Contents

Executive Summary:

1. Title: Leveraging the Science and Technology of Internet Mapping for Homeland Security

2. Prime Offerer: The Regents of the University of California; University of California, San Diego (UCSD)

The Regents of the University of California; University of California, San Diego on the behalf of San Diego Supercomputer under the program Cooperative Association for Internet Data Analysis (CAIDA) proposes to apply a decade of experience in Internet topology measurement, analysis, modeling, and visualization capabilities to DHS' immediate cybersecurity needs to understand and protect essential U.S. information infrastructure. Our dependence on the Internet in so many dimensions of our lives has rapidly grown much stronger than our comprehension of its underlying structure, performance limits, dynamics, and evolution. Further, the Internet's heritage as a cooperative network for government-funded researchers leaves it with fundamental vulnerabilities that are incongruent with its role as a global communications substrate, and ironically leaves it perpetually challenging to research and analyze, for technical as well as policy and economic reasons. CAIDA is now in a position to integrate the following six strategic measurement and analysis capabilities to improve DHS' situational awareness of Internet topology structure and behavior: a new architecture to support Internet topology measurement; application and assessment of techniques for deriving topologies at both router and service provider granularity from the IP path measurements; provider taxonomy and peering relationship inference; geolocation of IP resources; and interactive visualization of large annotated graphs. The result will be the capability to regularly provide richly annotated topology maps of observable Internet infrastructure, as well as a secure measurement platform capable of performing other types of Internet infrastructure assessments if needed.

1 Proposal

The Regents of the University of California; University of California, San Diego on the behalf of San Diego Supercomputer Center under the program Cooperative Association for Internet Data Analysis (CAIDA) offer the technical proposal that has as its main deliverable, periodic updates for router- and AS-level Internet topologies integrated into the dual-layer router+AS-level topologies, and richly annotated with AS business relationships, geographic, latency, etc., attributes. To achieve this main task, the project will also deliver a new Internet topology data acquisition infrastructure and Internet topology data processing, analysis, annotation, and generations software. We describe the details of our technical approach in Section 3.

2 Performance Goals

The proposed work directly targets the goals and deliverables outlined in the BAA, in particular in TTA5: Internet Tomography/Topology. The resulting technologies and data will increase our ability to understand the structure, dynamics, and vulnerabilities of Internet topology that includes U.S. critical infrastructure.

3 Detailed Technical Approach

Our main deliverable will be periodic updates of richly annotated dual router+AS-level topologies of Internet infrastructure. To achieve this goal we will use a tightly integrated methodology consisting of the following building blocks, described in detail in the following section.

1. **Data acquisition and analysis.**

   (a) scamper data collection (IP-level) on archipelago. Use our traceroute-based *scamper* measurement tool on *archipelago*, our new active measurement platform, to continuously collect raw traceroute data.

   (b) Alias resolution (IP→router-level). Resolve IP interface addresses observed in scamper topology measurements to common routers using state-of-the-art alias resolution techniques. Provide an evaluation of the relative accuracy of available techniques.

(c) AS mapping (IP→AS-level). Map IP-links (pairs of observed adjacent IP interfaces) to AS-links (pairs of BGP-peering ASes) via BGP tables using the CAIDA-developed *ASFinder* tool.

(d) AS+router-level merge. Merge the router-level topologies from step 1b and AS-level topologies from step 1c into integrated dual-layer AS+router-level topologies using our powerful dK-series framework.

(e) Topology characteristics. Release software that analyzes the most important and definitive topology characteristics, including those indicative of vulnerability to various forms of attacks.

(f) IPv6 measurements prototype. Prototype IPv6 topology data collection using *scamper* and *archipelago*

2. **Topology annotations and visualizations.**

(a) AS relationships and taxonomy. Annotate the AS-level topologies obtained at step 1c with AS business relationships and AS types.

(b) Geolocation. Annotate merged router+AS-level topologies with node geolocations, using the best available techniques and our dK-series method.

(c) Latencies (optional). Annotate merged router+AS-level topologies with link latencies, using the best available traceroute data and the dK-series method.

(d) Visualizations. Develop new techniques and adapt existing ones, e.g., Walrus, to visualize our richly annotated topologies.

3. **Analysis.**

(a) Topology comparisons. Using the set of important metrics defined in step 1e, compare the archipelago-, skitter-, and DIMES-extracted topologies and relate the differences to specifics of the associated datasets.

(b) AS-ranking++ Improve and enrich our AS-ranking suite, as-rank.caida.org

(c) Topology generator. Release a topology generator producing annotated AS+router-level topologies for analyses and "what-if" experiments.

(d) Telco hotel data integration. If the telco hotel data mentioned in the BAA is made available, integrate the knowledge into the AS-level map.

We describe these building blocks in detail next, but emphasize that the greatest challenge is automation of and interfaces between these architectural components.

3.1 Archipelago (step 1a)

The design goals of our new measurement architecture are flexibility, extensibility, and fine-grained control over active measurement experiments. We achieve these goals using software design concepts that are quite different from those we used in our previous measurement infrastructure skitter[1]. Our new architecture, Archipelago (or Ark) [2], is framed around the fundamental activity of coordination of measurements across the infrastructure. Coordination allows the many pieces of the infrastructure to work together efficiently toward a common goal. Archipelago provides a coordination facility inspired by Gelernter's tuple-space based Linda coordination language [3]. Archipelago extends Gelernter's model to support a globally distributed infrastructure that hosts heterogeneous measurements by a community of vetted users.

Another distinct feature of our new topology measurement architecture is that by default we will not use predefined probe destination lists, which are prohibitively costly to keep current, and as they fall out of date introduce significant error in the results. Instead, we split all routable IPv4 address space into /24 prefixes and for each probe we select an IP address, uniformly at random, from a random /24 portion. We have recently verified using preliminary topology measurements in June 2006 that this simpler approach delivers significantly more topology coverage than skitter. Ark can also support other destination selection algorithms, including predefined lists of addresses to probe (the skitter approach) which is useful for specific infrastructure assessments.

The Archipelago architecture uses a star topology in which the central coordination server (located on our machine room floor in San Diego) communicates with remote monitors located throughout the global Internet. The remote monitors download lists of probe destination IP addresses from the central server. In addition to destination list distribution, the ark coordination facilities provide scheduling, process control, dependency and order checking, and data preparation and cleanup. The tuple space model explicitly supports concurrency, putting the burden of locking shared space on the system rather than the end user, and providing automatic mutual exclusion to guarantee atomic transactions across multiple monitor processes. The tuple space model also supports decoupling of coordination in time (reader/writer processes may have overlapping lifetimes) and space (reader/writer processes need not know the identity, location, or even existence of each other). This design allows dynamically changing, open-ended sets of partic-

ipants and monitors over the course of each experiment. In addition to destination lists, the tuple space can also handle arbitrary lists, queues, trees, graphs, and other data structures existing independently of the control and measurement processes that can concurrently manipulate these data structures. These properties provide a strong foundation for problem solving using distributed measurements.

