables available, uniquely allow characterization of incongruity sources [6, 7]. The final task focuses on modeling peering policy dynamics to capture the effects of the millions of peering computations that occur simultaneously and continuously in the Internet routing system. Peering affects the fundamental connectivity of the infrastructure, however one can currently observe only a tiny fraction of the effects. Current models of Internet structure are deceptive in simplicity – models that ignore backup paths cannot accurately portray the behavior and resilience of the Internet to damage.

The intellectual merit of this agenda is substantial, combining mathematical approaches from graph theory with extensive knowledge of BGP engineering practices to derive metrics and dimensions of the Internet's core. The vast quantity and utility of BGP data now available, as well as the ability to extend CAIDA's current topology framework, supports rigorous research. CAIDA's history of collaboration with operational network engineers provides the ability to calibrate analysis results against common engineering practice, as well as to offer insight on the effects of such practices on the larger infrastructure.

The broader impacts of this work on society are profound. An empirically-derived and reasonably detailed depiction of macroscopic Internet structure would support operational stability, national security, and development of future Internet architectures. This research will facilitate the abilities to: identify highly connected components that might jeopardize the overall infrastructure if damaged or removed, as well as strategic locations within the Internet for deploying limited resources; establish a process for incremental, hierarchical emergency notification of ISPs in the face of a critical security vulnerability, in order to minimize damage; illustrate infrastructural tradeoffs of routing and address allocation policies; and provide realistic topology input into network modeling research and simulation, rather than relying on the current limited data [8, 9, 10]. Understanding the underlying connectivity of the relentlessly evolving Internet, an increasingly critical national resource, is prerequisite to securing a more robust, stable, and resilient infrastructure.

# 1 Project Description

## 1.1 Introduction: the problem to be solved

Internet routing poses complex data collection and analysis problems. In an attempt to separate facts from Internet ‘myths’ about the routing system, CAIDA hosted an Internet Statistics and Metrics Analysis (ISMA) workshop focused on Internet routing and topology data analysis. This workshop, held in December 2001 at SDSC, brought together researchers from academia and commercial ISPs actively working in this emerging field. A key workshop deliverable was to clarify what research questions are relevant to maintaining the health of the global Internet routing system.

ISP operators have questions that researchers are now in a position to answer. Research efforts offer the potential to benefit the Internet infrastructure as a whole if researchers pay attention to the needs of the operations community<sup>1</sup>. Both communities must cooperate to produce the high quality data necessary to achieve sound results. To this end, the Internet community needs to deploy more real-time instrumentation that will not negatively impact performance. Data and methodologies to support real-time analysis of critical routing events will both improve operations and help researchers to assess trends over time.

Widespread concern exists regarding the appropriate steps to promote graceful evolution of the Internet inter-domain routing system. There is a dearth of empirical data and understanding of how Border Gateway Protocol (BGP) connectivity manifests itself at the router level. For example, there are often disparities between the topology suggested by the Autonomous System (AS) BGP paths active at a given routing table and the actual forward IP path traversed by a packet leaving that router. Since BGP semantics are both applied and transmitted independently at every router, and also non-BGP-based decisions can influence routing, such incongruity is not unexpected. However these disparities undermine the integrity of any Internet modeling methodology that relies solely on topology information at the AS granularity, such as those described in [11, 12, 13, 14].

Understanding of global inter-domain Internet routing dynamics lags significantly behind the growth in the routing system itself [15, 16, 17]. To improve the routing system, we need to better understand how it currently works, starting with calibration and refinement of various methods used to analyze routing dynamics. This proposal is founded on the premise that it takes an understanding of the behavior of the existing Internet to develop scalable solutions that allow graceful Internet evolution. CAIDA proposes the development of techniques for correlating massive volumes of Internet routing and topology data and the development and enhancement of tools aimed at improving the community’s ability to model realistic inter-domain Internet routing behavior.

The structure of the the proposal is as follows: Section 2 describes currently available data, analyses, and tools concerning Internet topology and routing. Sections 3.1, 3.2, 3.3 and 3.4 outline the proposed research topics. For each research topic, five detailed dimensions organize CAIDA’s proposed research methodology: data requirements; data collection; analysis; visualization; and broader impacts. The data collection aspect involves engineering (tool development, deployment, and data management) in addition to research. CAIDA performs engineering tasks as a foundation for rigorous inquiry, and integration of CAIDA’s engineering and research has proven successful in previous efforts [18, 6]. Section 4 discusses data collection, analysis, and visualization issues common to all research topics. Section 5 outlines the process of dissemination of results to the relevant communities. Section 6 describes broader social impacts of the work.

## 2 Survey of current data, tools and methods

### 2.1 Quality of topology data

Both research and commercial entities are involved in WAN modeling and simulation efforts (e.g. VINT [19], OPNET [20], SSFnet [21, 10, 22]). Recent efforts have focused on analysis of real-time updates from BGP

---

<sup>1</sup><http://www.nanog.org>

routing tables [23, 24, 25, 3], and forward path measurement data [26, 7, 27, 28, 29]. There are currently no realistic studies of global multicast topology in the literature.

CAIDA's macroscopic topology project [6], sponsored since 1998 by both DARPA and NSF, provides the most comprehensive Internet public topology analysis in the world. CAIDA has gathered four years of IP-level topology information and indicators of path-specific Internet performance (e.g., latency). CAIDA monitors more than 500,000 destinations [26, 7, 30] from twenty sources worldwide on a daily basis. Data from this effort are currently used by more than twenty research groups [27] around the world. CAIDA maintains daily operational summaries [28], including plots of distributions of IP hop count, global RTTs, RTT by continent, RTT vs longitude, subsets of poorly served (high RTT) destinations, and distribution and dispersion of paths by country and AS. CAIDA also performed a commissioned study for RSSAC of destinations with high RTTs from root nameservers [7] and will continue similar DNS system topology and performance analysis work commissioned by RSSAC in 2002-2003.

## 2.2 Quality of routing data

Interdomain BGP routing tables [31] from backbone providers provide some of the most intriguing sources of Internet data today. CAIDA's December 2001 ISMA workshop report [25] describes public sources of global routing data and its use in modeling and analysis, including applications to traffic engineering. The RouteViews project [17] provides BGP data about the global routing system from several backbone perspectives around the Internet. While the RouteViews project was initially intended to help operators assess operational routing problems, many researchers have found it a rich data source for routing analysis and visualizations [17, 32, 29, 26, 33]. In 2001 RouteViews began to increase the granularity of their data collection, archiving tables every two hours rather than daily and in November 2001 they began to store individual BGP updates as well. Much of CAIDA's topology and routing research relies on the RouteViews data archive.

RIPE, a European IP network management collaboration, also recently started collecting fine-grained routing data by archiving BGP announcements from multiple locations (10 by 2001) in a MySQL database [34]. Specifically, remote route collectors take snapshots of routing tables three times per day and store all BGP route announcements. While more peers would add incrementally to the data coverage, current routing data archives (RouteViews and RIPE) provide sufficient data to perform macroscopic analyses.

Finer-grained temporal analysis of routing dynamics require actual BGP updates (each BGP message) rather than routing table snapshots. In particular, BGP updates are necessary for studying BGP convergence [35] properties at up to a one-second granularity.<sup>2</sup> Similarly, for fine-grained IP topology analysis, BGP data is necessary but not sufficient; it must be augmented with high granularity traceroute data since the nature of the incongruity between actual topology and topology available from BGP is as yet unclear. In particular, CAIDA has determined definitively that macroscopic traceroute data captures considerably more connectivity information than the aggregation of many dozens of BGP tables from top-tiered ISPs [1].

## 2.3 Quality of analysis

Formidable challenges exist in extracting meaningful insights from measurements, even of homogeneous data sets. Data sets are typically extremely large (gigabytes to terabytes), in heterogeneous formats, and gathered at different times, granularities, and locations, both geographically and logically. Developing effective and efficient distillation methods and modes of visualization are significant components of analysis tasks. Section 3 describes how CAIDA will build upon its extensive experience with analytical techniques.

