

Networking....
successes, new challenges, and an
expanding *waist*,
as the field approaches 40

Jim Kurose

Department of Computer Science
University of Massachusetts
Amherst MA
<http://www.cs.umass.edu/~kurose>

CSU ISTeC Seminar
December 4, 2003

40 years ago....

Beginnings of today's computer networks, and
applications:

- P. Baran, "On Distributed Communications Networks," 1964
- L. Kleinrock, *Communication Nets*, 1964
- Licklider & Clark, "On-Line Man Computer Communication," 1962

*... arguably, networking is ~ 40 years old,
approaching middle age (in human years)*

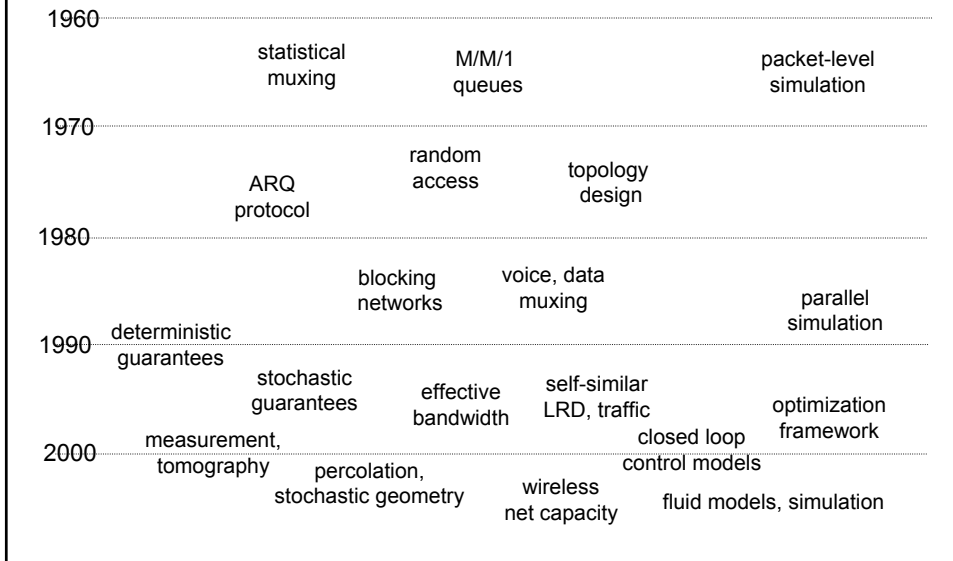
What happens as we turn 40?

- panic, denial ("I'm still 39")
- change in shape (expanding waist?)
- mid-life crisis (hopefully not!)
- more reasoned assessment:
 - ❖ where have I been, what have I accomplished?
 - ❖ what good things, exciting challenges lie ahead?

Overview

- where are we now?
- challenges on beyond today's Internet
 - ❖ sensing
 - ❖ overlays
 - ❖ managing
 - ❖ fundamentals: network science
- lessons learned

..... and understood ourselves better



Even in the last 10 years.....

□ 10 years ago: IP versus ATM

□ since then:


- QOS
- network calculi
- intserv diffserv
- signaling: RSVP
- high speed routers
- WWW protocols, caching
- CDNs
- streaming
- LRD traffic
- power laws
- TCP, closed loop control
- multicast
- measurement
- active networking
- traceback, DDoS
- BGP pathologies
- VoIP, SIP
- IP-over-*
-

many **successes** (success is in eyes of beholder)

Where are we now?

- Internet: global network infrastructure
 - ❖ huge social, economic, scientific impact
 - ❖ a huge success!
- resistance to change?
 - ❖ IPv6, native multicast, QoS?
 - ❖ backward compatibility, investments

Networking: expanding visions



Overview

- where are we now?
- challenges on beyond today's Internet
 - ❖ sensing
 - ❖ application overlays
 - ❖ fundamentals
 - ❖ education perspective
- lessons learned

Applications: drivers for change

- various forms of tele*
 - pervasive computing
 - mobility
 - homeland security
 - e-sciences
 - sensor nets
- } evolutionary?
revolutionary?
hand-maiden?