While the central coordination server maintains a global tuple space, each remote monitor has a local implementation of the tuple space. While processes from all nodes can access regions of the global tuple space for inter-node communications, only local processes on the monitors can access their local tuple regions. These facilities provide the mechanism for numerous communication models such as one-to-one, many-to-many, and all-to-all.

Ark monitors can be tuned to target an optimal rate of probing, which changes with resource availability. As each monitor completes a single assignment of probes, the monitor transfers the resulting data back to the central server, coordinated via the tuple space coordination facility.

All communication between the central coordinating server and the remote monitors is encrypted. Ark's security model includes requirements for a fine-grained authorization mechanism for reading and writing files, transferring data, access control, initiating communication channels, and installing and controlling measurement software. The model also recognizes the importance of scalability, and supports the ability to delegate authority for subsets of monitoring nodes as well as allow the monitoring sites to set and administer site-specific configurations and security policies.

3.2 Alias resolution (step 1b)

The topology data produced at step 1a is a collection of traced paths, i.e., sequences of IP addresses observed by the probing. To reconstruct the router-level topology from this data we need to group multiple IP addresses belonging to the same router. There exist several *IP alias resolution* techniques, i.e., heuristics to perform such grouping of IP interface addresses into the routers to which they belong. We will evaluate, compare, cross-validate (cf. Sections 3.7 and 3.8), and identify improvements to the available IP alias resolution techniques to assess their accuracy and performance. We will consider the following tools:

- `iffinder`.

The CAIDA's `iffinder` tool [4] implements one of the first IP alias resolutions techniques introduced in the Mercator project [5]. The tool sends UDP probe packets to all or a subset of IP addresses seen in the traces with destination UDP ports set to presumably unused values. If router R receives such a packet from prober P destined to one of R's IP interfaces, X, while R's route back to P goes via some other of R's IP interfaces, Y, then R is supposed to reply to P with an ICMP `Port Unreachable` message with its source address set to Y. Prober P can thus conclude that X and Y belong to the same router.

- `ally`.

  The `ally` tool is a part of the Rocketfuel tool [6]. It uses a combination of techniques, all providing some level of confidence that two IP addresses are configured on the same router. `Ally` resolution techniques include:

  - checking source addresses of responses to probe packets as described above;
  - checking if IP identifier fields in pairs of response packets are approximately consecutive;
  - utilizing the ICMP rate limiting function, i.e., the feature that the router responds only to the first probe of a burst of probe packets.
  - checking for differences in TTL values: significantly different TTL values indicate that a pair of IP addresses are not aliases.

  We note that `ally` is only one possible implementation of the method of identifying routers based on the content of the IP headers and other properties of the packets responding to probes. Other IP header fields not checked by `ally` may also offer identification and validation capabilities. We will evaluate the accuracy and efficiency of different variations and extensions [7] of this general approach.

- `AAR and APAR`.

  More recent are the the Analytical and Probe-based Alias Resolvers (AAR and APAR) described in [8] and [9]. The idea behind both techniques is to recognize the structure of the set of IP addresses observed in traces against common IP address assignment schemes. For example, IP addresses configured on point-to-point interfaces often belong to either /30 or /31 subnets.

Given this observation, we can check for the boundary IP addresses in such /30 and /31 subnets in the original and reverse traces, thus inferring which IP addresses are likely to be configured on the same router. For example, if the direct trace is two IP addresses X, Y, while the reverse trace is Y', X', and both pairs (X, X') and (Y, Y') belong to the same /30 or /31 subnets, then we can conclude that X and Y' are configured on the same router. The authors claim that this approach is more accurate, efficient, and simpler than all other existing techniques.

Other, less efficient IP alias resolution methods are reviewed in [8].

3.3 AS mapping (step 1c)

Mapping traceroute data to AS-level topologies is conceptually the simplest step. Traceroute-like techniques [10] such as implemented in scamper capture the sequence of IP interface hops along the forward path from the source to a given destination by sending either UDP or ICMP probe packets to the destination. Using the core BGP tables provided by RouteViews [11], we map the IP addresses in the gathered IP paths to the AS numbers that advertise the longest IP prefixes that match the corresponding IP addresses. If two consecutive IP hops in a trace resolve to two different ASes, we create a link between these two ASs. The set of these links constitute an AS-level topology graph.

Mapping traceroute-observed IP addresses to AS numbers using BGP routing tables involves potential distortion, e.g., due to AS-sets, private ASes, multi-origin ASes (the same prefixes advertised by multiple ASes [12]), and unresolved links. Both multi-origin ASes and AS-sets create ambiguous mappings between IP addresses and ASes, hence we filter them out. In addition, we filter private ASes as they create false links. Unresolved IP hops in the traceroute data give rise to indirect links, i.e., links that connect two resolved IP hops with one or more unresolved hops in between them. We discard indirect links as well. The total discarded and filtered links usually constitute approximately 5% of all links in the initial traceroute-generated AS graphs.

3.4 Dual AS+router-level topology construction (step 1d)

Although the router topologies at step 1b and the AS topologies at step 1c are obtained from the same raw traceroute data, they are intrinsically distinct. This

distinction stems from the completely different techniques used to create these topologies: heuristics to resolve IP addresses that are assigned to the same router for the former, and mapping IP addresses to AS numbers using RouteViews for the latter. Therefore, we must devise special heuristics to construct topologies that simultaneously and accurately represent the Internet at both the router and AS granularities.

To see that the simplistic approach to integrating a derived router and derived AS graph may lead to extremely inaccurate results, consider the following example. Suppose that we construct a router-level topology after executing IP alias resolution techniques at step 1b. We now wish annotate all routers in this topology with the AS who owns the router. Unfortunately, traceroute data contains no information that would identify the owning AS. All information available to us is represented in two maps:

- map1: IP address $\rightarrow$ router ID (step 1b).
- map2: IP address $\rightarrow$ AS number (step 1c);

These two maps do allow us to determine the set of ASes advertising the set of IP addresses assigned to the same router ID, but no more connectivity information than that. So we create a third map:

- map3: router ID $\rightarrow$ a set of AS numbers announcing IP addresses that are part of this router ID.

However, this map alone does not allow conclusions about which particular AS from the set actually owns the router.

If we proceed to assume that a router interconnects all ASes attached to it, i.e., if we create a full mesh of AS links among all attached ASes for a given router ID, then we will certainly introduce false AS links. For instance, in reality the router could be a customer access router belonging to a large ISP, and some (or all) of its interfaces facing customers could have addresses from the IP address space assigned to the ASes of those customers. These customer ASes may not interconnect to each other at all, so creating a full mesh of AS links for this router would create false links among all of them.

