top problems of the Internet and how to help solve them

kc claffy
kc@caida.org

the significant problems we face cannot be solved
by the same level of thinking that created them.
--Albert Einstein

outline of talk

- top engineering and operational problems
  - *measurement methodology, quick rundown*
- why they persistently resist solution: common threads (non-technical)
- how different communities are inauspiciously reacting to the above
- implications for research, policy, builders

16 operational internet problems

- authentication
- security
- spam
- scalable configuration management
- robust scalability of routing system
- compromise of e2e principle
- dumb network
- measurement
- patch management
- “normal accidents”
- growth trends in traffic and user expectations
- time management and prioritization of tasks
- governance
- intellectual property and digital rights
- interdomain qos/emergency services
- inter-provider vendor/business coordination

(http://www.caida.org/outreach/presentations/2003/netproblems_lisa03)

e.g.,internet measurement: state-of-art

for U.S. inter-domain internet science, the crash happened in 1994 when the nsfnet retired...

- can’t figure out where an IP address is
- can't measure topology effectively in either direction, at any layer
- can't track propagation of a routing update across the Internet.
- can't get router to give you all available routes, just 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
- can’t trust whois registry data
- no general tool for `what's causing my problem now?’
- privacy/legal issues deter research (& it was hard in a enlightened monarchy)
- makes science challenging -- discouraging to academics

jarring observation from history of science

The modern field of elementary particle physics depended crucially on the establishment of a huge volume of data gathered mainly in the period 1945-65. Only then was it possible for the synthesis of the Standard Model to take place, 1967-74.

-- Peter Galison, Professor of History of Science and Physics, Harvard

16 operational internet problems

- authentication
- security
- spam
- scalable configuration management
- robust scalability of routing system
- compromise of e2e principle
- dumb network
- measurement
- patch management
- “normal accidents”
- growth trends in traffic and user expectations
- time management and prioritization of tasks
- governance
- intellectual property and digital rights
- interdomain qos/emergency services
- inter-provider vendor/business coordination

(http://www.caida.org/outreach/presentations/2003/netproblems_lisa03)

why so persistently unsolvable?

- 1) rooted in non-technical issues: **economics, ownership, and trust**
- 2) not amenable to purely technical solutions -- require interdisciplinary investment
- 3) not amenable to < 4-year solutions (so even academia is out)
- 4) no one owns the problems

let me just repeat that

- top unsolved problems in internet operations and engineering are rooted in **economics, ownership, and trust (EOT)**
- solutions require path “paved with incentives to motivate diverse agents to adopt it, implement it, use it, interface with it, or just tolerate it.” without that path you’ve wasted your time.
- does not mean there are not useful technical problems to work on. but there will no technical solutions that don’t solve the EOT problems.

warning: there's a problem we left out

- the economics one runs deep
- best available data suggests that moving IP packets around is not even a for-profit enterprise. not just bernie factor.
- like most large scale transport networks, it turns out. classic public good.
- noone tasked with thinking about the 25-year internet provisioning problem
- just as hard to get sound data on the economics

how unregulated players survive operating in an inherently non-profit industry

- hide the fact that you lose money by using voice side of company to subsidize developing IP habit
- file bankruptcy every few years, includes billions spent on lobbying to keep incumbents in power rather than analyzing the macroproblem
- lie to the markets to get capital, confuse markets for a decade. or two.
- long term: complex vertical integration (bad for security), infrastructure control (bad for freedom)
- don't let anyone look at the data that would facilitate analysis of provisioning models for this commodity

what are we trying to do here?

The problem faced by the Internet industry is in ensuring that each provider of infrastructure is fairly compensated when its infrastructure is used. In essence, the problem is how to distribute the revenue gained from the retail sale of Internet access and services to the providers of carriage infrastructure. ....