Analysis of Internet router-level topology was pioneered by Paxson in his PhD thesis [36] and follow-up study [37]. Other topology studies have covered smaller data sets [38, 39, 40], data gathered from a single source [41, 42], the impact of topological properties on protocol performance [43], and web-based topology

---

<sup>2</sup>The effects of multi-hop eBGP on the timing of these updates seems unlikely to affect the integrity of resulting convergence analyses since BGP often has to deal with such raw (single hop) latency anyway, at least internationally.

(from URLs and HTTP GETs) [44]. CAIDA's studies have benefited from the availability of multiple data sources [29, 26, 45, 32, 1, 2]. Several studies on Internet connectivity have used publicly available routing tables [17, 46, 47, 48] described above in Section 2.2, in conjunction with other data for validating traffic engineering methodologies [49, 50, 51, 52, 53, 54, 55]. BGP data has been used to map IP addresses to ASes [56], and to analyze routing policies [1]. However, BGP connectivity does not portray redundancy in different parts of the network. BGP tables show only the selected (best) routes rather than all possible paths and also do not show public and private exchange points, short-term AS path variation, or AS load balancing. Significant distortions of network connectivity, including in the core, are produced when using only BGP data to develop any kind of topology map. BGP data is simply too sparse.

## **2.4 Quality of visualization**

Meaningful visualizations of Internet routing and topology are rare. Visualizing large network graphs has received significant attention [41, 57, 58] but no current tools can flexibly incorporate the semantics of BGP routing data. CAIDA has experience with developing and using tools for analysis and visualization of topology and BGP data [32, 59, 60, 41, 7], and has modified some of these tools to handle specific routing semantics or topology characteristics. CAIDA has also developed methods for mapping IP addresses to geographic (latitude/longitude) coordinates [61]. A successful visualization technique for AS relationships is found in CAIDA's AS-based Internet core map [62] which incorporates BGP and skitter data with geographic information.

# **3 Research Topics and Methodology**

## **3.1 Research Topic 1: Describing and modeling Internet structure and change**

Advantages of an empirically-derived depiction of macroscopic Internet structure include the abilities to: identify vulnerable or strategic points in the infrastructure; develop a process for incremental emergency notification of ISPs in the face of a critical security vulnerability; and realistic topology inputs for network modeling research and simulation.

### **3.1.1 Data requirements**

CAIDA has analyzed Internet graphs at various granularities and evaluated problems in gathering and analyzing Internet routing and topology data in graph form. Specifically, CAIDA reported the degree of connectivity information captured through BGP routing tables [17] versus forward topology (traceroute-like [63]) probes [1, 29]. These results show that CAIDA's macroscopic topology data collection captures significantly more bidirectional AS-level connectivity, by at least an order of magnitude, than any publicly available interdomain (BGP) routing tables. This divergence suggests that available routing table information may provide an inaccurate representation of Internet topology. In particular, Internet vulnerability, e.g., resilience to attacks, cannot be reasonably inferred from BGP data [11, 1, 29].

### **3.1.2 Data collection**

CAIDA's current topology data sets [28], in conjunction with the best available BGP routing data [17, 34], are sufficient for initial investigation of this topic. Deployment of additional topology probe monitors will provide more coherent sampling and allow more legitimate extrapolation to the global structure. This deployment will also support other topics of this proposal and research of other organizations.

Router-level topology data is critical for drawing global inferences about Internet core structure. CAIDA proposes to create a tool that focuses on finding new links (both core and edge) with emphasis on maintaining fresh topology state. CAIDA will also create a reliable method to determine router-level connections, (i.e.

interfaces on the same routers) so as to infer which of them belong to the same physical device, and map BGP information to the router level. CAIDA's experience with both collection and analysis of router-level topology data provides a unique advantage in developing this next-generation tool.

### 3.1.3 Analysis

During previous investigations of models of the Internet core, CAIDA has applied graph theory to macroscopic Internet topology data [1]. Using traceroute-based topology data and RouteViews BGP table data, CAIDA has analyzed, at both the IP and AS levels, the most highly connected components of the global Internet [1]. These portions represent strategic locations in the infrastructure. CAIDA has compared several combinatorial approaches, including:

- Extracting the combinatorial core by constructing a directed IP graph from forward path topology data [1] and then iteratively stripping nodes with outdegree 0 (no outgoing links). Nodes minimizing the maximum hop count represent the 'center of the combinatorial core'.
- Extracting the giant connected component: the largest portion of the combinatorial core with bidirectional connectivity between every node.
- Ordering nodes by the lengths of shortest paths originating from them.
- Comparing access points by the size of the 'access cone': the number of nodes/prefixes/ASes or the size of address space that depend on this access point for their global connectivity.

CAIDA's 'top 50' analysis [64] identified the most highly connected units of Internet topology, at three structural levels: IP addresses, IP address prefixes, and AS numbers. Despite the perceived heterogeneity of Internet topology and in the absence of any authority that shapes it, universal invariants apply to connectivity properties of almost every node in the core [1]. A particularly compelling invariant is the reachability function, i.e., how many nodes are within a number of hops from a given node [29].

### 3.1.4 Visualization

Challenges in visualizing pieces of the Internet's 'core' include those for visualization of any large graph; the core contains many thousands of nodes. While there is ongoing research in large graph manipulation [57, 65, 66, 67] none have focused on Internet-specific data. The challenge is to develop techniques to allow us to identify, navigate, and analyze core components. One of the few interactive navigation tools for routing data, created at Merit in 1997 [68], has demonstrated limited utility thus far due to its restricted functionality, efficiency, and discontinuation of support. CAIDA's dispersion graphs [56] depicting both AS and country granularities provide a useful visual mode for tracking topological diffusion into the infrastructure. Topic 1 will require extension of this analysis technique to an interactive mode, and incorporation of hierarchical as well as geographical layout to allow correlation between topology and geography.

### 3.1.5 Broader impact

An empirically-derived and reasonably detailed depiction of macroscopic Internet structure would support operational integrity, national security, and research on viable future architectures of the Internet. Specifically, this research will enable us to:

- Identify Internet vulnerabilities, i.e., highly connected components that might jeopardize the overall infrastructure if damaged or removed.
- Identify strategic locations within the Internet for deploying limited resources, e.g., root name servers.
- Establish a process for incremental, hierarchical emergency notification of ISPs in the face of a critical security vulnerability, in order to minimize the spread of damage.

- Provide realistic topology input into network modeling research and simulation, rather than relying on the fairly limited data available in the form of distilled routing tables [8, 9, 10].

A quantitative analysis of the hierarchical nature of the infrastructure is also likely to reduce ungrounded marketing or political assertions and increase scientific discussions on classification of service providers.<sup>3</sup>

### 3.2 Research Topic 2: Characterization of Internet routing system expansion, refinement, and churn

CAIDA has used the RouteViews data to support our topology analysis [1, 26, 7, 27, 28, 29, 25] as well as to characterize the sources of growth and instability of the routing system [2].

#### 3.2.1 Data requirements

Available BGP routing data from RouteViews and RIPE provide sufficient information for this topic.

#### 3.2.2 Data collection

As of November 2001 RouteViews has tables from 55 peering routers, many more than the 11 it had in November 1997. If an analysis spans a long time interval, some participants have likely joined and/or dropped out during the interval. In this case one must either use fewer tables or use tables from different peers throughout the interval. BGP routing tables available from RouteViews in November 2001 vary in size from 2 prefixes to more than 106K prefixes [47]. Table sizes vary noticeably by day with changes up to a few percent. There is a continuous gain and loss of prefixes, a property we refer to as *churn*.

#### 3.2.3 Analysis

BGP routing is performed at the granularity of IP address prefixes. Each prefix, specified by a starting address and a CIDR length, may not be nested within another prefix. A prefix,  $P_1$ , is *more specific* than another one,  $P_2$ , if  $P_1$  is nested in  $P_2$ . Similarly  $P_2$  is *less specific* than  $P_1$ . We categorize prefixes in the BGP routing tables as:

- standalone prefix* – has no subsets or supersets in the table
- root prefix* – a least specific prefix in the table; it has subsets, but no supersets
- more specific prefix* – a subset of some other prefix
- top prefix* – a standalone or root prefix (i.e. complement set of the more specific prefixes)

Note that all of these definitions are relative and depend on the particular prefix set.