Observations:

- limited success in predicting disruptive apps
- applications drive much (but not all) of networking
- interesting diversity ... even within app area

Sensor nets

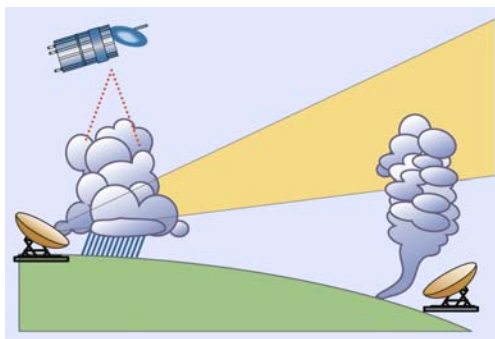


environmental sensing: ground, water, air
habitant monitoring
atmospheric sensing: meteorology
military sensing
structural monitoring



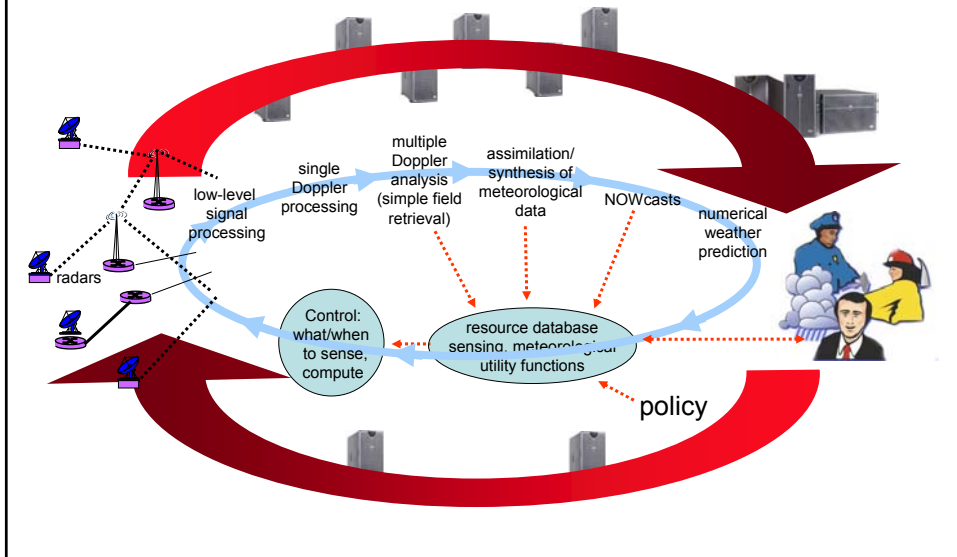
Atmospheric Sensing: application driver

Need: "insufficient knowledge about what is actually happening (or is likely to happen) at the Earth's surface where people live." [NRC 1998]




earth's curvature, widely spaced radars prevents troposphere observation below 2 km

Adaptive sensing of the atmosphere



Research Challenges:



- ❑ resource allocation: bandwidth, computation, storage
- ❑ deployable, computationally steerable sensors
 - ❖ mobile, ad-hoc
 - ❖ integration of fixed and mobile infrastructure: where are the "smarts"?
- ❑ low-power, low-duty cycle
- ❑ interdisciplinary: applications, sensing platforms
- ❑ architecture:
 - ❖ cross-layer integration
 - ❖ scale
 - ❖ security



Architecture: stovepipes

applications

habitat sensing geo sensing atmosp. sensing



habitat sensing net



atmospheric sensing nets



physical



Architecture: stovepipes or layers?

applications

habitat sensing geo sensing atmosp. sensing



habitat sensing net



atmospheric sensing nets

physical




Overview


- where are we now?
- **challenges**
 - ❖ sensing
 - ❖ **evolvability: overlays**
 - ❖ managing
 - ❖ fundamentals: network science
 - ❖ education perspective
- closing thoughts