Fortunately, we have developed techniques to overcome these challenges and derive accurate topologies that capture both dimensions. These techniques include:

- use the dK-series extended for dual graph (see Section 3.8 for more details);
- filter out all AS links not present in the AS-level graph;
- filter out AS links violating routing polices imposed by inferred AS relationships;
- assign routers to ASes maximizing the number of valid paths according to inferred AS relationships;
- assign routers to ASes based on AS sizes estimated by AS degree [13] or by the total number of observed IP addresses advertised by this AS;
- assign routers to ASes based on the number of IP addresses attached to the same router and advertised by the same AS;
- combinations of the above.

We will evaluate these and other techniques and use those that work best. The main outcome of this step is router-level topologies with links (and possibly nodes) annotated by AS numbers. We call these graphs *dual AS+router-level topologies*.

3.5 Topology analysis software (step 1e)

Vulnerability of a network to attacks depends drastically on calculable characteristics of the network topology. Consider the following three examples:

- *Spectrum.*

  The smaller the number of alternative paths between a pair of communicating nodes in the network, the easier it is for an attacker to sever all these paths to disrupt the communication. The average path diversity can be estimated as the minimum number of links one needs to cut to decompose the network into isolated pieces of approximately equal size. The *spectrum* of the topology graph representing a network, i.e., a collection of the eigenvalues of its adjacency matrix, provides tight estimates for this min-cut number [14].

- *Betweenness.*

  The number of shortest paths passing through a particular node or link, called *betweenness*, is a measure of its communication importance. The higher the

betweenness of a node or link, the more communication paths will be disrupted and/or re-routed if an attacker interrupts the normal operation of this node or link.

- *Assortativity.*

  All other things equal, networks with more links connecting high-degree nodes, i.e., nodes directly connected to many other nodes, to low-degree nodes are more vulnerable than networks with links interconnecting nodes of similar degrees. The former are called *disassortative*, while the latter are *assortative*. The reason that that the min-cut of disassortative network is smaller is that disassortative networks have fewer links in the network core than assortative networks do. Because links in the core interconnect high-degree network communication hubs, severance of such links decomposes a network into disconnected islands. The smaller the number of such links, the easier it is for an attacker to disrupt communications of the global network.

For other important topology characteristics and for explanations of why they are important, see our previous work [15].

We will release the software (per UC policy) to compute all these topology metrics. This software will be generic in nature, applicable to both router- and AS-level topologies.

3.6 Annotations and AS relationships (step 2a)

Inaccuracies associated with representing Internet topologies as simple undirected unweighted graphs result not only from potential sampling biases in topology measurements [16, 17, 18], but also from neglecting link and node annotations. By annotations we mean various types of links and nodes that abstract their intrinsic structural and functional differences. For the Internet topology at the Autonomous System (AS) level, link annotations represent different business relationship between ASes, e.g., customer-to-provider, peer-to-peer, etc. [19], while node annotations represent different types of ASes, e.g., large or small Internet Service Providers (ISPs), exchange points, universities, customer enterprises, etc. [20]. In router-level Internet topologies, link annotations can be different transmission speeds, latencies, packet loss rates, etc. Simply reproducing the topology of the Internet without any knowledge of the semantics of the links and nodes is insufficient; we must also understand and reproduce annotations.

We propose network annotations as a general framework to describe the functionality of individual links and nodes. Since links and nodes are constituents of a global network, increasing accuracy of description at this microscopic level will also increased accuracy of a variety of important macroscopic graph properties. In the AS topology case, for example, instead of considering only shortest paths, we will be able to study the structure of paths that respect constraints imposed by routing policies and AS business relationships; the same applies to path diversity and other properties important for accurate estimation of network resilience to accidents or intentional attacks.

AS relationships are annotations of links of the Internet AS-level topology. They represent business agreements between pairs of AS neighbors. There are three major types of AS relationships:

1. customer-to-provider (c2p), connecting customer and provider ASes;
2. peer-to-peer (p2p), connecting two peer ASes; and
3. sibling-to-sibling (s2s), connecting two sibling ASes.

This classification stems from the following BGP route export policies, dictated by business agreements between ASes:

- exporting routes to a provider or a peer, an AS advertises its local routes and routes received from its customer ASes only;
- exporting routes to a customer or a sibling, an AS advertises all its routes, i.e., its local routes and routes received from all its AS neighbors.

Even though there are only two distinct export policies, they lead to the three different AS relationship types when combined in an asymmetric (c2p) or symmetric (p2p or s2s) manner. If all ASes strictly adhere to these export policies, then one can easily check [21] that every AS path must be of the following valley-free or valid pattern:

1. uphill zero or more c2p links, followed by
2. top zero or one p2p links, followed by
3. downhill zero or more p2c links,

where by p2c links we mean c2p links in the direction from the provider to the customer.

Routing policies reflect business agreements and economic incentives which generally take precedence over (although can be related to) performance characteristics or other attributes of a link. As a result, suboptimal routing and inflated AS paths often occur – Gao and Wang's 2002 study [22] showed at least 45% of observed AS paths were inflated by at least one AS hop, with some paths inflated by up to 9 AS hops.

In our recent work [23] we further demonstrated that ignoring AS relationships leads to inaccuracies, which make the resulting, non-annotated, topologies look significantly more richly connected than the more realistic annotated topologies. In particular, we showed that if AS relationships are ignored, then:

- shortest paths between nodes are shorter than in reality;
- path diversity is larger than in reality;
- traffic load on nodes and links is lower than in reality.

In other words, ignoring AS relationship annotations results in topologies that, in all respects, appear more robust and less vulnerable that they are in operational reality.

To infer AS relationships, we will utilize our award-winning unique techniques based on state-of-the-art multiobjective optimization [19]. These techniques were found to provide unprecedented levels of accuracy both in our own validation [19] and in an independent study [24]. The main idea behind these inference heuristics is to optimize the trade-off between AS relationship information that can be extracted from AS degrees and maximization of the number of valid paths in the resulting annotated AS topology.

3.7 Geographic locations and latencies (step 2b and 2c

We will evaluate available geolocation inference tools and use those that most accurately annotate router nodes in our dual AS+router topologies with geographic and performance attributes. Geolocation inference techniques map individual IP addresses to their geographic positions. Among other evaluation methods, we will use cross-validation of geographic inferences with IP alias resolution results at step 1b: if a set of IP address resolves to the same router, then all these addresses should also map to (approximately) the same geographic location.

To optionally infer router-link latencies we will rely on traceroute data. Each IP hop in the data obtained at step 1a comes with the recorded round trip time (RTT). Unfortunately, relying solely on the RTT data may be misleading since processing of the ICMP `Time Exceeded` messages happens not in the data but in the control plane of the router. Processing time may in fact be significantly longer than the actual RTT between a router and the monitor. Worse, this processing time varies unpredictably from router to router. We will therefore cross-validate our latency inference results with our geolocation and, consequently, IP alias resolution results.