A prefix is *globally routed* if it is common to a comprehensive set of BGP tables. However results depend heavily on the particular BGP tables chosen for analysis. Analytic stability is achieved by abandoning the notion of globally routed prefixes and defining the concept of *semiglobal prefixes* instead. We call a prefix *semiglobal* if it is found in more than half of the RouteViews backbone tables. This choice of prefix set effectively rules out local prefixes seen in only one or two of the tables but also retains a prefix if it is dropped by a few peers. Use of semiglobal prefixes smooth out local variations that appear when using globally routed prefixes on different sets of tables and yield more robust results. (See [2] for details.)

Figure 1 shows the size of global and semiglobal prefix sets for 26 backbone peers in 724 snapshots taken every 2 hours for October-November 2001. We use the 711 of the tables that include more than 100K semiglobal prefixes for detailed analysis of routing system changes.

<sup>3</sup>The NANOG mailing list [69, 70, 71] has a recurring thread about the ‘tiering’ of service providers in the Internet hierarchy. Although there are informal business definitions, there is no quantitative definition for what constitutes inclusion in a given tier.

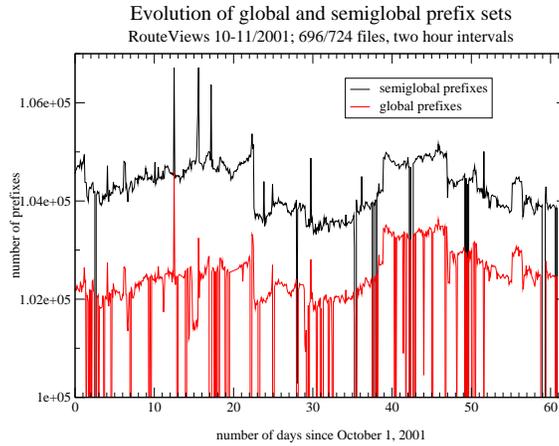


Figure 1: Semiglobal prefixes provide a more robust and stable set of prefixes for routing table analysis.

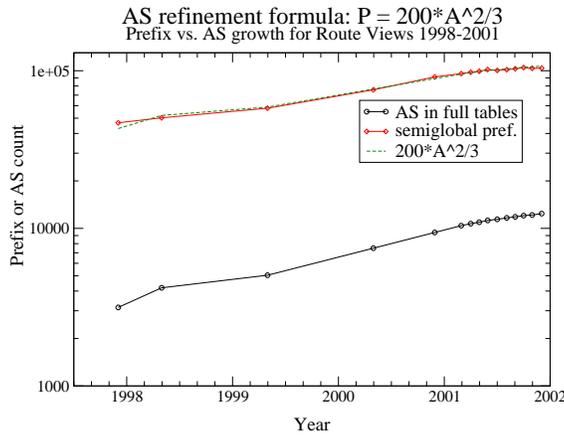


Figure 2: The average number of ASes increases faster (lower line) than prefixes (upper line), showing a refinement process.

**Example: AS refinement (prefix vs. AS growth)**

Internet growth depends upon interaction of economic, social and technological factors. Different metrics appear to grow independently, e.g., number of ASes and prefixes. However CAIDA has found a compellingly simple relation between the count of extant prefixes and ASes:

$$P = 200 A^{2/3} \tag{1}$$

where  $P$  is the number of semiglobal prefixes and  $A$  is the number of ASes.

Although oscillations in the error reflect the dynamic nature of Internet routing, the accuracy of Equation 1 is within a few percent. Figure 2 depicts this approximation against AS and prefix growth, demonstrating that Equation 1 provides an exceptional fit for semiglobal prefixes despite the variability in growth rates. Equation 1 also implies that the average number of prefix announcements per AS is shrinking at a rate of  $200/A^{1/3}$ . We call this process *AS refinement* since the number of prefixes per AS decreases over time. This refinement process accelerated in 2001; it remains to be seen whether this 2/3 relation will continue to hold in the future.

### 3.2.4 Visualization

Visualization challenges are similar to those of the Internet core, but with the additional challenge of handling depiction of macroscopic changes across time. CAIDA is currently investigating various visualization techniques to develop a method for animating such changes.

### 3.2.5 Broader impact

CAIDA's routing system analyses to date we have found several results that contradict prevailing wisdom. In particular, many measures of routing system complexity have demonstrated slow growth, dynamic equilibrium, and occasional contraction over the last several years. Contrary to existing operational opinion, most growth in prefixes between late 2000 and mid-2001 was in top prefixes (i.e., those prefixes that are not more specific) and caused by four sources: allocation (55%), deaggregation(37.5%), expansion(5%) and aggregation(2.6%). Other commonly held assumptions about Internet growth are refuted by our results:

- The number of semiglobal prefixes was stable from October to February 2002, compared to 37% growth between November 2000 and November 2001.
- AS path length, both the mean and the overall distribution, did not significantly change between 1999 and 2001. Link/node ratio (average degree), and peering richness of the BGP AS graph also did not significantly change between November 2000 and May 2001 although individual ASes often exhibited a high degree of change.
- Prefix set churn during 2001 was much higher than prefix growth rate. The churn was highest for more specific prefixes. AS and IP address churn was smaller but still comparable to their net growth. More specific prefixes constitute half of the entries in global BGP tables. Their proportion grew from 50% in November 1999, to 55% in November 2000, and then decreased to 52% by November 2001.
- As of November 2001, multihomed networks (transit or non-transit) are *not* significantly more likely to announce more specific prefixes than non-multihomed networks.
- The number of non-transit multihomed ASes grew from 46% to 49% from 2000 to 2001, but their share of global routes remained stable at around 30%.
- 40% of ASes originate only a single prefix. These ASes contribute 5% of all Internet routes. Only 1% of ASes originate 100 or more routes; these ASes contribute 32% of all routes in the global BGP table. The disparity between contributions by different classes of providers to the BGP table dramatically contradicts prevailing wisdom.
- Half of the routing instability in the form of withdrawal/reannouncement events in late 2001 is contributed by 1.2% of all (12,422 active) ASes. Government networks, telecoms in developing countries and major backbone ISPs are the top contributors. Small ASes (those originating a few prefixes) do not contribute more than their fair share to the BGP table size or to instability of the global routing system.

CAIDA's analyses suggest that many Internet metrics were stable and that Internet growth and instability originate mainly in large and medium-sized ISPs. The data implies that current operational thinking, which currently focuses on controlling the 'explosive' inter-domain routing growth, and views the contribution of smaller (and/or non-transit) ASes as dominant, deserves reexamination. Experts in the operational routing as well as research communities suggest that the recent contraction in routing table growth and churn is the direct result of more careful management of the routing system. Huston attributes global routing system changes of the last two years to developments such as: protocol implementation maturity; widespread deployment of flap dampening; and greater levels of circuit reliability [72]. We find all of these factors, even combined, offer a much less compelling explanation than the current global economic recession. Furthermore it is unclear how long this breathing room will last before we will need a longer-term solution to the interdomain routing problem [73].

### **3.3 Research Topic 3: Classifying and assessing the effect of incongruity between topology and routing**

Two types of incongruity in the Internet routing system remain largely unexplored. The first is the disparity between the topology suggested by the AS paths in a given BGP routing table and the actual forward IP paths taken by a packets leaving that router. BGP semantics are executed and transmitted independently at each hop with each peer. These semantics can be affected by route aggregation or overridden by traffic engineering tools or other non-BGP-based decisions, which makes such incongruity inevitable. Indeed, routing policy effects imply that just because a router can ‘see’ an AS advertising a given prefix P seven AS hops away, this router cannot necessarily reach prefix P along that path. Any intervening AS might have a different reachability perspective. CAIDA’s collection and active analysis of large topology graphs, some with routing tables co-located with the source, uniquely allows taxonomization of classes of incongruity between forward and announced paths [6, 7].

A second type of topological incongruity occurs between unicast and multicast routing. Multicast routing deployment evolved independently from the existing unicast system with fundamentally different routing models. Although policy-based routing has been developed for the multicast infrastructure, it is still possible (and typical) for multicast topologies to fundamentally differ from unicast topology. This problem is increasingly relevant in the face of growth in streaming media and expanded deployment of multicast technologies. Exploratory data analysis will allow for recommendations for measurement, analysis, and mitigation of problems generated by unicast versus multicast routing incongruities.