Middle age: loss of the hourglass figure?

middle age: a narrowing mind, a widening waist




IP "hourglass"



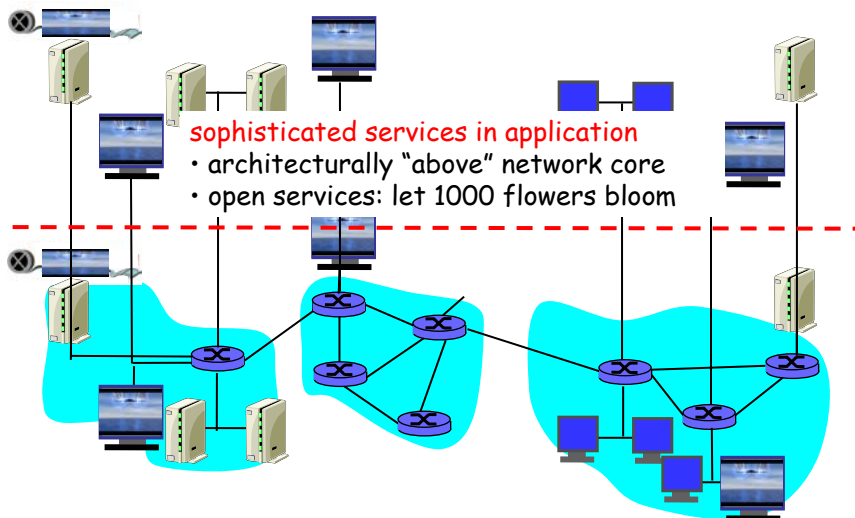
Middle-age IP "hourglass" ?

Evolvability




- difficulty: extending existing infrastructure
 - ❖ investment in legacy systems
 - ❖ economic incentives?
 - ❖ works well enough: "if it ain't broke, don't fix it"
- ossification as a result of success?
- approaches towards evolution:
 - ❖ active (extensible) networks: programmability in net "core"
 - ❖ overlays: programmability at the "edge"

Enhanced functionality: application-level



Middle-age: keeping an hourglass figure!

middle age: a expanding mind, a slim waist



Overlays:

- well-known story: layer n enhances layer $n-1$ functionality (e.g., reliable channel)
 - ❖ non-QoS to QoS-sensitive (RON)
 - ❖ non-multicast to multicast-enabled (ESM)
 - ❖ non-location-sensitive to location-sensitive
 - ❖ static to mobile
 - ❖ non-secure to secure (SoS)
 - ❖ content-blind to content-sensitive (pub/sub)

Overlays: challenges

- underlay services to support overlays
 - ❖ DHT
 - ❖ measurement
 - ❖ interface to layer below underlay
- competing control loops: underlay, overlay
- distributed versus centralized solutions?
 - ❖ various forms of p2p file sharing

Aside: appreciate your history

- much to learn from "fathers": telephone nets (signaling, reliability), ATM, DECNET, SNA
- question: lineage of p2p?