The outcome of this step is a dual AS+router-level topology with routers annotated with geolocations and, optionally, router-links annotated with latencies.

3.8 dK-series

Throughout the proposed work we will use the powerful dK-series framework that we introduced in [25]. The dK-series formalism provides a calculus for analysis of correlations within a network topology.

In its simplest form [25], dK-series are simply correlations among node degrees. More formally, the dK-series for a given network topology is a collection of dK-distributions defined as correlations of degrees of nodes within subnetworks of size d of the given network. Our most recent work with dK-series [23] extends the formalism to apply to correlations among any kind of node or link annotations, with node degrees being the simplest case of node annotations.

The dK-series formalism is elegantly simple and generic in nature. It provides a rigorous analytical framework and unprecedented power to create, analyze, and compare annotated dual AS+router-level topologies. Specifically, we consider the following three concrete examples of how we will use dK-series at different stages of the project:

- **Topology analysis (step 1e).**

  As discussed in Section 3.5 and in our previous work [15], structural topology characteristics such as spectrum, betweenness, distance distribution, assortativity, clustering, are related to the robustness of the topology, defined in terms of minimal number of links or nodes that must be removed in order to fragment a network. Our dK-series approach introduces a basis that unifies all these metrics—as well as any new metrics—into a single series of metrics. The dK-series method allows for evaluation of the robustness of

a given topology without having to compute all these metrics – good news since computation of some of them can be prohibitively hard. Instead, we can limit computations to appropriate dK-distributions and as soon as they match the dK-distribution of some point-of-reference topology whose robustness properties are already known, we can rely on the dK-randomness theory [15] to conclude that robustness properties of the given topology are approximately the same.

- **AS+router-level quasi-duality (step 1d).**

  The input to step 1d is scamper (traceroute) data processed at steps 1b and 1c to derive router- and AS-level topologies. We thus have information on:

  1. number $n(k)$ of routers of degree k;
  2. number $n(k, k')$ of router-links connecting routers of degrees k and k';
  3. number $N(K)$ of ASes of degree K;
  4. number $N(K, K')$ of AS-links connecting ASes of degrees K and K';
  5. number $N(K, K'; k)$ of AS-links connecting ASes of degrees K and K', and going through a router of degree k;
  6. number $n(k, k'; K)$ of router-links connecting routers of degrees k and k', and addressed from IP address space advertised by an AS of degree K;

  Statistics (1 and 2), and (3 and 4) are simply the $1K$- and $2K$-distributions [15] of the router- and AS- level topologies, correspondingly. The last two statistics are forms of dK-distributions extended for dual graphs as they describe correlations between AS- and router-level degrees. Informally, they glue together the AS- and router-level views of the Internet topology. We use them in combination with other methods mentioned in Section 3.4 to construct the dual AS+router-level topologies.

- **Geolocations and latencies (step 2b and 2c).**

  One type of correlation in topologies annotated with geolocations and latencies is between geographic distance and inferred latencies of links. Drastic anti-correlation between two statistics indicates problems with either geolocation mappings, or latency inferences, or both. We can use this type of correlation as a cross-validation tool to verify our inferences (cf. Section 3.7).

3.9 Visualizations, step 2d

For visualization of our topology data, we will initially experiment with Walrus [26], a CAIDA tool for interactively visualizing large directed graphs in three-dimensional space. Walrus can display graphs containing over a million nodes, but visual clutter, occlusion, and other factors diminish its effectiveness as the number and degree of nodes increases. In practice Walrus is best suited to visualizing moderately sized (up to a few hundred thousand nodes) graphs that are nearly trees. We will modify walrus to meet the needs of this project or investigate other tools for network visualization of large small-world graphs [27]. Van Ham [27], now at IBM, has expressed interest in collaborating with CAIDA on large network visualization.

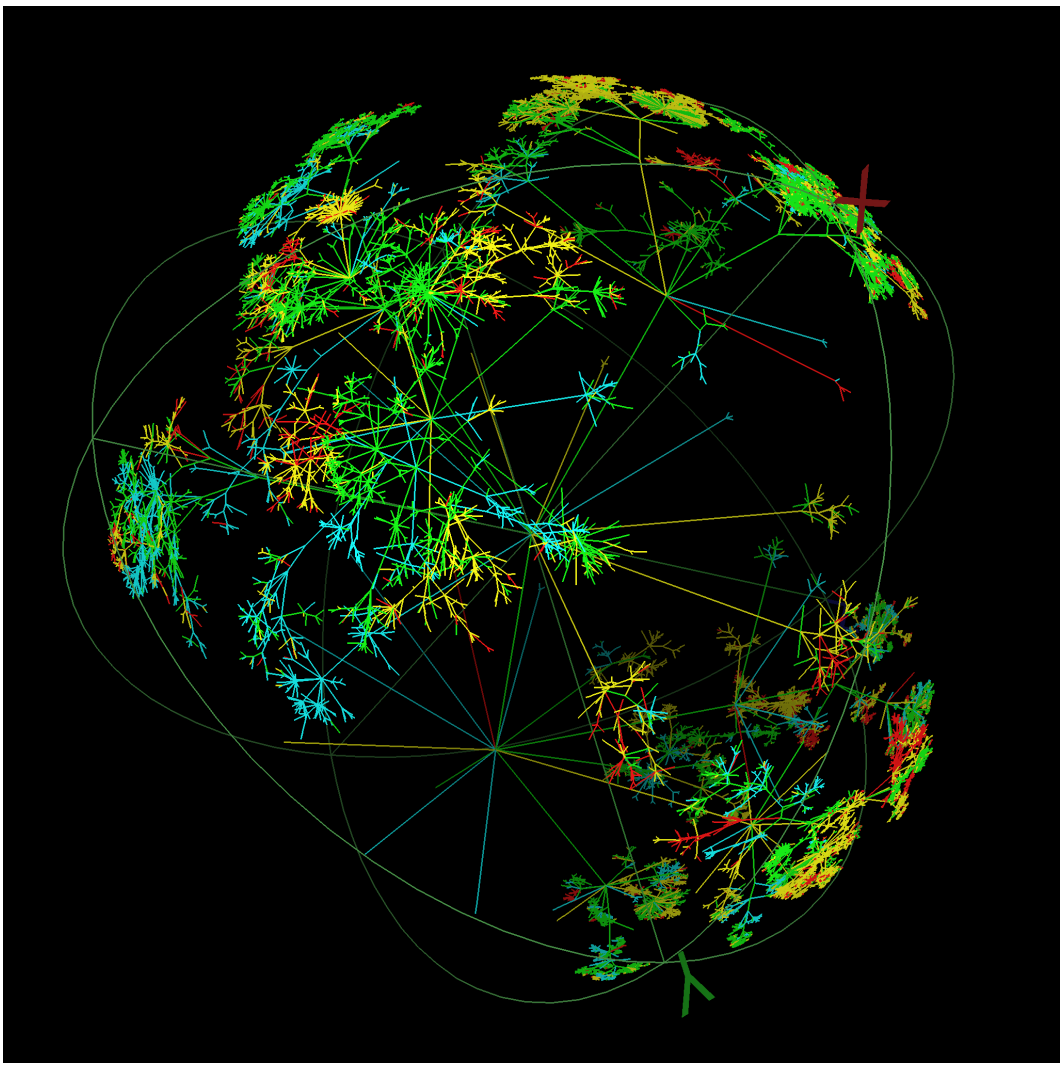

Figure 1: Walrus visualization of median RTTs from skitter measurements.