The significance of the effects of the incongruities that we have observed [4, 5] is currently unclear. Although initial investigations suggest that the consequences for network performance and for the integrity of topology modeling are substantial, quantifying the extent of the impact will be an important result of this research.

#### **3.3.1 Data requirements**

In addition to the topology data described in Section 3.1, incongruity analysis will require routing tables immediately adjacent to the host from which the topology map was gathered. CAIDA has permission from administrators at most CAIDA probe host sites to gather and archive their BGP data and the topology monitoring will continue under independent funding for the next three years.

#### **3.3.2 Data collection**

The most important data collection goal in this topic is to characterize the limitations in the necessarily coarse granularity forward path, BGP, and multicast routing and topology data. Determining the effects (costs) of these limitations will require finer time granularities in measurement instrumentation in order to capture the incongruities of interest. Specifically, quantifying differences between the forward AS path packets follow and the AS path stored in the router will require a real-time BGP feed with triggered active probing of the forward path as soon as an update occurs. This challenge requires extension of CAIDA’s current topology probing [6] software, which currently probes at a much lower granularity.

#### **3.3.3 Analysis**

CAIDA has conducted preliminary studies of forward versus announced AS path incongruity [4]. In several incongruous paths, the actual forwarded path was 2-3 ASes longer than that reported by BGP. These additional ‘hidden’ ASes have negative implications for network performance. CAIDA has confirmed these performance effects in its study comparing the predictive value of both forward AS path lengths and BGP AS path lengths to RTT (round-trip time) performance destinations [5]. The results suggest forward AS path lengths have a much greater predictive value than other metrics.

CAIDA has independently reported that forward path topology data is characterized by connectivity statistics that are significantly different from those derived from BGP data [29]. These results question the validity of research efforts that use only BGP data to model Internet structure or behavior. CAIDA has also observed anomalies of unintended transit, which opens the possibility of loops for multihomed providers [2]. CAIDA will be able to identify incongruent paths that use an AS for transit which was announced only as an origin AS in the routing table, by deriving a list of transit ASes from a routing table and comparing to actual measured forward paths.

### **3.3.4 Visualization**

This topic introduces a relatively unexplored area of both network and visualization research: the unique challenge of how to depict the differences between routing tables, sets of paths or both. CAIDA will initially adapt our existing powerful visualization tools Otter [32] and Walrus [74], along with our recently prototyped PlotPath [75], which can display forward paths between a source and a set of destinations. Further evolution of these tools will occur in conjunction with analysis of the data and concomitant refined understanding of the visualization requirements. Section 4.2 provides additional details about these tools and our goals.

### **3.3.5 Broader impacts**

The ability to compare different routing or forward topology graphs will allow network engineers to compare effects of router configurations and routing policy among various routers over time. These tools could be adapted for use on live BGP data feeds and for simulating the effects on AS topology of removing ASes, e.g., to emulate functional impact of damage to critical portions of the Internet. Analyzing the impact of congruity on performance can also help reveal limits of theoretical results using BGP data, illustrate infrastructural tradeoffs of certain policies, and suggest what other topological abstractions might be useful. Results will also apply to the problem of server selection, relevant to many current and proposed infrastructure management techniques.

## **3.4 Research Topic 4: Modeling peering policy dynamics**

As discussed in Section 3.1, a model of the Internet's core is subject to the effects of interdomain BGP peering. The observed infrastructure is the result of millions of peering computations occurring both simultaneously and continuously. Understanding routing policy and dynamics at a fine resolution is essential to understanding the structure of the Internet. Peering affects the fundamental connectivity of the infrastructure, however one can actually observe only a tiny fraction of the effects through currently feasible data collection methods. Policy observed is limited to the 'best path' computed at a router. Discovery of alternative paths only occurs through changes to the primary path. A realistic model of the Internet must derive topology based on all available paths (optimal and backup). Any other model of the structure may be deceptive in its simplicity. In particular, models ignoring backup paths are poor for understanding the behavior and resilience of the Internet to damage.

### **3.4.1 Data requirements**

Data requirements for understanding peering include the topology data described in Section 3.1 and a complete routing information base from routers. The only 'routing table' available for export from current routers is the forwarding information base ('FIB'), which is insufficient for this task as it includes only the best paths rather than the full set of possible paths.<sup>4</sup>

---

<sup>4</sup>There are extremely good reasons, having to do with system stability, for making it impossible to export any but the best path via BGP, but router vendors have never provided any other means for making the entire database available.

### 3.4.2 Data collection

A complete model of peering will require a mechanism to extract the entire routing information base rather than just the optimal paths used for forwarding. The engineering effort may require resources (from router vendors) beyond the scope of this project, but we will pursue other funding sources and collaborators to meet this goal, since it is central to so much of current Internet research. This objective is discussed more in Section 4.1. CAIDA's existing and planned topology measurement will fill gaps left by this lack of complete routing information base ('RIB') availability.

### 3.4.3 Analysis

Accurate identification of peering points will allow us to better classify peering relationships between ISPs into categories that qualify the underlying business relationships. The model of the Internet core introduced in Section 3.1 will provide an initial framework for parameterizing routing policy relationships among ISPs in the Internet today. CAIDA will create a model with enough flexibility to express routing policies of ISPs. Because CAIDA's topology data is richer in connectivity information than available routing table data, this data will enable researchers to further develop the models of peering initially proposed by Griffin and Wilfong [76] and extended by Gao and Rexford [49]. In both studies, the goal was to infer and characterize Autonomous System relationships in the Internet. Gao *et al.* [52] hierarchically categorized AS relationships and evaluated the accuracy of heuristic mechanisms for determining whether a given AS relationship is provider-customer, peer-peer (in which parties exchange customer routes) or sibling-sibling (in which parties exchange transit routes). The methodology provided a starting point for an investigation of the effects routing policy has on AS path lengths.

Refining this model will require extracting three parameters from active Internet measurements: (1) AS adjacencies, which can be extracted directly from BGP routing tables; (2) peering point prefixes, which can be extracted from traceroute data; and (3) routing policy estimates, which must be inferred from both BGP and traceroute data sources. AS adjacencies in BGP AS paths represent a record of some peering relationships between pairs of ISPs.

An effective study of routing policy cannot treat ASes as single points since ASes may span large geographical distances around the globe. CAIDA's preliminary results found several cases in which ASes do not exhibit the same per-prefix policy at all adjacencies. It is thus necessary to know which peering point the traffic crosses to understand which routing policy applies. Previous work concentrated on topology at the granularity of ASes but did not preserve sufficient per-prefix information to study the effects of routing policy in this topology.

A realistic model of the Internet core will thus require maintaining a more complete list of public prefixes used for peering between ASes in addition to a list of ISP peering sessions observed at BGP peering points. (Both tasks are already supported with some momentum but with scant funding [77, 78].) The current lists can be augmented by using prefixes from traceroute data and assuming that an IP address corresponds to an exchange point if it connects two other distinct ASes. Using traceroute forward path topology in conjunction with BGP data [17] one can map each IP address to an origin AS. Additionally, prefixes that occur on the paths between many diverse pairs of ASes are likely candidates for exchange point prefixes.

### BGP atoms

As routing tables continue to grow, consuming computational and physical resources, it becomes increasingly important to identify the primary sources of the growth in order to limit deleterious effects. In preliminary analyses CAIDA found that many different prefixes have exactly the same AS path. (In August 2000 RouteViews data, there were seven times as many prefixes as unique AS paths in a complete BGP table.) A compelling question is: how many prefixes share a common system of AS paths versus those that are actually routed differently? Using the RouteViews data, CAIDA researchers found it possible to redistribute

IP addresses such that global routing tables would contain less than 20,000 prefixes with no change to the AS path system in April 2001 [79].

In order to make inferences about path properties that allow improvement in routing performance, CAIDA wanted to find the most concise, stateful representation of global routing relationships. After extensive measurement and analysis, CAIDA defined an equivalence between IP prefixes in a BGP table, based on their forward AS paths: Two IP prefixes are AS-path-equivalent if their AS paths coincide in the BGP table, i.e. there is no way to distinguish between them by asking the question "What is the AS path to this prefix from peer (name)?" Such a class of AS-path-equivalent prefixes is called a *BGP atom*.