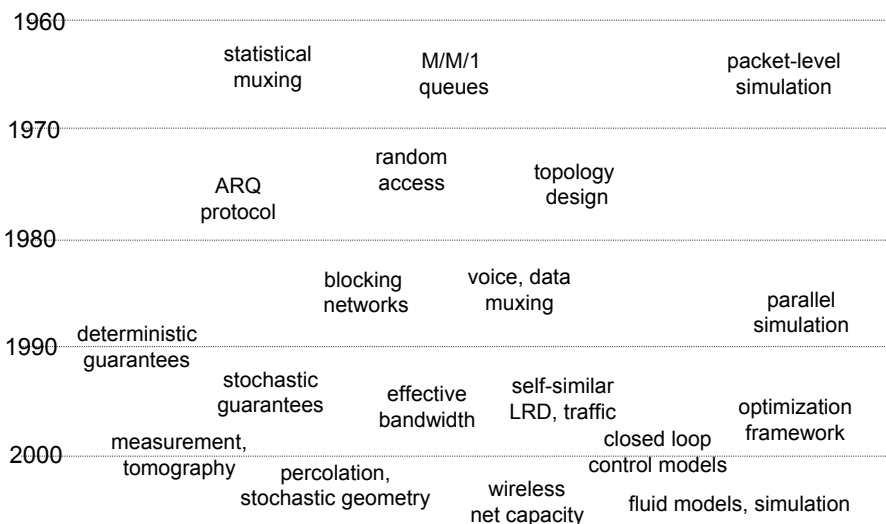
APPN

"An extension to SNA featuring (a) greater distributed network control that avoids critical hierarchical dependencies, thereby isolating the effects of single points of failure; (b) dynamic exchange of network topology information to foster ease of connection, reconfiguration, and adaptive route selection; (c) dynamic definition of network resources; and (d) automated resource registration and directory lookup" (late 1980's)

Overview

- where are we now?
- challenges on beyond today's Internet
 - ❖ sensing
 - ❖ overlays
 - ❖ fundamentals ("network science" ?)
 - ❖ education perspective
- closing thoughts

we've understood ourselves better




Challenge: on beyond the data plane

- Q: data plane performance really *the* major roadblock?
 - ❖ "robustness" ❖ adaptability
 - ❖ "complexity of control" ❖ reconfigurability
 - ❖ maintainability ❖ security
 - ❖ evolvability ❖ manageability
- the "...ities"
- Fundamental advances here are hard!
 - ❖ "efficiency" not always the most important measure
 - ❖ little/no past work on the "...ities"
 - ❖ metrics and models still to be defined

On beyond data plane: soft state control


- conventional wisdom: "soft-state is robust, less complex than hard-state signaling"
 - ❖ really?
 - ❖ how to define "robustness"?
 - ❖ how to define "complexity"?
- posing/answering such a question is:
 - ❖ *hard*: no well-accepted models, paradigms
 - ❖ *easy*: little/no past research
 - ❖ *important*: a fundamental question
 - ❖ *religious*: beliefs, rather than formal analysis

Soft-state signaling




- best effort signaling

Soft-state signaling



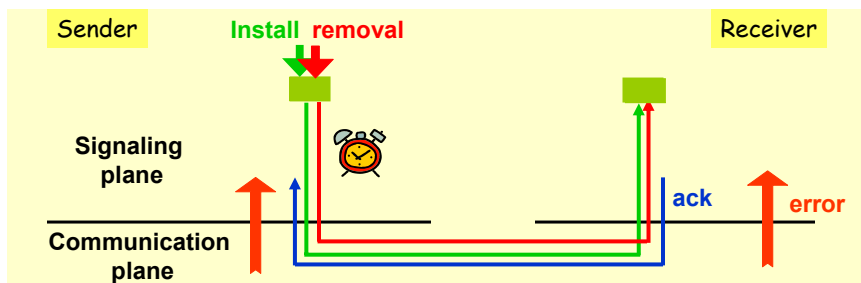
- best effort signaling
- refresh timer, periodic refresh

Soft-state signaling



- ❑ best effort signaling
- ❑ refresh timer, periodic refresh
- ❑ state time-out timer, state removal only by time-out

Hard-state signaling



- ❑ reliable signaling
- ❑ state removal by request
- ❑ requires additional error handling
 - ❖ e.g., sender failure
- ❑ telephony approach