4 Statement of Work (SOW), Schedule, and Milestones

This Section provides an integrated display for the proposed research, with tasks and major milestones. Corresponding deliverables are listed in the next Section.

4.1 Applied Research Phase - 18 months

Phase I of the project includes the following tasks:

1. Establish ongoing measurements of IPv4 topology using Archipelago measurement infrastructure.
   (a) Complete initial deployment and debugging of Archipelago monitors and software
   (b) Start an ongoing IPv4 topology data collection
   (c) Continue to expand the Archipelago measurement infrastructure
2. Build a router-level graph of the Internet
   (a) Evaluate existing IP-to-router resolution techniques
   (b) Select the best tool and collect data for aliases resolution
   (c) Derive a router-level graph from Ark data and aliases data
3. Build a dual AS-router level graph of the Internet
   (a) Derive an AS-level graph of the Internet from Ark data and BGP data
   (b) Develop methodology of merging the router- and AS-level graphs into a dual topology graph of the Internet
   (c) Produce an experimental dual graph of the Internet topology
   (d) Validate the resulting graph vs. other internationally recognized sources of Internet topology data
   (e) Release software for calculation and comprehensive analysis of topology characteristics

Milestones (time is shown from the start date of the project)

3 mo	ongoing IPv4 topology measurements on Ark platform
6 mo	select best IP alias resolution techniques
12 mo	dual AS-router level Internet graph
12 mo	15 additional monitors added to Ark platform
15 mo	recommendations and possible modifications for topology measurement improvements
18 mo	release comprehensive software suite for topology characteristics

4.2 Development Phase - 12 months

Phase II of the project includes the following tasks:

1. Continue to improve the Archipelago measurements
   (a) Deploy 15 additional monitors
   (b) Prototype IPv6 topology measurements
2. Develop software for automated merging of router- and AS-level graphs into a dual topology
   (a) Develop software for automated construction of router-level topology graphs
   (b) Update software for automated construction of AS-level graphs
   (c) Develop software for building dual AS-router level topology graphs
3. Develop software for annotating dual graphs of the Internet
   (a) Provide automated annotation of AS-graphs with AS types and business relationships
   (b) Compare existing geolocation tools
   (c) Develop software for adding geolocation annotations to dual graphs
   (d) (optional) Develop software for adding latencies annotations to dual graph
4. Develop visualization methods for annotated dual AS-router Internet topology

Milestones (Phase II begins 18 months since the project start)

24 mo 15 additional monitors added to Ark platform
26 mo start regular updates of dual AS-router level graphs
27 mo recommendations on best geolocation tools
30 mo experimental IPv6 topology graph
30 mo annotated dual graph
30 mo visualization of annotated graphs

4.3 Deployment Phase - 6 months

Phase III of the project includes the following tasks:

1. Continue to improve the Archipelago measurements
   (a) Implement recommendations for improving Internet topology measurement learned during Phase I and II of the project
   (b) Deploy 10 additional monitors
2. Advise sponsors regarding use of data to support understanding of critical infrastructure for national security needs
   (a) Enrich our AS-ranking suite using all available measurement data and annotations
   (b) Validate our automated annotated dual graphs vs. other topology sources
   (c) Implement topology generator using annotated dual graphs methodology
   (d) Integrate telco hotel datasets into our data

Milestones (Phase III begins 30 months since the project start)

32 mo 10 additional monitors added to Ark platform
34 mo release topology generator
36 mo release improved AS-ranking