While it is not completely clear that the deregulated open market nature of the Internet can sustain a diverse, efficient and effective service provider industry, it is also unclear what form of regulatory constraints or intervention are appropriate, if any.
.... There is a strong risk that regulatory involvement, if applied inappropriately, will trigger structural inefficiencies that ultimately will be reflected at the consumer level in higher prices and inferior services. Competition is not an end in itself, nor is regulatory impost. The challenge here is to foster the conditions that allow the Internet to be a productive and efficient platform for all. That, for me, appears to be at the heart of the challenge of the Information Society. Geoff Huston January 2005

(more succintly:)
“We never learned how to route money.” --David Clark, MIT

public vs private provision

- we don't really have a good grasp of the economics, social, or cultural impact, so aren't in a position to really say yet how the commodity of digital information transport should be best delivered to society
- govt has kept relative distance from Internet, for better and/or worse. changing.
- we guessed a lot of things wrong
- tremendous struggles for next few decades as we learn the economics the hard way (amidst multiple sources of measurement error)

who can we get to help?

- ietf: can help write standards/BCPs but lacks operational experience/mission
- academia: can bring rigor/science to bear, but lacks access to the infrastructure as well as funding structure to tackle 10-year problems
- operators (nanog): lots of operational clue (and consumed w applying it) but lack financial, legal, and incentive/legal structure to support cooperative work
- rest of industry: lacks incentive
- government: means well but often operates out of ignorance or misapplied philosophical bent, and beholden to ‘stability of infrastructure’ (aka corporate interests)

many have insights into pieces of problem;
all just doing their job in historical context;
noone has oversight for coordination

constituency responses to situation

- publicly funded research community: battles clock; lowers standards
- u.s. nat'l science foundation: battles disciplinary focus, funding cuts, backpressure from darpa's battle with the terrorist ether
- operational community: battles malware & growth w/ minimal cooperation. it's not just a job...
- telecom: battles antique regulation, unprofitability/mergers, citizens
- copyright-owners: battles its own customers with state backing (or not)
- u.s. fcc: battles breasts, freedom, & immeasurably flailing industry
- rest of world: battles digital divide, u.s. hegemony
- UN/ITU: battles reality (it isn't pretty)
- military: battlefield ISP (profoundly important)
- people who can afford it: build their own (google, aol, DOE, .edu, pa.ca)

doesn't look good in the near term

- actually downright depressing in the short term
- solution: think long-term!
- "think about economic and business fundamentals, not accounting considerations"
- think about architectural north stars, then work backwards through economic feasibility
- requires stepping out of most contexts we operate in. also requires research funding.

implications for empirical internet research

- need to start asking questions we ask of critical infrastructure
- need vehicles to inform policy
- need vehicles to protect & analyze data
- measurement integrity is essential

implications for internet applications research

- now it's all about mechanism design
- 'goodness' of a computational artifact can't be captured by time and space performance, or even reliability, usability. need path to relevancy amidst complex socio-economic context (christos' 'algorithms, games, and the Internet')
- underlying science (game theory, complex systems, nanoeconomics) still quite young

implications for internet architecture research

- goal: design `in the light' (first time ever!)
- question first principles, e.g., cooperation vs competition (cooperative distributed systems researchers, the next decade may be fun)
- need interdisciplinary attention & investment
- need to make it way cheaper and way simpler (afaict this is blocked on yochai's 'agoraphobia')
- need longer-term thinking than any current vehicle supports, "open architecture" necessary but not sufficient
- heads up: in the meantime, incumbents are on it

implications for regulatory research

- goal: bring regulation and economic models in congruity with technology and empirical data
- anticipating future here not good enough
- separate transmission from content (classic EOT issue, decades of experience), only way to protect freedom
- cost-benefit analysis of complexity, robustness, functionality & cost. consider multiple price-points.
- universities can play a valuable role here
- so can estonia
  - http://www.privacyinternational.org/survey/phr2003/countries/estonia.htm The 1992 Estonia Constitution recognizes the right of privacy, secrecy of communications, and data protection.

implications for intellectual property

- assumption: in 25 years, everyone has unmediated connectivity to everyone else. (political certainty: all alternatives are worse.)
- stronger copyright protection for cyberspace: "desirable, inevitable, and irrelevant." (andrew odlyzko 9 years ago)
- cost of distribution -> 0. business models that depend on controlling distribution will go away
- some DRM cycles better go into information science....

implications for all of us

- policy really needs information right now
- push fiscal and other lessons outside our communities to city, state, world. teach high school students to build own nets!
- enlightened connectivity models will serve 'fittest' societies in the 21st century
- north star: most economic way to promote freedom of access to "all"
- promote open analyses of cost models: how much is unfettered p2p digital access worth?

conclusions

- in order to make any real progress on the biggest Internet problems, we're going to have to engage deeply with economics, ownership and trust issues
- leadership void regarding long-term stewardship of digital communications infrastructure
- need open, technically literate discussion of objectives, then work w/economists what ownership, regulation, & access structures come closest to meeting objectives
- cooperative architectures merit R&D

q&a

or (supplementary material)...