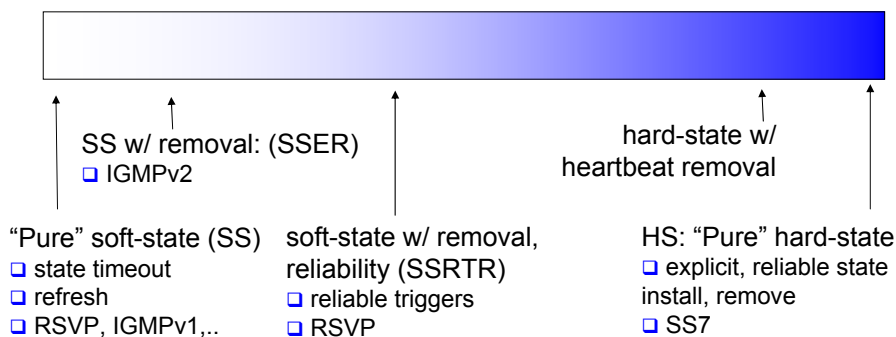
On beyond data plane: soft state control

- ❑ conventional wisdom: "soft-state is robust, less complex than hard-state signaling"
 - ❖ really?
 - ❖ how to define "robustness"?
 - ❖ how to define "complexity"?
- ❑ posing/answering such a question is:
 - ❖ *hard*: no well-accepted models, paradigms
 - ❖ *easy*: little/no past research
 - ❖ *important*: a fundamental question
 - ❖ *religious*: beliefs, rather than formal analysis

Soft-state versus hard-state

spectrum of signaling choices (not totally ordered)

- ❖ exchange of (control) messages to establish, maintain, tear down network services
- ❖ used to manage *state* in network nodes




Soft-state versus hard-state: a performance comparison

unified Markov model for approaches

- ❖ state installer, state holder: (in)consistency
- ❖ message loss
- ❖ signaling overhead cost
- ❖ cost of inconsistent state: e.g., receipt of unwanted multicast data

J. Ping, G. Zihui, J. Kurose, D. Towsley
"A comparison of hard-state and soft-state signaling protocols,"
ACM Sigcomm 2003

Soft-state versus hard-state: a performance comparison



Soft-state versus hard-state

- so far: performance only, no “..ities”
- theory for design of soft-state systems
 - ❖ self-stabilizing algorithms
- implementation, operational complexity
 - ❖ resilience to bugs, misconfiguration, attacks?
 - AT&T black holing
 - ❖ operational overhead/management required?
 - ❖ measures of such complexity?
 - ❖ quantifying performance/complexity tradeoff

On beyond data plane: traffic, routing

overheard at a major NJ research lab:

Q: “Given the network topology and the traffic matrix, how do you optimize the routes?”

A1: “Uh....”

A2: “We don't really think about it that way....”

A3: “Well, we don't know the topology, we don't know the traffic matrix, the routers don't automatically adapt the routes to the traffic, and we don't know how to optimize the routing configuration. But, other than that, we're all set!”

On beyond performance: traffic, routing

□ load-sensitive automated adaptive routing

□ topology inference
□ topology generators

□ SD traffic matrix estimation from link counts

□ optimal routing

A3: "Well, we don't know the topology, we don't know the traffic matrix, the routers don't automatically adapt the routes to the traffic, and we don't know how to optimize the routing configuration. But, other than that, we're all set!"

On beyond performance: traffic, routing

Q: To what extent are these questions important?


A: Some *very* important, others still unknown

[These problems are solvable!](#)

A3: "Well, we don't know the topology, we don't know the traffic matrix, the routers don't automatically adapt the routes to the traffic, and we don't know how to optimize the routing configuration. But, other than that, we're all set!"

The right level of complexity


Q: What process determines the “right” level of complexity?



[adapted from Hluchy 2001]

On being the right size

“For every type of animal there is a most convenient size, and a large change in size inevitably carries with it a change of form” [J. Haldane, 1928]



On being the right complexity?

For every type of networked system, there is a most convenient complexity of control, and a large change in size or function inevitably carries with it a change of form of control...

Network theory and practice: ... many many challenges

Additionally...

