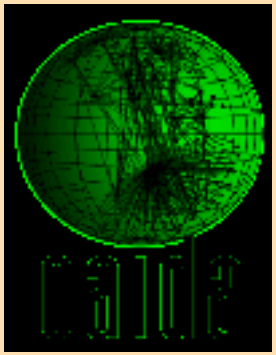

caida 2004-2006 view

ucsd/sdsc/caida
kc@caida.org
june 2004
http://www.caida.org/projects/progplan/

caida activities: 2004 upate

research programs

- active: macroscopic topology project
- passive: (real-time) traffic workload characterization
- DNS analysis
- routing analysis and modeling
- performance/bandwidth estimation methods and tools
- Internet Measurement Data Catalogue (IMDC)
- security issues

other areas

- tools development
- new network visualization metaphors
- policy
- outreach & education

macroscopic topology project

2003 activities

- massive macroscopic traceroute data - most comprehensive in world
  - established legitimate framework for IP topology analysis
- mapping IP -> AS -> organization -> latitude, longitude
  - largest publically available database (still hard problem -needs funding)
- distilled AS topology data available to community
  - derived from skitter probes and BGP data
  - weekly
  - Internet toplogy data kit (ITDK) 2003
    - april 2003 data: topology, routing, meta-data
    - hopefully wide use of this carefully selected data set
- AS ranking (in/out degree)

2004-5

- extending ASrank to organizational granularity
- correlation with routing tables
- IPv6 topology map (scamper data, WIDE funding)
- pop-level map of the Internet (need funding)