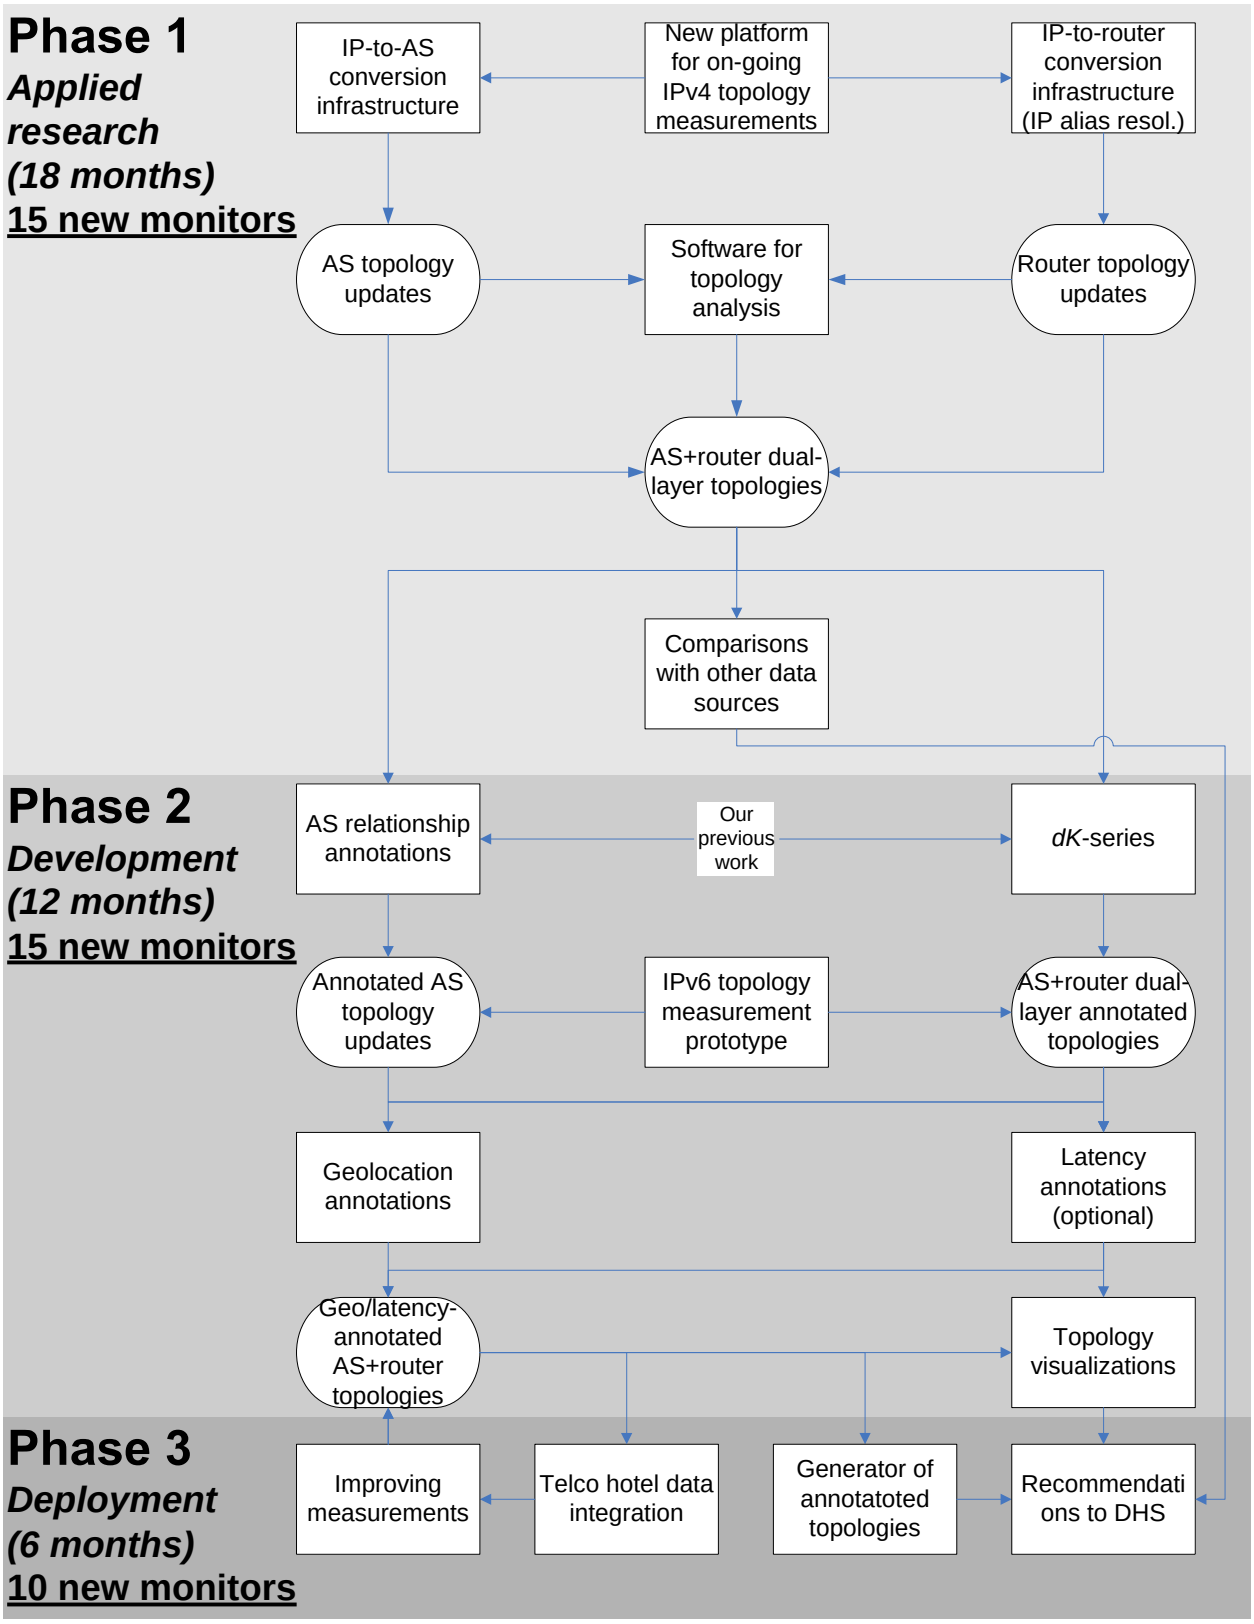

Figure 2: Work plan.

5 Deliverables

This Section provides a brief summary of all deliverables proposed under this effort, along with suggested due dates (in months after the start date of award). This fundamental applied research is being performed on a reasonable efforts basis and the data and reports delivered will be a summary of the data.

5.1 Applied Research Phase

#	Deliverable	Type	Date
1	row IPv4 topology data collected on Ark platform	data	4 mo
2	recommendations for best IP aliases resolution techniques	report	7 mo
3	data for IP-to-router resolution	data	9 mo
4	Ark-based router-level topologies	data	10 mo
5	Ark-based AS-level topologies	data	10 mo
6	Ark-based dual AS-router topologies	data	12 mo
7	caveats and recommendations regarding Ark-based topology measurements	report	15 mo
8	comprehensive software suit for topology characteristics	software	18 mo

5.2 Development Phase

#	Deliverable	Type	Date
1	improved Ark-based topology data	data	20 mo
2	regular updates of router-level graphs	data	22 mo
3	regular updates of annotated AS-level graphs	data	22 mo
4	software for automated merging of router- and AS-level topologies	software	24 mo
5	regular updates of dual Internet topologies	data	26 mo
6	recommendations for best geolocation tools	report	27 mo
7	annotated dual AS-router graphs	data	30 mo
8	preliminary IPv6 topology data	data	30 mo
9	visualization of annotated dual AS-router graphs	report	30 mo

5.3 Deployment Phase

#	Deliverable	Type	Date
1	improved Internet topology data	data	32 mo
2	generator for annotated dual Internet topologies at the AS- and router-level	software	34 mo
3	Ark/skitter/DIMES topology comparisons at different levels of granularity	report	36 mo
4	AS-ranking++	software	36 mo
5	recommendations for the next generation of Internet topology measurement platforms	report	36 mo

6 Management Plan

In this Section we provide a brief summary of the management plan, including an explicit description of what role each participant will play in the project, and their past experience in technical areas related to this proposal.

CAIDA's qualifications for this work are outstanding. CAIDA maintains the longest longitudinal Internet topology measurements (8 years) known to be available. Over these years CAIDA researchers published the most detailed, thorough analyses of the available data and pioneered a number of new approaches to topology data interpretation and utilization.

Two Principal Investigators will lead a team of support personnel to carry out the proposed work. The University of California, San Diego, and the San Diego Supercomputer Center will provide oversight, as well as facilities and support services for this research.

We will hold weekly meetings to ensure coordination among personnel working on the project and to track the status of the scheduled tasks and corresponding deliverables. Minutes will be recorded and emailed to all participants of the project. Since all personnel involved in the proposed effort belongs to the same research group, we can have ad-hoc meetings immediately if specific needs arise.

We will provide a web site for the project and will post regular updates announcing new data availability and other major milestones of the project. Monthly Program Reports will be submitted to our Program Manager as required. The PIs or other senior personnel immediately working on the project will participate in project meetings and reviews.