- ❑ wireless networks: capacity, percolation
- ❑ resilience to faults, misconfiguration, bugs
- ❑ pricing, economics
- ❑ measurement:
 - ❖ techniques: sampling, inference, signal analysis
 - ❖ understanding behavior: congestion, control reconfiguration
- ❑ management
 - ❖ auto-configuration, rapid deployment
 - ❖ diagnosing system/service-level failure
- ❑ mobility
- ❑ security

Overview

- ❑ where are we now?
- ❑ challenges
 - ❖ sensing
 - ❖ evolvability: overlays
 - ❖ fundamentals
 - ❖ education perspective
- ❑ closing thoughts

Networking Education: teach your children well

- intro networking course: standardized
- 2nd course?
 - ❖ internet internals (Stevens)
 - ❖ hot (single) topics (p2p, security, wireless, multimedia, modeling)
 - ❖ other fields: standardized graduate core(?)
- graduate core: is there a canon?
- topics with a long half life
 - ❖ useful to your children and grandchildren
 - ❖ beyond Internet: circuit-switching, signaling
 - your parents (Bellheads) know more than you may think

A graduate core: a modest proposal

- | | |
|---|---|
| □ design principles: <ul style="list-style-type: none">❖ multiplexing❖ virtualization❖ signaling, state❖ scale❖ randomization❖ indirection | □ architecture: the big picture |
| | □ implementation: <ul style="list-style-type: none">❖ principles❖ walk-through |
| | □ protocol specific, verif |
| | □ network algorithmics |
| | □ simulation |

Overview

- where are we now?
- challenges
 - ❖ sensing
 - ❖ evolvability: overlays
 - ❖ managing
 - ❖ fundamentals: network science
 - ❖ education perspective
- closing thoughts

The constancy of change

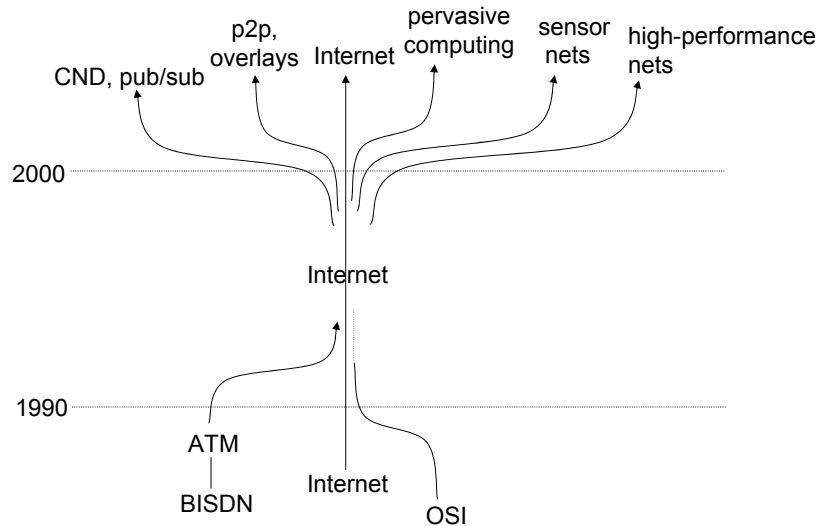
Q: What will do to today's (Internet) network what the Internet did to the telephone network?

- ❖ "Bellheads": today's network of the past
- ❖ "IP hourglass heads:" designing tomorrows' network of the past today?

Q: is a new community needed to think radically?

- ❖ A: not necessarily, but need to think out-of-the-box, driven by application needs

Networking: an exciting time!



Summary: advice to students

- ❑ *lots* of successes to be proud of!
- ❑ *lots* of interesting on-going efforts
- ❑ *lots* of interesting unanswered questions
 - ❖ fundamental questions with large half life (thinking outside the box)
 - ❖ "on beyond data plane": the "...ities"
 - ❖ applications, applications, applications

The end

Thanks!

- **acknowledgements:** D. Towsley, H. Schulzrinne, M. Ammar, L. Zhang, several generations of students
- **slides available:**
http://gaia.cs.umass.edu/kurose/talks/gaia.cs.umass.edu/csu_istec_2003.pdf