macroscopic topology project

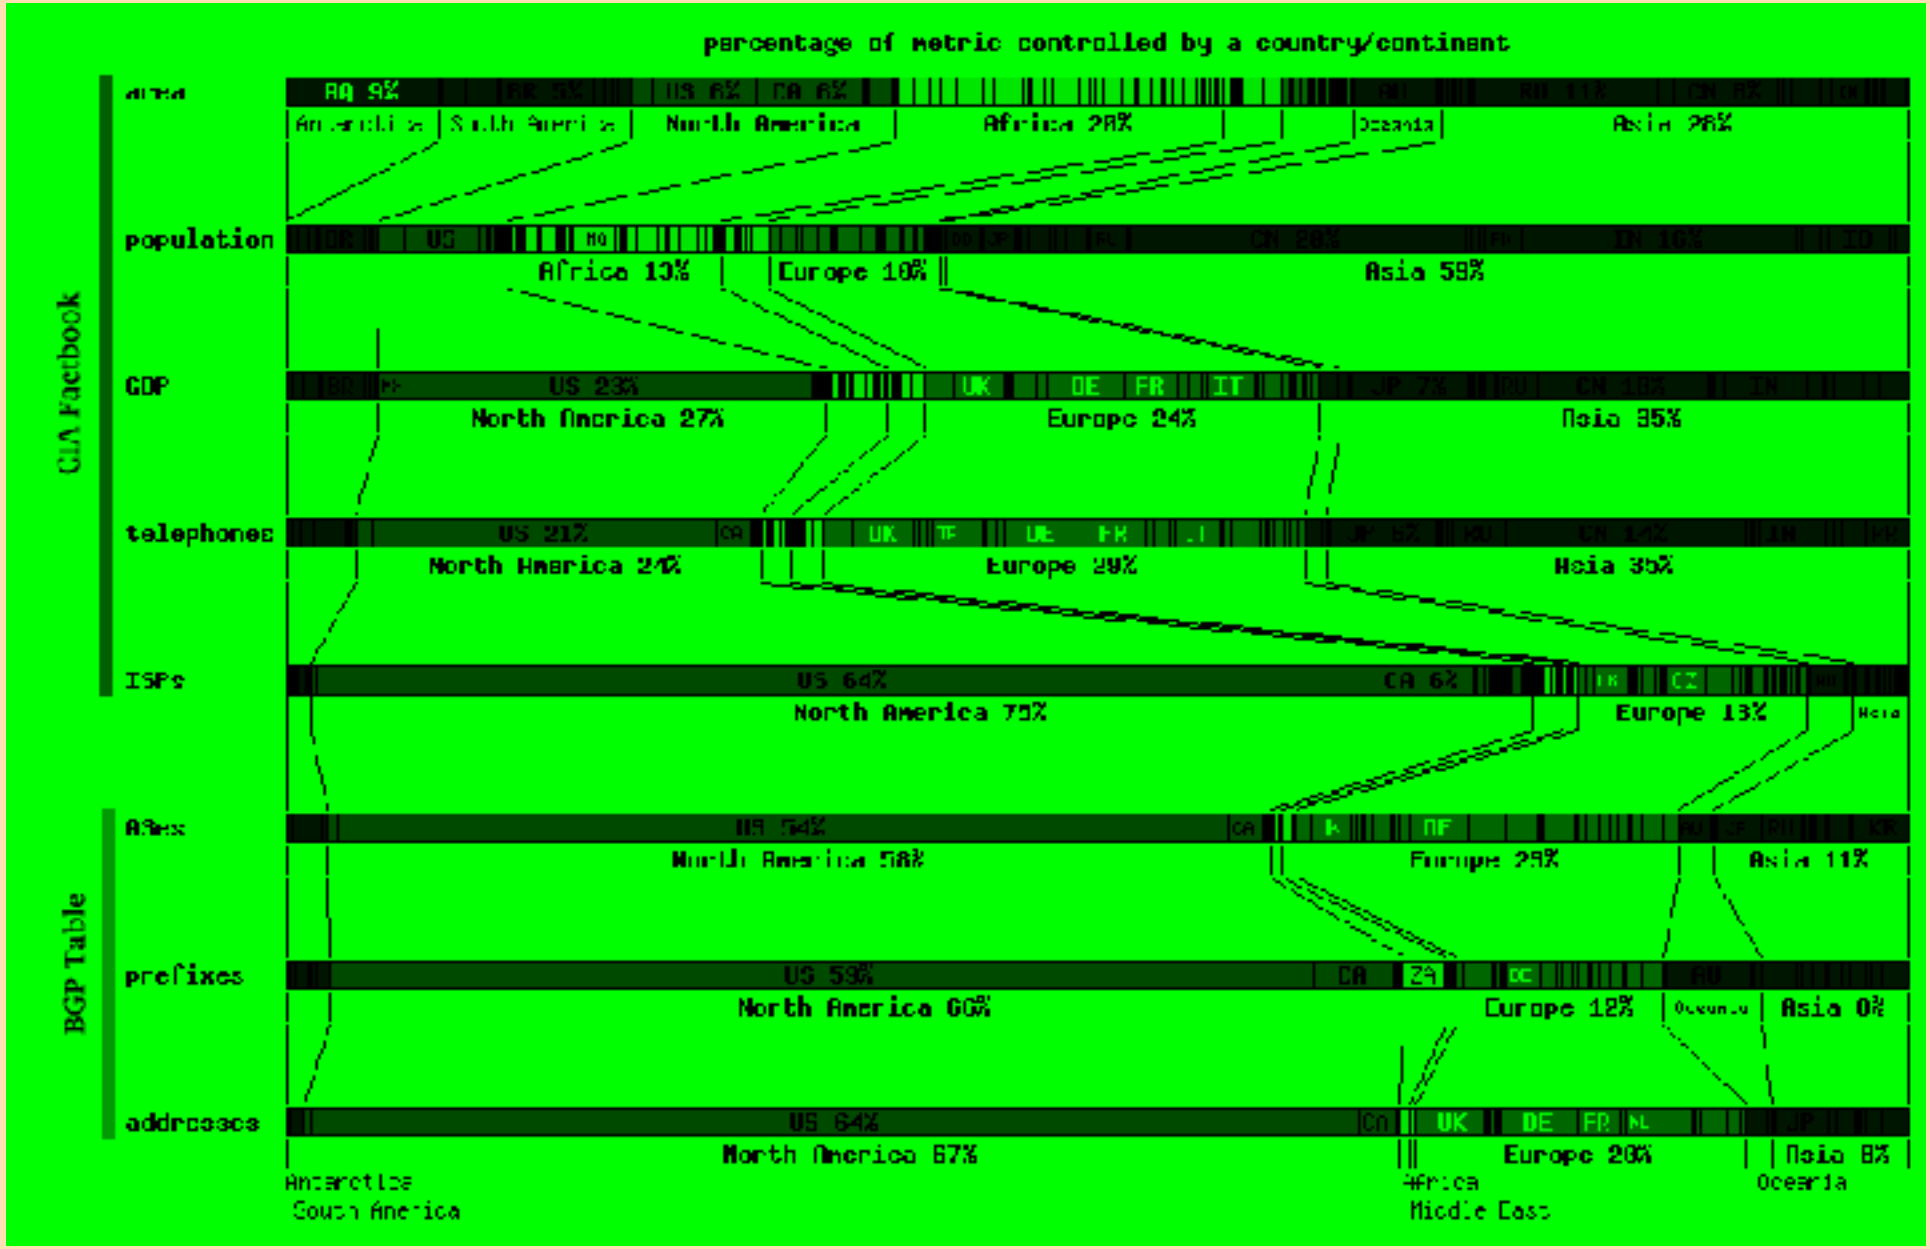

traffic workload characterization

2003-5

- continued passive measurements of Internet data
- techniques for high speed traffic sampling/aggregation
- only OC48 backbone traces available to researchers (so far as we know)
- also only network telescope available to researchers (so far as we know)
  - backscatter, worms, scanning traffic
  - invaluable source of data to security researchers
- various levels of anonymization available to community
  - under AUP
- study how user activities produce torrents of bytes
  - testing models for TCP in presence of bursty cross traffic
  - detection of long-running streams
  - tracking Internet usage patterns, e.g., p2p
  - PAM 2004 paper: 'their share: diversity and disparity in IP traffic'
  - PAM 2004 paper: 'measurements & lab simulations of upper dns hierarchy'

2004-5

- co-chairing IETF WG developing standards for flow measurements
- traffic spectroscopy (andre broido)
- 2005 goal: 24 hour packet trace from the core

Domain Name System (DNS) data analysis

DNS = indispensable Internet component

- new technologies (e.g., anycast, DNSSEC) being deployed at highest (point of failure) levels without instrumentation to debug

2003

- real-time public monitor of root/gTLD performance
- studies of garbage at root servers
- modeling of DNS resolver behavior
  - trace-based simulation

2004-2005

- analysis of F-root (ISC) data for caching resolver pollution
  - submitted paper to Sigcomm workshop (Duane will discuss today)
- support ICANN's Security and Stability Committee (SSAC) with data
  - empirical analysis to support policy recommendations
- proposed CAIDA/OARC project to NSF
  - getting sound DNS data to researchers
  - preliminary OARC support (w WIDE help)

interdomain routing

new routing researcher: Dima Krioukov

- theoretical background in routing
- IRTF chair of working group on scalable interdomain routing
- will talk tomorrow on compact routing
  - infocom 2004 paper
- submitted proposal to NSF for follow-up funding
  - explore applicability of surprising theoretical results from 2003

2003

- completed atoms project. no follow-up for now
- atoms PI patrick verkaik will be joining UCSD PhD program in the fall

2004-5

- supporting data for pop-level map
- compact routing research for inter-domain
- maccroscopic AS topology available weekly

performance tools and analyses

bandwidth estimation

- collaboration with GA tech - they creating new bwest tools
  - pathrate: packet pair technique: dispersion of two back-to-back packets
  - pathload: SLOPS methodology: looks at one-way delays of a periodic packet stream
    - non-intrusive but requires cooperation of both endpoints
- tools methodology, evaluation
- comparing and calibrating available tools
  - pathload, pathrate, pathchirp, ABw, igi, netest2, iperf
- experiments in CalNGI reference lab
  - full control of environment & conditions
  - 100 Mbp and GigE links
- next stage: experiments against real traffic

2004-5 (ga tech lead, pending funding)

- convenient user interface to these tools
- integration with other network middleware

performance data

skitter and scamper delay data

- intermediate RTTs now being collected
- brad and matthew to analyze this year

beluga per hop latency tool

- unfunded

2003

- AS rank
- skitter daily summary

2004

- AS rank by organization
- IPv6 topology map over time
- improve operational integrity of measurement and analysis software

I’net Measurement Data Catalog (IMDC)

‘trends’ project

- year 2 of three-year project funded (partially) by NSF

  "Correlating Heterogeneous Measurement Data to Achieve System-Level Analysis of Internet Traffic Trends"

- design a universal annotation system (meta-data)
  - how to describe heterogeneous Internet data sets?
- build meta-data repository to store "data about data"
- do cross-correlational analysis
- start building ‘community memory’
  - recommendations for long-term archiving of measurement data
- collaboration with IMRG (Internet measurement research group)

It is time for a substantial increase in attention toward the task of conducting Globally Relevant Measurements of Internet phenomena and trends

challenge: characterize Internet traffic trends

motivation: lack of data since 1995
another motivation: way too much data

- admissions about dealing with Internet data
  - vern's 2001 talk www.icir.org/vern/talks/vp-nrdm01.ps.gz
  - david moore's 2002 talk www.caida.org/outreach/presentations/2002/ipam0203/
- longitudinal data are highly ad hoc
- measurement tools lie to us
  - packet filters, clocks, "simple" tools...
  - no culture of calibration
- measurements carry no indication of quality
  - lack of auxiliary information
- measurements are not representative
  - there is no such thing as typical
- analysis results are not reproducible
- large-scale measurements are required
  - that overwhelm our home-brew data management
- we do not know how to measure real traffic

just so i don’t understate the case

- for the most part we really have no idea what’s on the network
- can’t measure topology effectively in either direction. at any layer.
- can’t track propagation of a bgp update across the Internet
- can’t get router to give you its whole RIB, just FIB (best routes)
- can’t get precise one-way delay from two places on the Internet
- can’t get an hour of packets from the core
- can’t get accurate flow counts from the core
- can’t get anything from the core with real addresses in it
- can’t get topology of core
- can’t get accurate bandwidth or capacity info
  - not even along a path much less per link
- SNMP just an albatross (enough to inspire telco envy)
- no ’why’ tool: what’s causing my current problem?
- privacy/legal issues disincent research
- result --> meager shadow of careening ecosystem
- result --> discouraged (or worse) academics

if you’re not scared i’m not explaining this right

obstacles to Internet/network research

where is the data?