6.1 Principal Investigators

The PI Dr. K. Claffy is founder and director of the Cooperative Association for Internet Data Analysis (CAIDA), a program at The Regents of the University of California; University of California, San Diego, has led topology measurement, analysis, and visualization projects for over ten years. She will devote 10% (1.2 cal months in each year) effort to provide direction to the research staff working on this project. kc received her Ph. D. in Computer Science from UCSD in 1994.

The Co-PI Dr. Dima Krioukov is an internationally recognized network researcher with first hand operational, engineering, and development experience with Internet routing and engineering. During his tenure at CAIDA he developed new methodologies to capture accurate realistic Internet topology structural

characteristics as well as dynamics. He will spend 10% (1.2 cal months in each year) of his time to provide scientific guidance to other researchers involved in this project. He will devise the mathematical analysis needed for the proposed tasks.

6.2 Other Personnel

Ken Keys, Programmer Analyst, has extensive experience with software development and security-related Internet measurements. He will help evaluate and select tools for accomplishing IP-to-router and alias resolution and implement the software required to build a router-level graph. He will also design and implement a prototype experiment for IPv6 topology measurements using the Ark infrastructure. Ken received his B.S in Computer Science from UCSD in 1993.

Bradley Huffaker, Programmer Analyst, has led most of CAIDA's AS topology analysis and visualization efforts for the last eight years. He brings experience planning, organizing and implementing efforts to develop tools supporting Internet engineering, traffic measurement and analysis, and visualization of data in a variety of programming languages. Brad will develop the methodology for merging the router- and AS-level graphs to build the dual AS-router level graph and implement prototypical visualizations. He will also develop software for automatically annotating the dual graphs. Finally, Brad will implement topology generation software using the resultant methodology for the annotated dual graphs. Brad received his M.S. in Computer Science from UCSD.

Daniel Andersen, System Administrator, has extensive knowledge and experience with the administration and support of large numbers of computer systems in a production environment running a heterogeneous mix of operating systems and acting as desktops, racked servers, remote monitors, data processors, and networked storage devices. He has the skills to support a high availability hardware and software infrastructure including locally and remotely administered systems in support of long-term data collection, processing, storage, and archival. He will be integral to the support as well as further deployment and improvement of the Archipelago infrastructure.

Marina Fomenkov, Project Manager, has a Ph. D. in Information Technologies and extensive background in project management. She will be responsible for management tasks for this project, overall activities coordination and reporting. She will also provide data analysis expertise for the project and will supervise proposed data validation experiments.

Josh Polterock, Manager, Scientific Projects, will coordinate the deployment of additional Ark monitors in each phase of the project and supervise data collection and distribution. His background includes web, application, and database development and technical personnel and project management. He received his M.A. in Communications Management from USC.

The project requires **a Programmer Analyst**, to be named. S/he will require network software design and development experience with programming expertise in Unix environments with various programming language. S/he will evaluate, design and implement software related to the dual router- and AS-level Internet graphs to; 1) do validation and comparison with other sources of Internet topology data, 2) calculate and analyze the topology characteristics, 3) integrate IP address geolocation annotation, and 4) optionally add latencies annotations. This role generally requires an advanced degree in Math, Computer Science, or related field and five or more years of professional experience.

An administrative Specialist, TBN, will assist with project coordination and administration. The budget for this support is a mandatory SDSC policy.

7 Commercialization Plan

CAIDA has experience in licensing software technologies to commercial spinoffs, and executing its own spinoffs. In 2000, CAIDA employees spun off Caimis, Inc., exclusively licensing software CAIDA developed in the five previous years: CoralReef, skitter, and NetGeo, CAIDA's publically available database for geolocation of Internet resources. The company was successfully acquired in 2001. The software licenses were returned to the university by 2002, though we did not have resources to further develop the tools until recently. In the meantime, Digital Envoy, Inc., agreed to donate to CAIDA use of their NetAcuity™ geolocation software for research purposes. CAIDA regularly communicates with Internet-related companies to exchange research and operational expertise.

Fringe benefits have been listed at the current composite rate for each salaried employee. Cost of living and merit increases have been projected in years 2 and 3 based on UCSD system requirements.

7.1 Other costs

Equipment. This project includes no major equipment costs.

Supplies and Materials. We ask for 15 Ark monitors in years 1 and 2, and for 10 monitors in year 3 to continue expansion of Ark measurement infrastructure. We hope to have Ark topology measurement nodes hosted at 100+ sites by the end of Year 3 (through cost sharing with other projects). Other computation resources required for analysis, simulations, visualization, and data processing will come from existing CAIDA resources.

Charges for copying, mail, telephone lines and tolls, and other project specific costs when directly related to this work have also been included. Supplies may also include computer related hardware and software upgrades, and other computer related supplies to be used in conjunction with this research.

Travel. We plan four domestic trips per year for travel. Travel support is requested to attend Project Meetings and reviews (in Washington DC) and to attend relevant domestic research conferences or workshops. Meetings with other Internet researchers will stimulate the exchange of ideas and provide useful feedbacks for our findings and approaches. Candidate venues include ACM SIGCOMM conference, Internet Measurement Conference (IMC), and Passive and Active Measurement (PAM) workshops.

Other Direct Costs. Per UCSD policy, Computer Services and Communications costs have been included for telephone and associated voice and data communications charges directly related to the individuals working on the project.

Indirect Costs. UCSD current indirect cost rate is 54.5costs.

8 Technology Transfer Plan

CAIDA has demonstrated experience in licensing software technologies to commercial spinoffs, executing its own spinoffs, and transferring software technologies via open source licenses. In 2000, CAIDA employees spun off Caimis, Inc., exclusively licensing software CAIDA developed in the five previous years: CoralReef, skitter, and NetGeo, CAIDA's publically available database for geolocation of Internet resources. The company was successfully acquired in 2001. The software licenses were returned to the university by 2002, though we did not have resources to further develop the tools until recently. In 2001, Digital Envoy, Inc., agreed to donate to CAIDA use of their NetAcuity™ geolocation software for research purposes, and we have been using NetAcuity since that time. CAIDA regularly communicates with Internet- related companies to exchange research and operational expertise.

9 Facilities

A general description of SDSC/UCSD facilities and equipment follows.

The resources available through the San Diego Supercomputer Center include supercomputers, archival storage systems, data-handling platforms, high-bandwidth networking, and advanced visualization systems. The capabilities of the center are being upgraded continually to include higher-capability systems that provide a robust environment for cyberinfrastructure research, development and deployment.

Among the hardware resources at SDSC, the foremost is DataStar, an IBM system with peak performance of 15.6 teraflops. DataStar has 2,528 Power4+ processors in nearly 283 nodes connected to a high-speed Federation switch and parallel file system. SDSC also hosts an IBM/Intel cluster associated with the TeraGrid containing 512 compute processors with a peak performance of 3.1 teraflops. Most recently, SDSC put into production the first IBM Blue Gene/L system at an academic institution. This unique architecture boasts 2,048 compute processors plus 128 nodes for the maximum I/O performance possible.