CAIDA's definition of BGP atom introduces an intermediate granularity between prefix and AS, useful for studying Internet routing and topology. In the actively probed topology data [27], there are twice as many IP address equivalence classes (atoms) as ASes, and about five times as many prefixes as atoms. Because each prefix in a BGP atom shares common characteristics, e.g. packet kinematics (RTT) and dynamics (loss), path properties currently only available via extensive and exhaustive data collection and analysis could be inferred from BGP data. Thus BGP atoms could considerably reduce the range and frequency of Internet measurements while providing input necessary to improve router performance.

When the fundamental unit is the AS path equivalence relationship represented by the BGP atom, a 'routing map' would then illustrate relationships among atoms, CIDR prefixes, and lists of other attributes. Notification of changes in those attributes could occur via TCP or some other (new) protocol to articulate changes in bandwidth, reachability, performance across a given AS path. To conserve traffic the message could be compressed by encoding the AS path once and adding label, prefix, and reachability data elements. The map could be updated with information regarding a given atom rather than for every affected prefix. Using BGP atom granularity for information exchange could conceivably dramatically decrease convergence times without violating policy.

#### 3.4.4 Visualization

Visualization of global peering involves challenges similar to those involved in mapping the core of the Internet (Section 3.1.4). In addition, dynamic components are needed to incorporate real-time determination of global peering relationships using both graphical and text-based interfaces. Specific research goals include: selection of meaningful mappings (both spatial and logical); visualization at the router as well as AS level; identification of the age and source of peering information; and depiction of bilateral versus multi-way peering.

#### 3.4.5 Broader impact

Enhanced understanding of peering dynamics will enable CAIDA to address questions that relate to the core mapping described in Section 3.1. A dynamically acquired/maintained map of Internet peering points will also allow us to address fundamental questions of theoretical, operational and public policy interest:

- For a given AS, which ASes are downstream and upstream? (e.g., how many customers does this AS have, and from whom does this AS receive peering/transit?)
- What percent of address space can be reached only via a given set of ASes or network prefixes? How much connectivity disappears if an AS or a prefix is removed?
- For a given AS number or prefix, which ASes/prefixes are reached only through that AS or prefix?
- What is the frequency of AS path changes for specific destinations over a day/week/month.

Peering points, unlike ASes, are physically localizable. While a transit AS in the Internet core is distributed over a continental or global region, a peering point is typically a single physical location where multiple routers are co-located. IP addresses inside these peering points form a small subset of all IP addresses that can be used to definitively map a path through the Internet to a series of geographical locations.

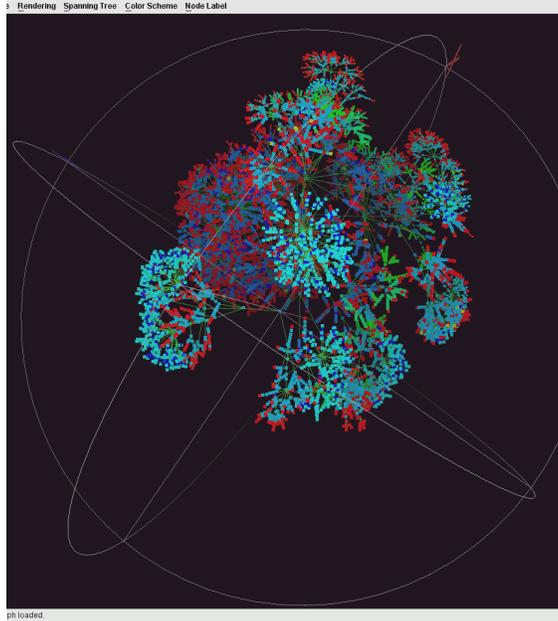


Figure 3: Walrus image, showing nodes colored by AS using data from San Diego monitor (13,000 IP addresses).

The location of each peering point can be the basis for a method of mapping network devices and paths to geographical locations with a high degree of accuracy. This approach will allow CAIDA to quickly report statistics of public policy and regulatory interest, such as what percent of routes transit the U.S. or other countries and in particular what percentage of intra-country routes (where the source and destination are both in the same country) transit another (third-party) country.

## 4 Common threads to overall mission

### 4.1 Data collection requirements

For each topic described in Section 3, CAIDA recognizes the need for a router-level graph of the ‘core’ of the Internet. It is computationally impossible to monitor every Internet router and derive a complete graph, but CAIDA’s topology data collection in conjunction with CAIDA’s heuristic-based interface-to-router aggregation tool provide the best feasible approximation.

CAIDA has years of experience gathering topology data, monitoring from more than 20 source monitors to more than half a million IP destinations around the world. Further, CAIDA has evolved procedures for remote administration of boxes as well as transfer and management of the resulting large data sets. Last but perhaps most important, CAIDA has also achieved an effective method for dealing with complaints from probed sites which minimizes the misunderstandings inherent to topology probing projects [42, 41].

CAIDA collaborates with the University of Oregon’s Advanced Network Technology Center [17] to expand the functionality and operational analysis of RouteViews, based on the needs and ideas from the community and our empirical datasets. Participants in the CAIDA ISMA 2000 workshop [80] offered useful feedback on the utility of currently collected routing data to the research community, which has led ANTC to investigate moving toward a non-IOS solution, improving the archiving infrastructure, and integrating with the RIPE-NCC’s Information Service [34] archiving project.

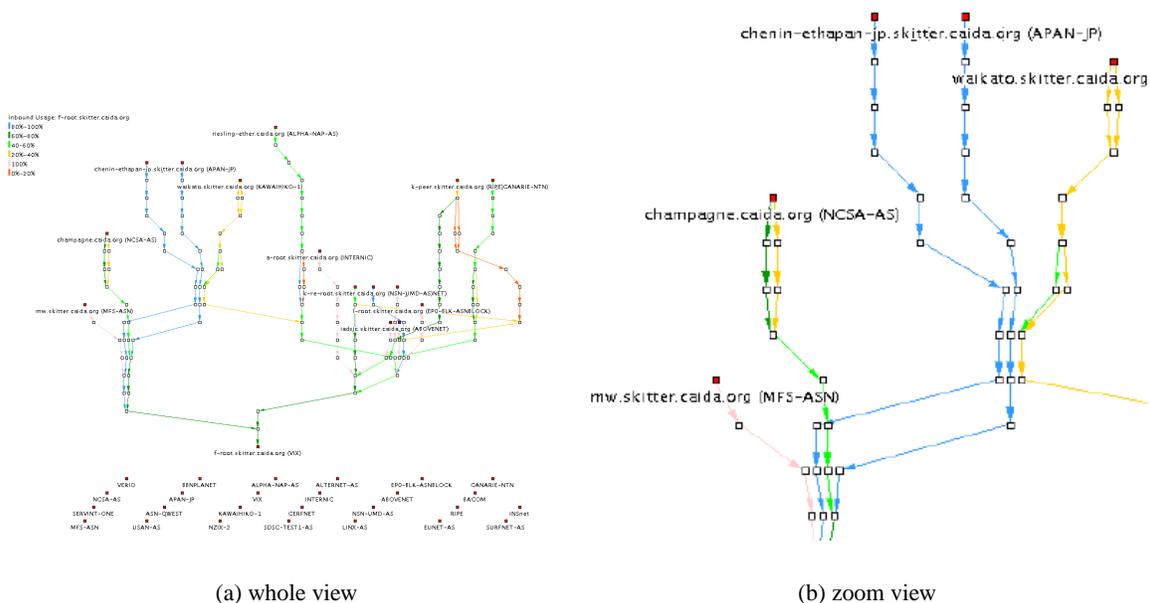


Figure 4: PlotPath layout of router-level topology between 15 global ly deployed monitors and the monitor co-located with the f-root DNS server. Colors show the percentage of measurements which took a given path.

## 4.2 Analysis and visualization issues

The fundamental data types in this proposal are BGP routing tables and router topologies, which can be conceptualized as a complex, semantically rich graph. Researchers need visualization tools to represent this information at a higher level. Existing tools for visualizing these data sets are currently insufficient to allow researchers to filter and segregate interesting components from these graphs.