- Internet grew organically, incorporating useful technologies as less useful ones obsolesced
- scientifically rigorous monitoring & instrumentation not included in post-NSFNET Internet
- data often proprietary; research use outside owning administrative domain is rare
- researchers can't find out about what little data is available
- Internet research fundamentally different from physics/biology/chemistry -- although we have their problems as well
  - why wouldn't we? -- it's a dynamic, organic system, composed of interactions we don't understand, amon particles we can't access individually
- more like astronomy w/no national virtual observatory or even decent telescopes
- or early quantum mechanics
  - in that you can't measure the particles when you need to
- add a bunch of lawyers -> recipe for bleak future

requires sophisticated tools And special access to data

obstacles to Internet/network research

problems caused by lack of data

- results with predictive power elusive since every link/node has its own idiosyncracies/policies
- makes it hard to assess the quality of any result
- fundamental research cannot be accomplished
- tools designed to combat major problems cannot be tested
  - DoS attack mitigation
  - virus/worm spread
- can't validate theory, model, or simulation against real network
  - not to mention code bugs, methodology flaws

result: weak Internet science

- it's not just soft, it's slippery
- and stunted
- no revolutionary progress in the field for years
- and most of us are partial to revolution
  - so if we're sometimes cranky, that might be why

the view from here

the data we do have

- disparate
- incoherent
- limited in scope
- scattered
- unindexed

what we need

- globally relevant measurements
  - rational architectures for data collection
  - instrumentation suitable for above OC48 links (that number tends to grow..)
  - archiving and disseminating capabilities
  - data mining and visualization tools for use in (nearly) real time?
  - historic data for baseline
  - cross-domain analysis of multiple independent data sets
  - local phenomena vs. global behavior

what can be done

find way to fund researchers to share data

- time and resources are required to share public data with other researchers
- make a data catalog of available data sources -- a single clearinghouse for information on available data sets

need ‘well-curated’ Internet measurement data repository

- measurements need pedigrees describing them, how to navigate
- audit trails, portable analysis scripting language to support reproducibility
- well-managed meta-data (machine readable and searchable)
- software tools to analyze
- understand sampling implications and technology better
- anonymization tools & reduction agents
- long-term and sustained support of such repositories

btw, much here already been/being solved by google, amazon, orkut

- tech transfer might should go both ways

IMDC project: tasks

- deploy strategic Internet measurement instrumentation
- improve measurement tools
  - advanced hardware for monitoring OC48 links
  - advanced software for pre-processing the data various levels of aggregation
  - modules for storage and manipulation of data
  - expand security related monitoring
    - ability to capture DoS attacks in progress
- develop and support a large data storage infrastructure at SDSC
- coordinate movement of traffic measurement data
- create multi-faceted sets of data (datakits)
- universal annotation system (next slide)

IMDC project: universal annotation system

requirements

- accomodate heterogeneous raw data sets
- handle data sets distributed among many sites
- facilitate community access to data repositories
  - data sharing and comparative analysis
- flexible and extensible
  - define meaningful data cross-mappings
- community-based approach to develop common formats
- encourage wide use of common formats
- leave control and security issues to data owners
- ? what else ?

present state of knowledge

- none for the Internet community
- draw from other sciences
  - biology, physics, astronomy

IMDC project: universal annotation system (2)

tasks

- create front-end user interface
  - Internet access to data
  - APIs
  - AUPs
  - compatibility with collection-based sofrtware
- create back end information management system
  - automatic methods of indexing
  - include: data, tools, analysis requests
  - distributed data collection and publication
- maintain and develop compelling tools
  - responsive to user needs
- solicit input from concerned research and standards groups
  - Grid Forum, IETF (IPFIX, IPPM, PSAMP), IRTF (IMRG)
  - NANOG, ISP community (security issues)

expected users of IMDC

- CAIDA currently receives dozens of queries for data every week
- CAAIDA makes available hundreds of gigabytes of data, including:
- anonymized and unanonymized OC48 backbone traces
- network telescope data including:
  - host scan dynamics
  - the spread of Internet worms
  - Denial-of-Service backscatter
- making CAIDA data searchable via IMDC will encourage people to use

we've attempted a compromise between requiring so much context for contributed data that no one will contribute, and requiring so little background that searches don't provide meaningful information

IMDC: research problems (cont.)

example: workload trends

- patterns of usage over time
- pace of new protocols' deployment
- growth of tunneling technologies
  - impact on fragmentation
- more users or more traffic per user?
  - per host, prefix, site, AS
- behavioral characteristics
  - for classification
  - for engineering purposes
- comparison of various flow models
- traffic load and geography
  - local
  - regional
  - international
- tracking distributed denial-of-service activity

expected uses of IMDC

exploding myths

- e.g., RIAA claimed in august "P2P traffic dropped"
  - http://www.pewinternet.org/reports/pdfs/PIP_File_Swapping_Memo_0104.pdf
  - march/may 2003 -> december 2003 brought 29% -> 14% "usage"
  - data sources: telephone surveys nov18->dec14 (huh?); software downloads
  - not data sources: Internet data (wth?)

real data

- have never seen a trace at time t with less p2p traffic than at time t-1
  - frankly i don't see that happening soon

being able to verify/refute this claim is actually a huge deal

- (and not just about changing how we must think of ownership of everything that comes out of our brains)
- will change Internet engineering as we know it today
- current stability and profitability/usability assumptions of asymmetric utilization
  - (btw also driving community to re-evaluate issues of privacy and anonymity;
  - won't ever see a p2p protocol again that doesn't support encryption)

IMDC project: meta-commentary

end game: legitimate tracking of trends

- caveat: trends really not good
- the more we see, the less we like
- kc's 2004 talk 'top problems of the Internet & how researchers can help'
- grep for 'garbage' in bruce sterlings's nsf april 2004 grand challenge workshop keynote talk
  - http://www.cra.org/Activities/grand.challenges/sterling.html
- "digital imprimateur" -- john walker
  - http://www.fourmilab.ch/documents/digital-imprimatur/
  - "how big brother and big media can put the Internet genie back in the bottle"
  - rich 'optimistic pessimism'
- geoff huston's nznog talk
  - video http://s2.r2.co.nz/20040129/
  - slides http://www.nznog.org/ghuston-trashing.pdf
  - not so much with the optimism

this project's website (neutral about falling sky)

- http://www.caida.org/project/trends/

IMDC: interim progress (20/36 months in)