Data-handling resources include a storage-area network (SAN) of 1.4 petabytes (1,400 terabytes) of disk and a 6-petabyte tape-storage archive. Managed by a Sun Fire 15K server with 72 processors and 288 GB of shared memory, SDSC's data-handling environment supports databases, data management, and data mining. Data-intensive computing software includes the Storage Resource Broker, a distributed data-handling system developed at SDSC, digital library technology acquired through collaborations with MIT and Cornell, parallel object-relational database technology acquired in collaboration with IBM, and the High-Performance Storage System (HPSS) archival storage software being developed and tested in conjunction with IBM and LLNL. SDSC also maintains and works with Sun on the SAM-QFS online/archival storage environment. The archival storage systems at SDSC sustain up to 30 terabytes of data movement per day.

SDSC's core program supports scientific data collections for disciplines including seismology, neuroscience, molecular science, Earth systems science, and astronomy. The combination of information management technology, scientific data collections, and the data-handling platforms that support rapid access to the data provides an excellent testbed for evaluating new infrastructure for managing scientific data and scientific algorithms.

The SDSC Synthesis Center supports collaborative viewing of scientific data and advanced scientific visualization capabilities. A complete video and audio

production suite is used to produce publication quality animations. The video lab is network accessible and can be used to render scientific images.

UCSD provides office space and access to telephones, photocopying resources and computer networks. Offices and adjacent meeting spaces have teleconferencing facilities to provide an appropriate venue for interaction with other industrial and academic researchers, if necessary.

10 Government-Furnished Resources

Data and software distribution resulting from this work will make use of the framework and resources established by the DHS Protected Repository for the Defense of Infrastructure Against Cyber Threats (PREDICT) project under DOI contract NBCHC040159 and continued under the new DOI contract NBCHC070133 (currently under negotiations).

11 Assertion of Data Rights

Data and software resulting from this work make use of the framework previously established by the DHS PREDICT project under DOI contract NBCHC040159 and continued under the new DOI contract NBCHC070133 (currently under negotiations).

The offeror asserts for itself, or the persons identified below, that the Government's rights to access, use, modify, reproduce, release, perform, display, or disclose only the following technical data or computer software should be restricted.

The offerer has reviewed the requirements for the delivery of data or software and states:

Data proposed for fulfilling such requirements qualify as limited rights data or restricted computer software and are identified as follows:

1. IPv4 topology data collected on Ark platform
2. IPv6 topology data collected on Ark platform
3. Data for IP-to-router resolution (derived)
4. Ark-based router-level topologies and graphs (derived)
5. Ark-based AS-level topologies and graphs (derived)
6. Ark-based annotated dual AS-router topologies and graphs (derived)

These data come with limited distribution rights as they contain IP addresses that may be used to reveal details about end users including names, geographic and network location, organization, and other personal and private information and should not be subject to unauthorized access. Except for the above, the Offeror (UCSD) can provide the government with unlimited rights for government purposes regarding this proposal.

References

[1] CAIDA, "Macroscopic Topology Measurements." Research Project. http://www.caida.org/analysis/topology/macroscopic/.

[2] CAIDA, "Archipelago: new active measurement architecture and platform." Research & Development Project. http://www.caida.org/publications/presentations/2006/young_wide0611_ark/young_wide0611_ark.pdf.

[3] D. Gelernter, "Generative communication in Linda," *TOPLAS*, vol. 7, no. 1, pp. 80–112, 1985.

[4] CAIDA, "iffinder." http://www.caida.org/tools/measurement/iffinder/.

[5] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet map discovery," in *INFOCOM*, 2000.

[6] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with Rocketfuel," *Transactions on Networking*, vol. 12, no. 1, pp. 2–16, 2004.

[7] R. Sherwood and N. Spring, "Touring the Internet in a TCP Sidecar," in *IMC*, 2006.

[8] M. Gunes and K. Sarac, "Analytical IP alias resolution," in *ICC*, 2006.

[9] M. Gunes and K. Sarac, "Resolving IP aliases in building traceroute-based Internet maps," Technical Report UTDCS-62-06, University of Texas at Dallas, 2006.

[10] "traceroute." http://www.traceroute.org/#source%20code.

[11] "University of Oregon RouteViews Project." http://www.routeviews.org/.

[12] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz, "Towards an accurate AS-level traceroute tool," in *SIGCOMM*, 2003.

[13] H. Tangmunarunkit, J. Doyle, R. Govindan, S. Jamin, W. Willinger, and S. Shenker, "Does AS size determine AS degree?," *Comput Commun Rev*, vol. 31, no. 5, 2001.

[14] F. K. R. Chung, *Spectral Graph Theory*, vol. 92 of *Regional Conference Series in Mathematics*. Providence, RI: American Mathematical Society, 1997.

[15] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, kc claffy, and A. Vahdat, "The Internet AS-level topology: Three data sources and one definitive metric," *Comput Commun Rev*, vol. 36, no. 1, pp. 17–26, 2006.

[16] A. Lakhina, J. Byers, M. Crovella, and P. Xie, "Sampling biases in IP topology measurements," in *INFOCOM*, 2003.

[17] A. Clauset and C. Moore, "Accuracy and scaling phenomena in Internet mapping," *Phys Rev Lett*, vol. 94, p. 018701, 2005.

[18] L. Dall'Asta, I. Alvarez-Hamelin, A. Barrat, A. Vázquez, and A. Vespignani, "Exploring networks with traceroute-like probes: Theory and simulations," *Theor Comput Sci, Complex Networks*, vol. 355, no. 1, pp. 6–24, 2006.

[19] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley, "AS relationships: Inference and validation," *Comput Commun Rev*, vol. 37, no. 1, 2007.

[20] X. Dimitropoulos, D. Krioukov, G. Riley, and kc claffy, "Revealing the Autonomous System taxonomy: The machine learning approach," in *PAM*, 2006.

[21] L. Gao, "On inferring Autonomous System relationships in the Internet," *Transactions on Networking*, vol. 9, no. 6, 2001.

[22] L. Gao and F. Wang, "The extent of AS path inflation by routing policies," in *GLOBECOM*, 2002.

[23] X. Dimitropoulos, D. Krioukov, G. Riley, and A. Vahdat, "Graph annotations in modeling complex network topologies," 2007. `arXiv:0708.3879`.

[24] W. Muehlbauer, S. Uhlig, B. Fu, M. Meulle, and O. Maennel, "In search for an appropriate granularity to model routing policy," in *SIGCOMM*, 2007.

[25] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat, "Systematic topology analysis and generation using degree correlations," in *SIGCOMM*, 2006.

[26] Y. Hyun, "Walrus visualization tool." `http://www.caida.org/tools/visualization/walrus/`.

[27] F. van Ham, "Interactive visualization of large graphs," 2005.