- what we could do now
- reading recommendations
- words to live by

california's information freeway?

- moving bits just as fundamental in 21st century as moving people and cars (hopefully more so than moving cars)
- california has history of pushing policy in the direction of science
  - e.g. proposition 71 (stem cell); proposition 215 (medical cannabis)
  - "using best available data, narrow the congruity gap between policy and science." best available data suggests that incentive structure for spectrum research is badly out of balance.
  - in character, would california demonstrate open spectrum has higher social value than cost (an EOT argument), then just open it?

pessimistic [near-term] reading

(aka “business as usual”)

- "the Internet is dying" -- Karl Auerbach provocative article
- between spam, anti-spam blacklists, rogue packets, never-forgetting search engines, viruses, old machines, bad regulatory bodies, bad implementations, Internet will lose half its users in 6 months
- in its place a much more controlled set of communications
- “digital imprimateur" -- john walker
  - how big brother and big media can put the Internet genie back in the bottle
  - rich optimistic pessimism
- geoff huston’s nznog talkt
- trashing the Internet commons: implications for ISPs’
- feb 2003, http://s2.r2.co.nz/20040129/ www.nznog.org/ghuston-trashing.pdf
- not so much with the optimism
- bruce sterling keynote at NSF workshop feb 2002
- http://www.cra.org/Activities/grand.challenges/sterling.html
- ubicomp, ultrawideband, machines-building-machines will save us

optimistic [long-term] reading

(aka "living up to the net's potential")

- eben moglen, columbia
  - freedom of thought
- lawrence lessig, stanford
  - code is law, future of ideas
- yochai benkler, yale
  - "property , commons, and the first amendment: towards a core common infrastructure"
- your kids
  - they like wiki and google more than they hate spam; they know they can code the societies they want to live in.

// It is crucial that we understand this fundamental change in the material conditions of production in the networked information economy. We find ourselves faced with policy and design questions that assume that the role of market production is fixed, rather than technologically contingent. We observe in many contexts policy choices and design impulses that take assumptions appropriate to the capital requirements of industrial economies and try to force behavior in the networked information economy into the social and market behavioral patterns that were appropriate for that technological stage and capital structure, rather than for the one we live in today. We must learn instead how to adjust our expectations, assumptions, and, ultimately, policy prescriptions to accommodate the emerging importance of social relations in general, and sharing in particular, as a modality of economic production.
- yochai benkler, sharing nicely//

If I had to unify all these battles with a single theme, it would be openness: of software, of culture, of infrastructure, of ideas. It was common to criticize the Net in the early days because of the amount of garbage that was out there. But the Net's openness to all comers also fostered pockets of insane brilliance -- and we need those more than we need overall bland mediocrity and the conservative choices that are made when the majority rules. Our world is, after all , governed by Sturgeon's law: "90 percent of everything is crap." It is those pockets of brilliance, amplified by the Net, that have the power to change the world in accordance with all our old democratic dreams of equality. Ultimately, the power of the Net is in the ideas of the people who use it.

-- Wendy Grossman,
From Anarchy to Power: The Net Comes of Age, 2001

seekers of the larger view

seven steps to trilogy of action

- discover unity
- draw together pieces of science and technology to create a system
- find the economic feasibility for a new technology by virtue of a wide grasp of the worlds of man and matter
- reach harmony through intuition, by meditating on the base of a wide and deep knowledge of the field so as to arrive at a new result
- build a model, a simplified representation of the problem at issue, subject to experimental and calculational analysis
- serve as a science-technologist generalist who, not once or twice his life, but many times in a year, and generally in the service of others, extracts the single, simple missing point out of a complicate situation
- make decisions or help others make decisions, by imaginative interaction with alternative scenarios calculated as consequent on those decisions

– john archibald wheeler