- short answer: not done yet
  - design process complete, including user interface
  - database configured and functional
  - prototype implementation in progress
- medium answer: impediments on our minds
  - ineffective data cataloging
  - disparate formats
  - inadequate documentation
  - inadequate or missing information or quality control
  - inadequate analysis tools
  - inadequate local storage for data analysis
- long answer: workshop in early june 2004
  - co-chair with IRTF's IMRG chair to maximize community input
  - introduce community to and solicit feedback on architecture and user interface
    - get architecture to fit data, not vice-versa
    - discuss typical user modes for researchers, engineers
  - discuss logistical issues
    - supporting processing tools
    - anonymization techniques
    - security of database
  - future workshop 'reverse engineering the Internet' theme (--neil spring's paper)
  - relationship to and support for distributed observatory

CAIDA: security research

global denial of service activity

- CAIDA invented backscatter methodology
  - detecting denial-of-service (DOS) activity on the global Internet
  - monitoring spread of worms in the networks
    - Nimda, Code Red, Sapphire, ... (to be continued)
- the only publicly available data quantifying DOS

main results

- understand nature of current DOS threat
- longer-term analysis of recurring patterns of attacks
  - number, duration, focus, behavior
- modeling quarantine systems to block self-propagating code
  - use real data from epidemics & macroscopic topology probing
  - explore systems in terms of abstract properties
    - speed of detection, granularity of blocking, breadth of deployment

disturbing discovery: no way to react in time!
automated detection of worms and response are essential

network telescope observation station

network telescope

- a chunk of globally routed IP address space
  - e.g., UCSD's has a /8 and /16 network
    - (1/256th plus 1/65539th of all IP version 4 addresses)
- little or no legitimate traffic (or easily filtered legitimate traffic)
- unexpected traffic arriving at the network telescope can imply remote network/security events
- generally good for seeing explosions, not small events
- depends on random component in spread
- has given vital data on: codered*, sapphire, SCO attacks, witty worm

UCSD's network telescope team:

David Moore & Colleen Shannon

security: Internet worm attacks (3)

sapphire effects

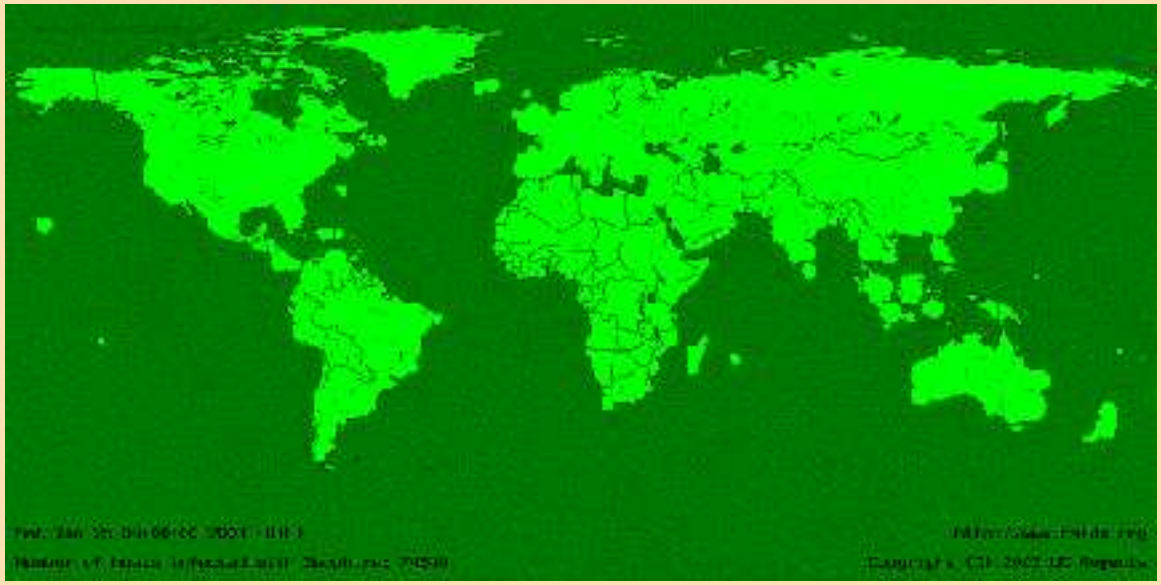

- over 75,000 hosts infected in ten *minutes*
- sent more than 55 million probes per second worldwide
- collateral damage:
  - bank of america ATMS
  - 911 disruptions
  - continental airlines cancelled flights
- unstoppable; relatively benign to hosts