Any given BGP table holds routing information in the form of AS paths. These paths are sequences of autonomous systems a packet transits from source to destination. Elementary visualization approaches split AS path information into node pairs (i.e. two nodes and one link) to display adjacencies. However these adjacencies are meaningless in isolation from the peering session behind the link. Pure AS connectivity is misleading as it represents unilateral local peering links between two specific ASes. Therefore each AS path must be displayed as an ordered and directed ‘route-ful’ set of AS adjacencies.

CAIDA has enhanced its Otter tool to include a “path” object, which allows a more accurate representation of BGP data. The interactive applet allows highlighting of the entire path that a given peering session believes. A second enhancement modifies Otter to allow pre-computed layout coordinates, scaling the tool from accepting graphs of only a few hundred links to those several orders of magnitude larger. This functionality has been a critical step in macroscopic topology visualization.

CAIDA’s Walrus[74] program was designed to handle large data sets (more than 500,000 nodes). Walrus makes use of 3D hyperbolic space (Figure 3) which allows the user to interactively navigate the data set while the majority of the graph remains within view. Walrus incorporates node and link color to aid in segregation of areas of the data of particular interest. Walrus is currently undergoing alpha testing with 60 users around the world.

PlotPath [75] was developed by CAIDA as a visualization tool that depicts paths between many different source-destination forward and reverse paths simultaneously. Information about individual links is retained. PlotPath incorporates a general layout algorithm allowing incorporation of disparate data sets while retaining a simplified (digestable) view of complex network topology. (See Figure 4.)

CAIDA recognizes that these tools will undergo development in conjunction with evolving mathematical

and statistical analyses. Network topology research will require integration of disparate data sources from multiple perspectives, e.g. integration of forward paths with BGP table updates, as well as longitudinal changes in topology.

## **5 Connection to R&E community and industry**

### **5.1 Collaboration**

CAIDA has a long history of collaboration with vendors, ISPs, routing registries, and other researchers. This project will be no exception. Collaborators will include: UO Advanced Network Technology Center, RIPE NCC, ARIN, Cisco Systems, Inc. Juniper Networks, Inc., and MFN. In addition, we will make use of open-source routing collection software from projects such as Zebra, used in RIPE NCC's RIS Collector project.

### **5.2 Impact on R&E and industrial community**

CAIDA has a strong history of successful outreach and information dissemination through workshops [81], web sites, animations and videos [82], presentations, and presentations to commercial and research organizations. CAIDA is also known for making collected data available to researchers while protecting the privacy of users. Results from this project will lead to routing-related curriculum materials for undergraduate and graduate use as part of CAIDA's NSF-supported Internet Engineering Curriculum (IEC) repository [83] and Internet Teaching Labs (<http://iec.caida.org/>).

Results of this effort will also yield development of tools for navigation, analysis, and correlated visualization of routing tables and massive volumes of traffic data and path-specific performance and routing data critical to advancing both research and operational efforts of the evolving commercial Internet. CAIDA will create a knowledge base that will facilitate suggestions to routing vendors with respect to what instrumentation within the router would facilitate diagnosing and fixing problems in closer to real-time. This research also has implicit relevance to public policy and regulatory questions regarding concentration of administration of Internet infrastructure. Specifically, the developed tools will allow immediate feedback, concerning interrupted service or malicious damage affecting Internet health, to appropriate agencies so as to facilitate rapid implementation of global remedies and/or countermeasures.

## **6 Scientific impact on society**

There is widespread community concern regarding the appropriate next steps to take to support graceful evolution of the Internet inter-domain routing system. A more flexible Internet inter-domain routing system could allow for calculation of optimal route (with variety of metrics), be load-sensitive/adaptive, and even support negotiation of inter-domain SLAs. There is little chance of success in design of a new routing protocol/system without greater understanding of the dynamics of the current system. No good Internet routing models currently exist nor are current methods to analyze routing dynamics calibrated or particularly refined. Both the operational and research communities need to better understand relationships between Internet size and routing tables, topology and policy, and policy and traffic demands.

There is growing concern about the complexity of the Internet (in terms of paths, topology, convergence properties, etc.) due to more complex routing practices (multihoming, backup, load-balancing etc.). Current factors causing BGP table growth and dynamics will not be alleviated by another round of CIDR deployment pressure within the operator community. The community must consider how to manage a BGP routing table which has millions of small entries, rather than the original expectation of tens of thousands of larger entries.

A new routing architecture will require integrated economic feedback such as in-band micropayments for route announcements, as well as inter-domain propagation of performance-based metrics, e.g., RTT, packet loss, jitter, bandwidth estimates. Yet there is little chance of success in designing a new routing

protocol/system without a greater understanding of the dynamics of the existing system. Not only do we currently lack any good inter-domain Internet routing models, but current methods to analyze routing dynamics are neither calibrated nor refined.

The current economic slowdown has given the global Internet routing system a bit of breathing room [3], but we do not know how long this respite will last, nor whether ‘responsible’ use of the routing system can effectively hold things in check. (Consider countries in Asia, with a few billion people, huge markets, and essentially not much dependence upon the US-European economic dynamics).<sup>5</sup> We should use this time wisely to develop models of Internet structure and change that are consistent with available and growing data archives. Payoffs to the research outlined in this proposal include greater understanding of: incongruity between routing policy and topology; peering policy dynamics; relationships between address allocation policy and connectivity; and which metrics best reflect routing dynamics and growth. Progress on understanding these fundamental aspects of the relentlessly evolving Internet is absolutely prerequisite to securing a more robust, stable, and resilient infrastructure.

## 7 Results from Prior NSF Support

1. “CAIDA: Cooperative Association for Internet Data Analysis”. NCR-9711092. \$3,143,580. Oct 1997 - Jul 2001. (Claffy, Moore, Huffaker, Plummer) This collaborative undertaking brings together organizations in the commercial, government, and research sectors. CAIDA provides a neutral framework to support cooperative technical endeavors, and encourages the creation and dissemination of Internet traffic metrics and measurement methodologies [85]. Results of this collaborative research and analytic environment can be seen on published web pages [86]. CAIDA also develops advanced Internet measurement and visualization tools [87].
2. “IEC: Internet Engineering Curriculum Repository”. ANI-97-06181. \$90,555. Aug 1997 - Sep 2001. (Claffy, Moore, Huffaker) This CAIDA project helps educators and others interested in Internet technology to keep up with developments in the field. A repository of collected teaching materials is published on the web [83].
3. “Internet Atlas”. ANI-99-96248 \$304,816. Jan 1999 - Dec 2001. (Claffy) This effort involves developing techniques and tools for mapping the Internet, focusing on Internet topology, performance, workload, and routing data. A gallery that assesses state-of-the-art in this nascent sector is published on the web [88].

## 8 Milestones and Timeframes for funding and research completion

<i>Research Topic</i>	<i>Year 1</i>	<i>Year 2</i>	<i>Year 3</i>
Description of Internet structure	compare models of core; analyze exhaustion of ASNs, IPs, and IP space fragmentation	calibrate and refine topology inputs.	propose emergency notification process for ISPs
Routing system dynamics	track growth, refinement and churn; determine trends. in background update generation vs full-table transmission due to core router stability.	categorize contributors to BGP dynamics. design algorithms to find major BGP update sources and their propagation range	track growth, refinement, and churn. offer data-based policy guidance
Assess incongruity between topology and routing policy	identify examples of incongruity between forward and announced paths	create taxonomy of incongr. types; estimate degree of Internet hidden NATs/VPNs on host and traffic basis	evaluate effects of incongruity on perf. and stability
Model peering dynamics	provide initial framework for parametrizing peering policy	measure statistics of public policy interest	publish model of peering dynamics

<sup>5</sup>It is true that by some metrics the United States indirectly controls the vast majority of the address space and routing system[84], it is important to keep in mind that a single /8 allocated to APNIC is 65,000 /24’s, a significant chunk of a core BGP routing table. And with market pressure mounting, ISPs may soon find incentive to drop the premise of filtering /32’s. Or, we may see entirely different inter-domain paradigm emerge in other parts of the world, e.g., many /32s attached to NATs.