telescope: worm attacks

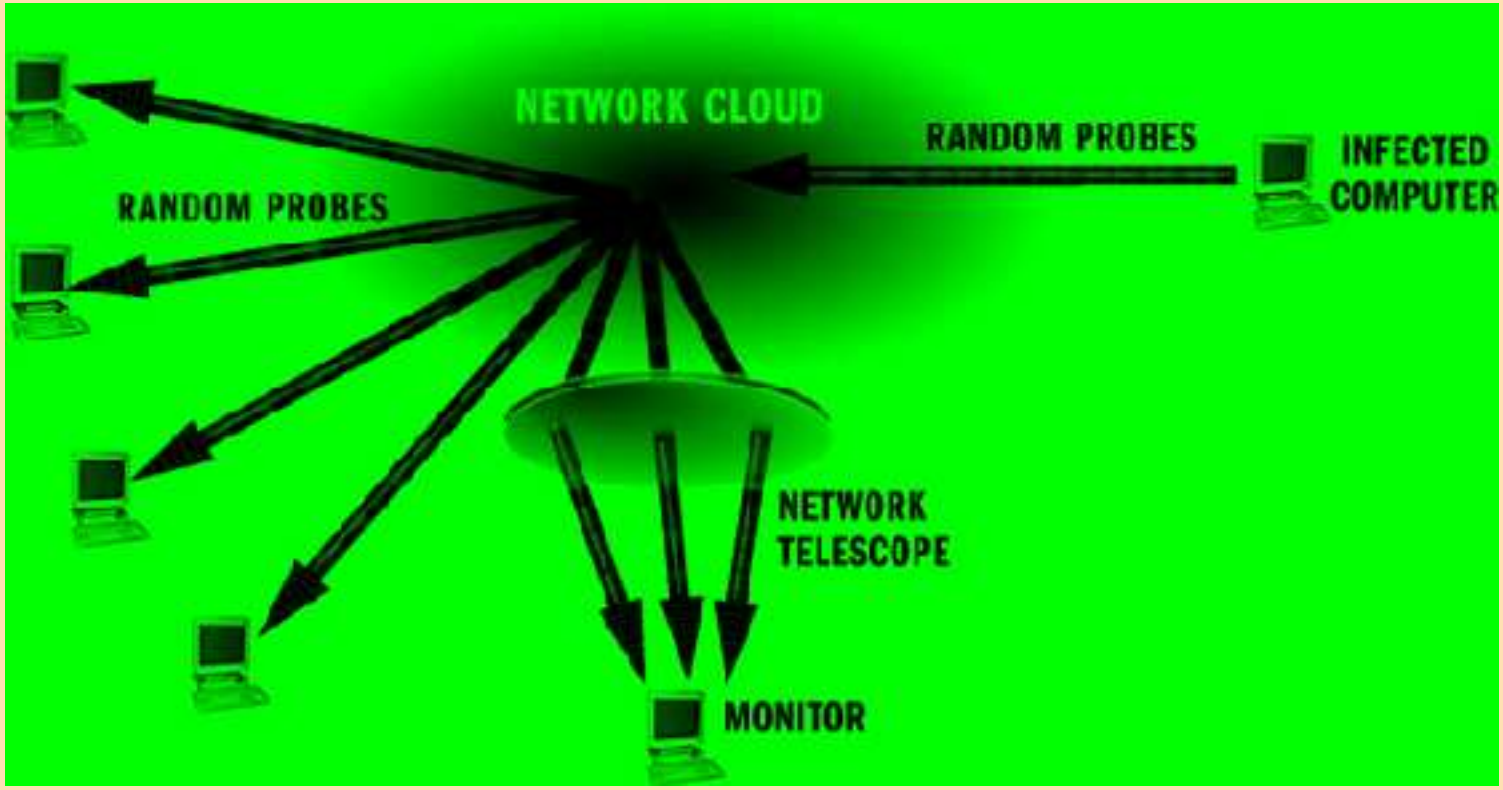

- open research questions
  - random number generation and spread rates
  - effective countermeasures
  - victim classification/hitlists

telescope observation station goals

- continuous data collection with rotating data files:
  - full packet trace kept for 24 hours
  - complete packet header trace kept for 1 week
  - aggregated data (flow tables) stored indefinitely
- sanitized data publicly available to research community
  - under NDA
  - intend to integrate with doug's data collection efforts
- expansion to include monitoring distributed address space
  - countermeasures include to #define telescope prefixes out of scripts
  - countercountermeasures include distributed lenses and moving lenses (requires ARIN support)

telescope: user interface

NTOS graphical interface

- publicly accessible realtime graphical monitor
  - denial-of-sercice attacks
  - worm activity
  - port scanning
- authorized users
  - drilldown technology
    - timescale
    - transport protocol
    - application ports
    - subnets
- ability to save (manually or automatically) data of interest
- email alerts for trigger events

NTOS graphical interface

ICMP host scanning

- 5 october 2003
- some attacks are apparent, but others are difficult to identify

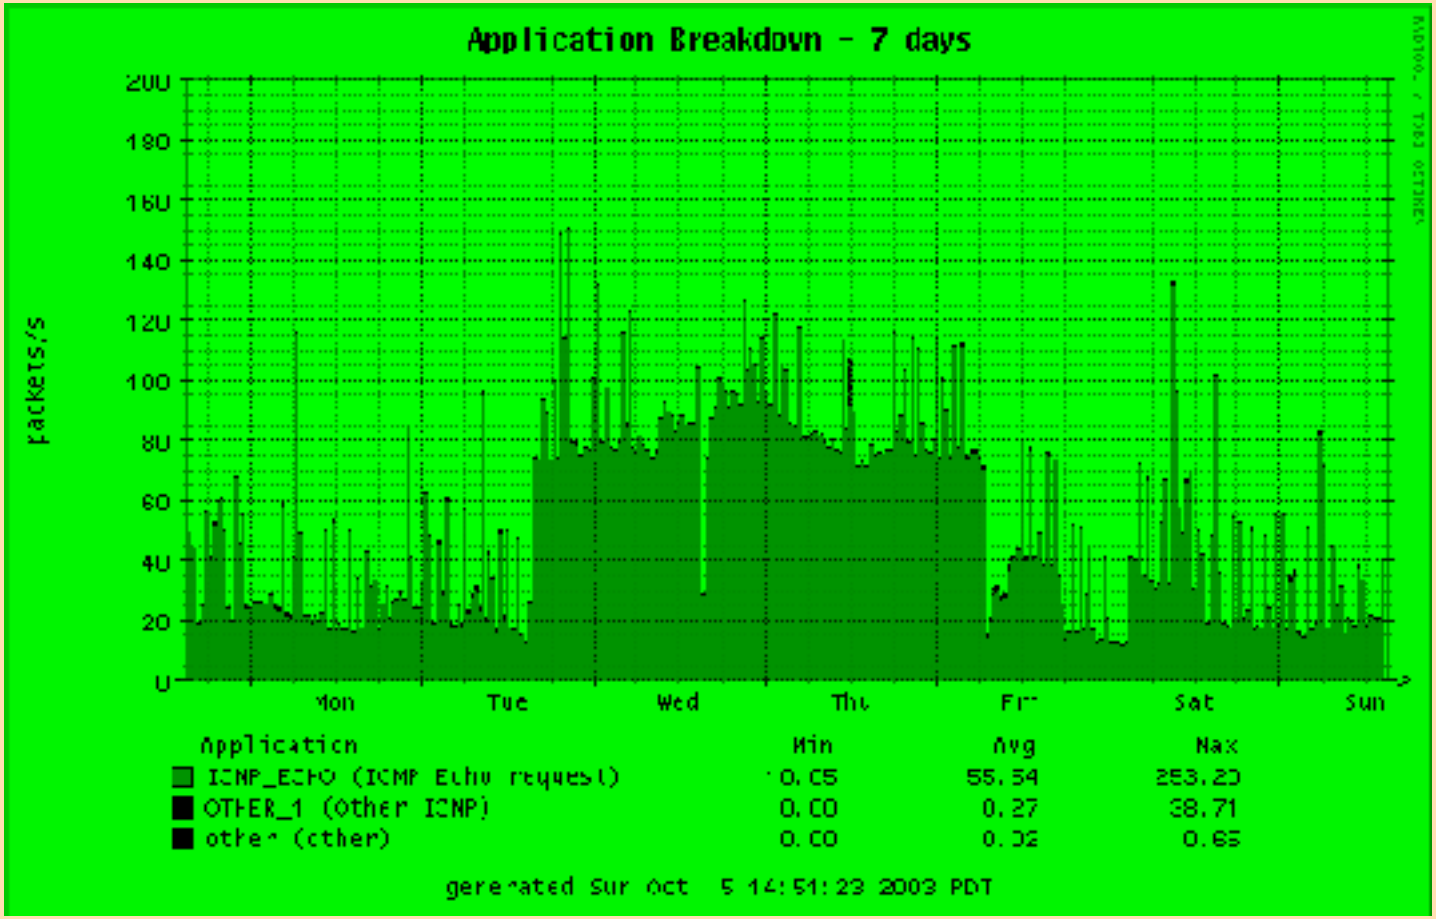

NTOS graphical interface

ICMP host scanning

- 5 october 2003
- viewing attacks by source country helps to differentiate them

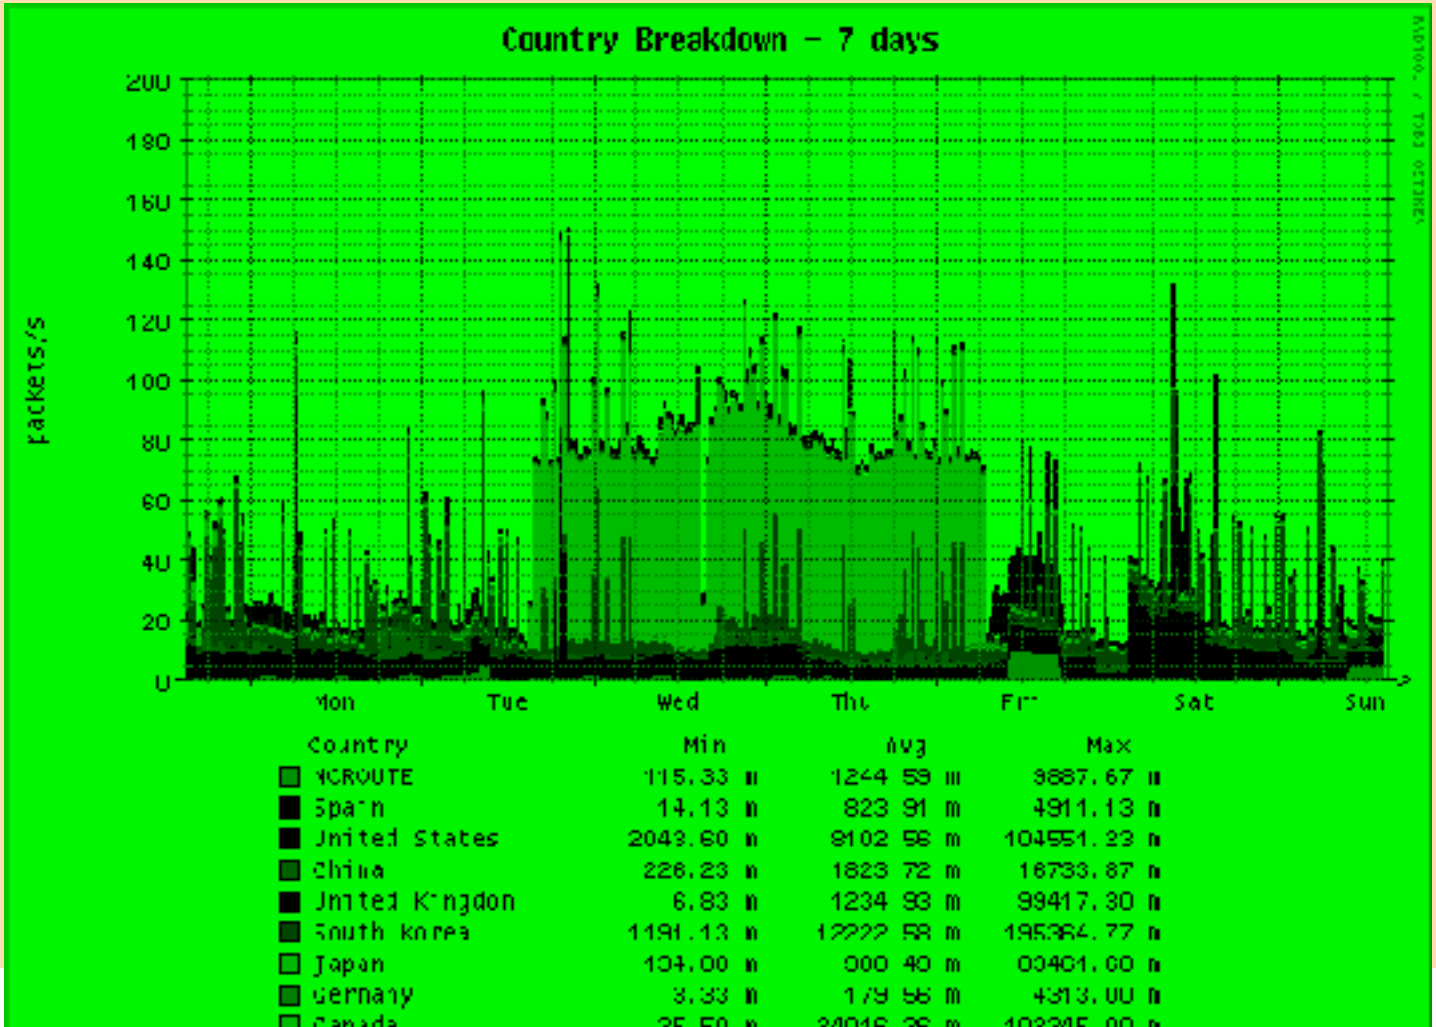

NTOS graphical interface

ongoing denial of service attacks

- 7 october 2003
- breakdown of attacked services

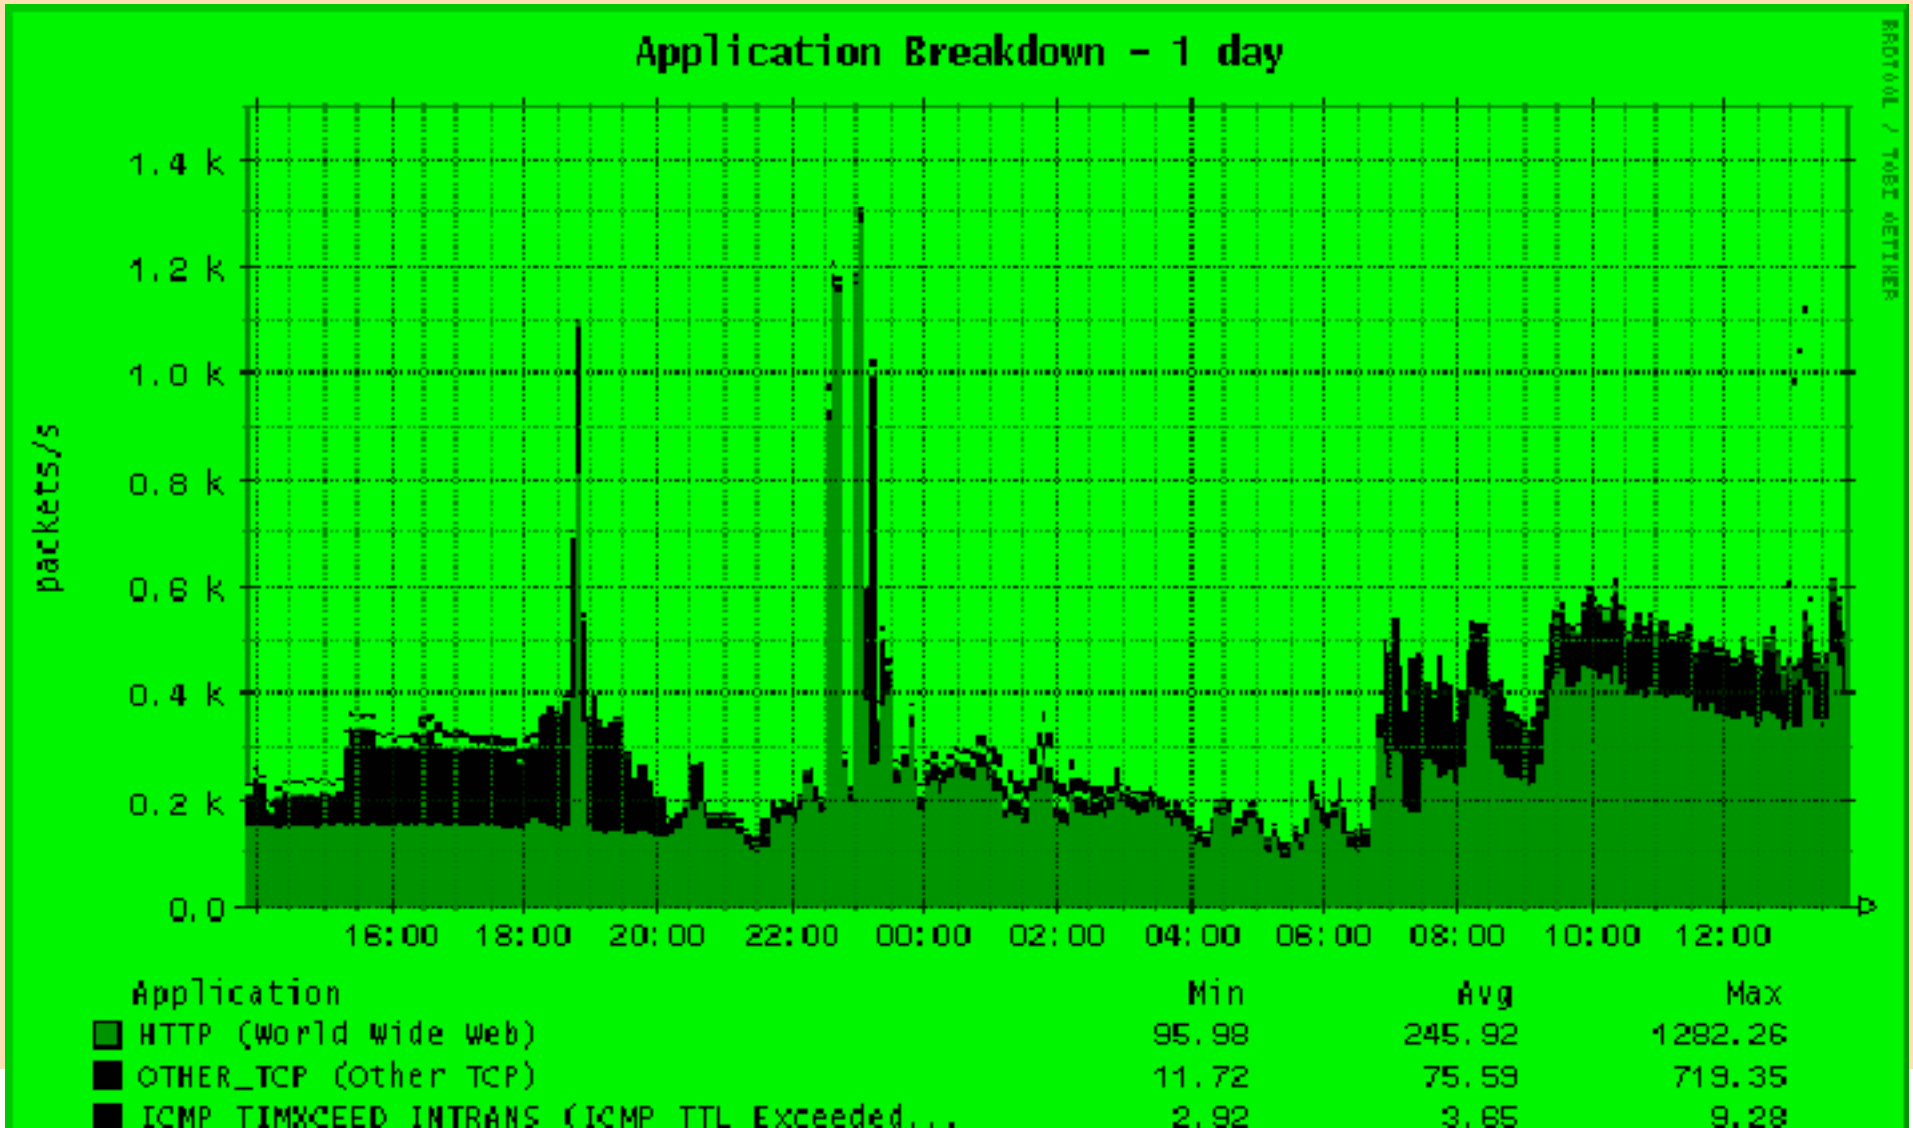

NTOS graphical interface

ongoing denial of service attacks

- 7 october 2003
- breakdown by victim location

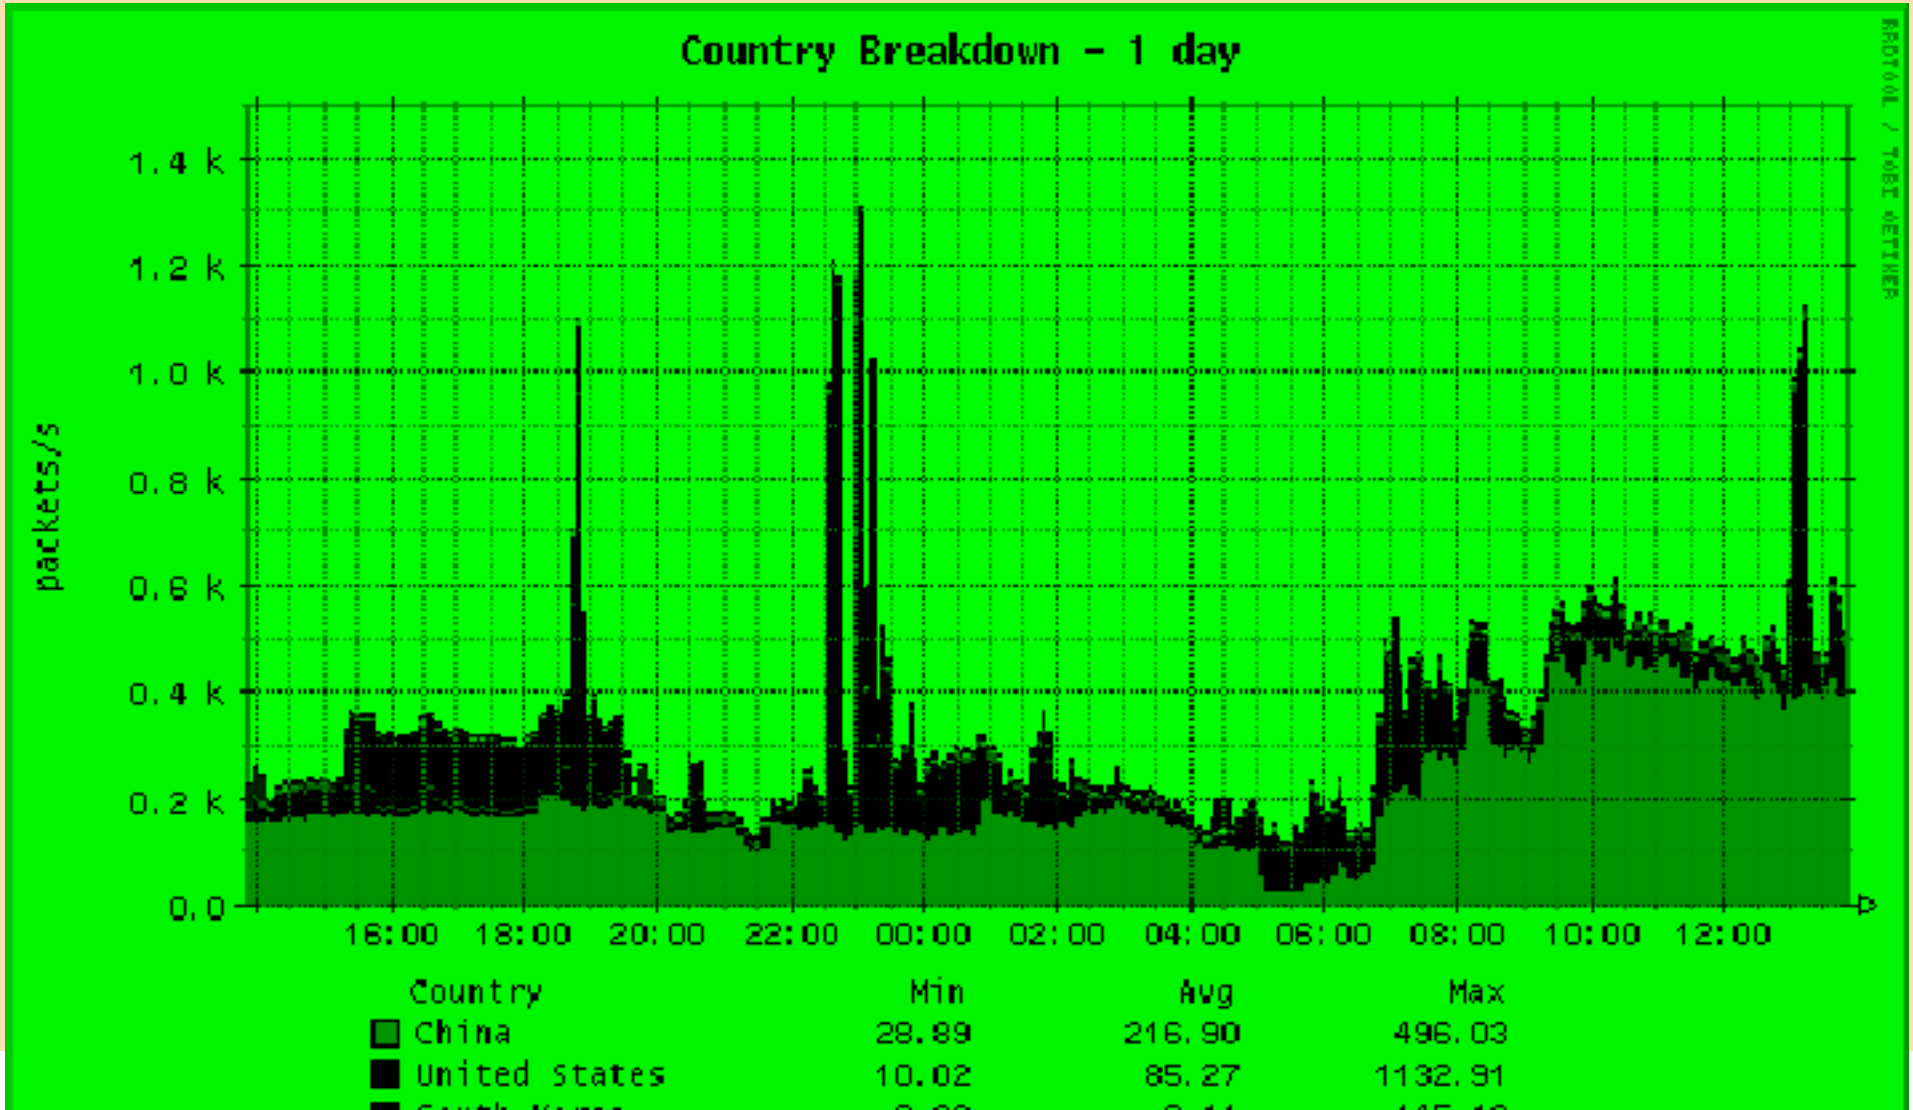

telescope: conclusions

Network Telescope Observation Station will continuously monitor worm and denial-of-service activity worldwide, archiving data for in-depth analysis.

NTOS furthers CAIDA's mission to foster communication and cooperation via collection, dissemination, and visualization of Internet data.

caida other activities

tools

- Internet measurement tool taxonomy:www.caida.org/tools/taxonomy/
  - used extensively by research and operational community