## References

- [1] A. Broido and k claffy, "Internet Topology: Connectivity of IP Graphs," in *SPIE ITCOM: Scalability and Traffic Control in IP Networks*, Aug 2001. <http://www.caida.org/outreach/papers/2001/OSD/>.
- [2] A. Broido, E. Nemeth, and k claffy, "Internet Expansion, Refinement and Churn," in *ETT - European Transactions on Telecommunications*, 2002. <http://www.caida.org/outreach/papers/2002/EGR/>.
- [3] A. Broido, E. Nemeth, and K. Claffy, "Internet Expansion, Refinement, and Churn," Jan 2002. <http://www.nanog.org/mtg-0202/broido.html>.
- [4] A. Broido and k claffy, "Incongruities Between Announced BGP AS Path and Forward Path." <http://www.caida.org/analysis/routing/incongruity/>.
- [5] B. Huffaker, "Distance Metrics in the Internet." in preparation, <http://www.caida.org/outreach/papers/2002/distance/>.
- [6] CAIDA, "Skitter Active Measurement Tool." <http://www.caida.org/tools/measurement/skitter>.
- [7] M. Fomenkov, k claffy, B. Huffaker, and D. Moore, "Macroscopic Analyses of the Infrastructure: Measurement and Visualization of Internet Connectivity and Performance," in *USENIX, San Diego, CA*, Dec 2001. <http://www.caida.org/outreach/papers/2001/Rssac2001a/>.
- [8] A. Medina, I. Matta, and J. Byers, "On the Origin of Power-laws in Internet Topologies." *ACM Computer Communication Review*. April 2000. <http://www.cs.bu.edu/faculty/matta/Research/BRITE/#Papers>.
- [9] Georgia Tech College of Computing, "Modeling Topology of Large Internetworks." <http://www.cc.gatech.edu/projects/gtitm/>.
- [10] B. Premore, "BGP4 Revised Implementation with Internal BGP, Route Reflection, and Policy Specification." SSFnet Project. <http://www.ssfnet.org/homePage.html>.
- [11] R. Albert, H. Jeong, and A.-L. Barabasi, "Error and Attack Tolerance of Complex Networks," *Nature*, vol. 406, pp. 378 – 382, Jul 2000. [http://www.nature.com/cgi-taf/DynaPage.taf?file=/nature/journal/v406/n6794/abs/406378a0\\_fs.html](http://www.nature.com/cgi-taf/DynaPage.taf?file=/nature/journal/v406/n6794/abs/406378a0_fs.html).
- [12] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin, "Resilience of the Internet to Random Breakdowns." *Physical Review Letters*, 85 (21), Nov 2000.
- [13] D. Callaway, M. Newman, S. Strogatz, and D. Watts, "Network Robustness and Fragility: Percolation on Random Graphs." *Physical Review Letters*, 85 (25), Dec 2000.
- [14] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-law Relationships of the Internet Topology." *Proceedings of Sigcomm 1999*. <http://www.acm.org/sigcomm/sigcomm99/papers/session7-2.html>.
- [15] T. Li, "Hardware Implications of Internet Routing Table Growth," Feb 2001. <http://www.nanog.org/mtg-0102/ppt/li/>.
- [16] G. Huston, "Analyzing the Internet's BGP Routing Table," Mar 2001. <http://www.telstra.net/gih/papers/ipj/4-1-bgp.pdf>.
- [17] J. Jaggli and D. Meyer, "RouteViews, U. Oregon's RouteViews Project." <http://www.routeviews.org/>.

- [18] D. Moore and K. Keys, "CoralReef Software Package." <http://www.caida.org/tools/measurement/coralreef/>.
- [19] DARPA, "Virtual InterNetwork Testbed (VINT)." <http://isi.edu/nsnam/vint/>.
- [20] OPNET Technologies, "OPNET Modeler: Modeling Methodology and Model Hierarchy." <http://www.opnet.com/products/modeler/method.html>.
- [21] O. Goldschmidt, "ISP Backbone Traffic Inference Methods to Support Traffic Engineering." <http://www.caida.org/outreach/isma/0012/talks/olivier/>.
- [22] B. Premore, "Investigation of Global Network Routing Behavior." <http://www.caida.org/outreach/isma/0012/talks/premore/>.
- [23] C. Alaettinoglu, "Analysis of RIPE/RIS Project's BGP Data: CIDR at Work," Dec 2001. <http://www.caida.org/outreach/isma/0112/talks/cengiz/index.pdf>.
- [24] A. Ogielski, "Global Routing Instabilities during Code Red II and Nimda Worm Propagation," Dec 2001. <http://www.caida.org/outreach/isma/0112/talks/andyo/index.pdf>.
- [25] k claffy and M. Murray, "ISMA 2000 Workshop Final Report." San Diego Supercomputer Center. 17-19 Dec 2001. <http://www.caida.org/outreach/isma/0112/finalReport.xml>.
- [26] B. Huffaker, M. Fomenkov, D. Moore, E. Nemeth, and k claffy, "Measurements of the Internet Topology in the Asia-Pacific Region." Proceedings of Inet 2000. [http://www.caida.org/outreach/papers/asia\\_paper/](http://www.caida.org/outreach/papers/asia_paper/).
- [27] "Research Community Use of skitter Data sets." CAIDA, 2000. [http://www.caida.org/projects/ngi/content/reports/skitter\\_comuse.xml](http://www.caida.org/projects/ngi/content/reports/skitter_comuse.xml).
- [28] CAIDA, "Skitter Daily Summaries." [http://www.caida.org/cgi-bin/skitter\\_summary/main.pl](http://www.caida.org/cgi-bin/skitter_summary/main.pl).
- [29] A. Broido, "Connectivity of IP Graphs: Comparing BGP and skitter Data." <http://www.caida.org/~broido/bgpskitter.html>.
- [30] D. McRobb, k claffy, and N. Bachman, "skitter: Macroscopic Internet Topology Mapping and Visualization," in *XIWT workshop on Critical Infrastructure: The Path Ahead Symposium on Cross-Industry Activities for Information Infrastructure Robustness*, May 1998.
- [31] J. W. Stewart, *BGP4: Inter-Domain Routing in the Internet*. Addison-Wesley, 1999.
- [32] B. Huffaker, k claffy, and E. Nemeth, "Otter: A General-purpose Network Visualization Tool." INET 1999, Jun 1999, <http://www.caida.org/tools/visualization/otter/paper/>.
- [33] P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the Marginal Utility of Deploying Internet Measurement Infrastructure." <http://www.cs.bu.edu/techreports/2000-018-internet-measurement-utility.ps.Z>.
- [34] RIPE-NCC, "Routing Information Service (RIS) Project." <http://www.ripe.net/ripenc/pub-services/np/ris-index.html>. Henk Uijterwaal, RIPE-NCC Director.
- [35] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet Routing Convergence." Proceedings of ACM SIGCOMM 2000. <http://www.acm.org/sigcomm/sigcomm2000/conf/paper/sigcomm2000-5-2.pdf>.