- Taxonomy of public and private performance measurement infrastructures: www.caida.org/analysis/performance/measinfra/
- CAIDA-developed tools:
  - workload: CoralReef, NeTraMet cflowd
  - topology: skitter, iffinder, gtrace
  - performance: beluga
  - IP data management utilities: arts++, netgeo
  - viz: chart:graph, walrus, rrdtool, geoplot, mapnet, otter, libsea, plot-latlong
  - dns: dnsstat, dnstop
  - mbone: mantra

caida outreach

- conference and journal publications
  - http://www.caida.org/outreach/papers/
- national and international presentations
  - http://www.caida.org/outreach/presentations/
- provide data to researchers
  - http://www.caida.org/outreach/data/
- ISMA workshops
  - http://www.caida.org/outreach/isma/
- security analysis
  - http://www.caida.org/dynamic/analysis/security/
- Internet course curriculum materials
  - http://iec.caida.org
- Internet tools taxonomy
  - http://www.caida.org/tools/taxonomy/
- Internet Atlas gallery
  - http://www.caida.org/projects/internetatlas/gallery/
- Internet measurement infrastructures
  - http://www.caida.org/analysis/performance/measinfra/
- networking research/analysis at UCSD
  - http://www.caida.org/home/about/research/

conclusions

current caida projects (apr 2004)

- [UCSD-RAMP] DARPA RAMP (UCSD CSE collaboration)
- [DOE-SciDAC] Bandwidth Estimation (bwest) [ends in 2004]
- [NSF-Trends] Correlating heterogeneous measurement data to achieve system-level analysis of Internet traffic trends
- [NSF-NCS] Inference of Internet structure (routing/topology)
- [Mbrs] Outreach to commercial ISPs and vendors
- [Cisco URB] Routing and Topology Analysis (AS ranking)
- [Cisco URB] Security: DOS attack and countermeasure analysis
- [DNS-WIDE] analysis of DNS root and gTLD nameserver system

Measurement is the link between
mathematics and science.
-Brian Ellis, 1968

kc
ucsd/sdsc/caida
kc@caida.org
www.caida.org

http://www.caida.org/outreach/presentations/