- [36] V. Paxson, "End-to-End Routing Behavior in the Internet," in *IEEE/ACM Transactions on Networking*, vol. 5, pp. 601–615, Oct 1997. <ftp://ftp.ee.lbl.gov/papers/vp-routing-TON.ps.Z>.
- [37] Y. Zhang, V. Paxson, and S. Shenker, "The Stationarity of Internet Path Properties: Routing, Loss and Throughput," in *ACIRI Technical Report*, May 2000. <http://www.aciri.org/vern/papers/stationarity-May00.ps.gz>.
- [38] J.-J. Pansiot and D. Grad, "On Routes and Multicast Trees in the Internet," *ACM SIGCOMM Computer Communication Review*, vol. 28, Jan 1998.
- [39] R. Siamwalla, R. Sharma, and S. Keshav, "Discovering Internet Topology," 1998.
- [40] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Colling, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan, "Detour: a Case for Informed Internet Routing and Transport," *IEEE Micro*, vol. 19, pp. 50–59, Jan 1999.
- [41] H. Burch and B. Cheswick, "Mapping the Internet," *IEEE Computer*, vol. 32, Apr 1999.
- [42] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet Map Discovery," in *Proceedings of IEEE Infocom*, 2000. Also Technical Report 99-717, USC Computer Science Dept, Oct 1999.
- [43] P. Radolavov, H. Tangmunarunkit, H. Yu, R. Govindan, S. Shenker, and D. Estrin, "On Characterizing Network Topologies and Analyzing their Impact on Protocol Design," 2000. <http://www.isi.edu/~hongsuda/publication/>.
- [44] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener, "Graph Structure in the Web," May 2000. *Comput.Netw.*, 33, 2000.
- [45] B. Huffaker, k claffy, and E. Nemeth, "Tools to Visualize the Internet Multicast Backbone." INET 1999. Jun 1999. <http://www.caida.org/outreach/papers/manta/Paper/>.
- [46] G. Huston, "BGP Monitoring System." <http://telstra.net/ops/bgp/>.
- [47] S. McCreary and B. Woodcock, "PCH RouteViews Archive." <http://www.pch.net/documents/data/routing-tables>.
- [48] T. Bates, "Classless Inter-Domain Routing (CIDR) Report." <http://www.employees.org/~tbates/cidr-report.html>.
- [49] L. Gao and J. Rexford, "Stable Internet Routing without Global Coordination." ACM SIGMETRICS. Jun 2000.
- [50] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving Traffic Demands for Operational IP Networks: Methodology and Experience." <http://www.research.att.com/~jrex/papers/sigcomm00.ps>.
- [51] H. Tangmunarankit, R. Govindan, S. Shenker, and D.Estrin, "The Impact of Routing Policy on Internet Paths," in *Proceedings of INFOCOM*, Apr 2001.
- [52] L. Gao, "On Inferring Autonomous System Relationships in the Internet." <http://www.caida.org/outreach/isma/0012/talks/lixin/>.
- [53] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network Topologies, Power Laws, and Hierarchy," in *Technical Report 01-746, USC Computer Science Dept*, 2001. Submitted for publication.

- [54] K. Nayak, “Measuring Provider Path Diversity from traceroute Data: Work in Progress,” Dec 2001. <http://www.caida.org/outreach/isma/0112/talks/krishna/index.pdf>.
- [55] M. Lloyd, “On the performance characteristics of BGP routing decisions,” Dec 2001. <http://www.caida.org/outreach/isma/0112/talks/mike/>.
- [56] B. Huffaker, M. Fomenkov, D. Moore, and k claffy, “Macroscopic analyses of the infrastructure: measurement and visualization of internet connectivity and performance.” PAM2001 - A workshop on Passive and Active Measurements, (Amsterdam, Netherlands, Apr 2001), RIPE NCC. <http://www.caida.org/outreach/papers/2001/SkitViz/>.
- [57] S. Eick and G.J.Wills, “Navigating Large Networks with Hierarchies,” in *Proceedings of 1993 IEEE Visualization Conf.*, pp. 204–210, IEEE Computer Society, 1993.
- [58] B. P. Madden. Tom Sawyer Software. <http://www.tomsawyer.com/corporate.html>.
- [59] CAIDA, “Geoplot tool.” <http://www.caida.org/tools/visualization/geoplot/>.
- [60] “Mapnet: Macroscopic Internet Infrastructure Visualization and Measurement Tool.” <http://www.caida.org/tools/visualization/mapnet/>.
- [61] D. Moore, R. Periakaruppan, J. Donohoe, and k claffy, “Where in the World is netgeo.caida.org?.” *Proceedings of Inet 2000*, [http://www.caida.org/outreach/papers/2000/inet\\_netgeo/](http://www.caida.org/outreach/papers/2000/inet_netgeo/).
- [62] B. Huffaker, A. Broido, k claffy, M. Fomenkov, K. Keys, E. Lagache, and D.Moore, “skitter AS Internet Graph,” Oct 2000. [http://www.caida.org/analysis/topology/as\\_core\\_network/](http://www.caida.org/analysis/topology/as_core_network/).
- [63] V. Jacobson, “‘traceroute.’” <ftp://ftp.ee.lbl.gov/traceroute.tar.Z>.
- [64] A. Broido, “Finding the ‘Internet Core’: the Top-most IP Addresses, BGP Prefixes and AS Numbers.” <http://www.caida.org/~broido/dest/top50.html>.
- [65] R. A. Becker, S. G. Eick, and A. R. Wilks, “Visualizing Network Data,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 1, pp. 16–28, March 1995.
- [66] T. Munzner, *Interactive Visualization of Large Graphs and Networks*. PhD thesis, Stanford University, Computer Science Department, 2000. [http://graphics.stanford.EDU/papers/munzner\\_thesis/](http://graphics.stanford.EDU/papers/munzner_thesis/).
- [67] T. Munzner, “H3: Laying Out Large Directed Graphs in 3D Hyperbolic Space,” in *1997 IEEE Symposium on Information Visualization*, pp. 2–10, 1997. <http://graphics.stanford.EDU/papers/>.
- [68] Merit Network, Inc., “Internet performance measurement and analysis (ipma) tool: As explorer.” <http://www.merit.edu/ipma/tools/>.
- [69] “NANOG mailing list archives.” <http://www.nanog.org/>.
- [70] A. Broido and k claffy, “Distribution of Peering Sessions Carried by Routers (Core versus Non-Core).” Email to NANOG list, Feb 2001, <http://www.caida.org/~broido/bgp/asoutdeg.html>.

- [71] A. Broido and k claffy, "Response to Question: Of the 'stub AS's', how many inject only few (j10) network prefixes?." Email to NANOG list, Feb 2001, <http://www.caida.org/~broido/bgp/stubas.html>.
- [72] G. Huston, "Commentary on Inter-Domain Routing in the Internet." <http://www.potaroo.net/papers/ietf/draft-iab-bgparch-02.html>.
- [73] S. Hares and Y. Rekhter, "IETF Inter-Domain Routing Working Group." <http://www.ietf.org/html.charters/idr-charter.html>.
- [74] Y. Hyun, "Walrus - Graph Visualization Tool." <http://www.caida.org/tools/visualization/walrus/>.
- [75] J. Gallagher and B. Huffaker, "PlotPath (aka Otter Pup)." [http://www.caida.org/~john/otter\\_pup.html](http://www.caida.org/~john/otter_pup.html).
- [76] T. Griffin and G. Wilfong, "An Analysis of BGP Convergence Properties." Proceedings of the ACM Conference of the Special Interest Group on Data Communication at SIGCOMM 1999. <http://www.acm.org/sigcomm/sigcomm99/papers/session8-1.html>.
- [77] B. Manning, "Exchange Point Information." <http://www.ep.net/>.
- [78] B. Woodcock and B. Arasmith, "Growth of Public Internet Exchanges over Time," Dec 2001. [http://www.caida.org/outreach/isma/0112/talks/bill\\_ix/](http://www.caida.org/outreach/isma/0112/talks/bill_ix/).
- [79] A. Broido and k claffy, "Internet Global Routing: BGP Atoms." Proceedings of Network-related data management (NRDM) workshop, Santa Barbara, May 25, 2001, p.18.
- [80] CAIDA, "ISMA 2000 Workshop Final Report." San Diego Supercomputer Center. 7-8 Dec 2000. <http://www.caida.org/outreach/isma/0012/finalReport.xml>.
- [81] CAIDA, "Internet Statistics and Metrics Analysis Workshops." <http://www.caida.org/outreach/isma/>.
- [82] CAIDA, "CAIDA Animations / Flicks." <http://www.caida.org/outreach/resources/animations/>.
- [83] E. Nemeth, T. Ott, K. Thompson, and k claffy, "The Internet Engineering Curriculum Repository," in *Proceedings of Inet 2000*, 2000. [http://www.caida.org/outreach/papers/2000/inet\\_iec.html](http://www.caida.org/outreach/papers/2000/inet_iec.html).
- [84] B. Huffaker, "Bgp geopolitical analysis of internet administrative resources." <http://www.caida.org/analysis/geopolitical/bgp2country/>.
- [85] N. Brownlee, "Metrics Frequently Asked Questions (FAQ)," Jan 2000. <http://www.caida.org/outreach/metricswg/faq.xml>.
- [86] CAIDA, "CAIDA Analysis Projects." <http://www.caida.org/analysis/>.
- [87] CAIDA, "CAIDA Internet Measurement, Analysis, and Visualization Tools." <http://www.caida.org/tools/>.
- [88] CAIDA, "Internet Atlas Project." <http://www.caida.org/projects/internetatlas/>.