

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>

IEEE Infocom 2004 Keynote
Hong Kong 2004

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)*

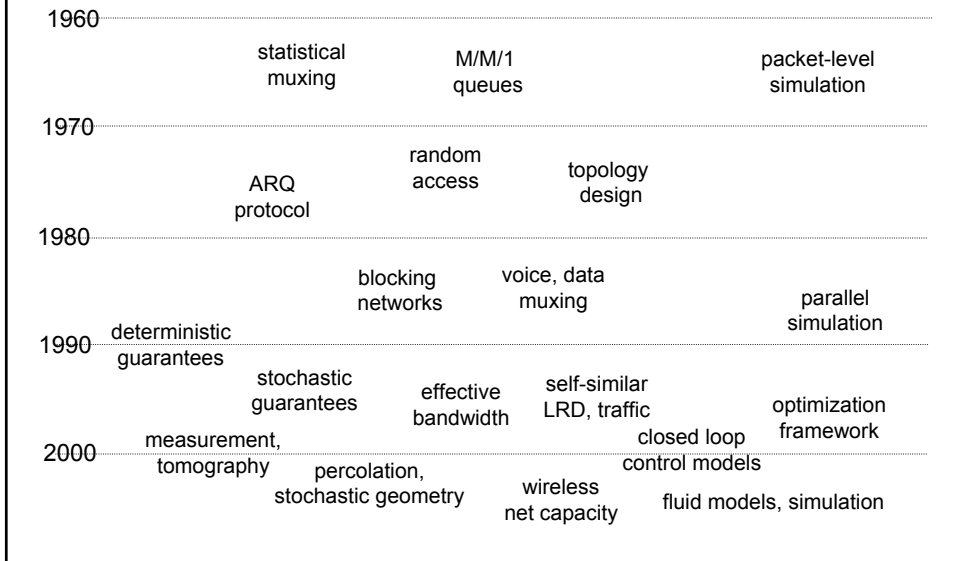
What happens as we turn 40?

- ❑ 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?
 - ❖ how to spend my time: what good things, exciting challenges lie ahead?

Overview

- ❑ where are we now?
- ❑ challenges on beyond today’s Internet
 - ❖ driven by applications
 - ❖ driven by fundamentals: network science
- ❑ closing thoughts

..... and understood ourselves better



Even in the last 10 years.....

□ 10 years ago: IP versus ATM

□ since then:

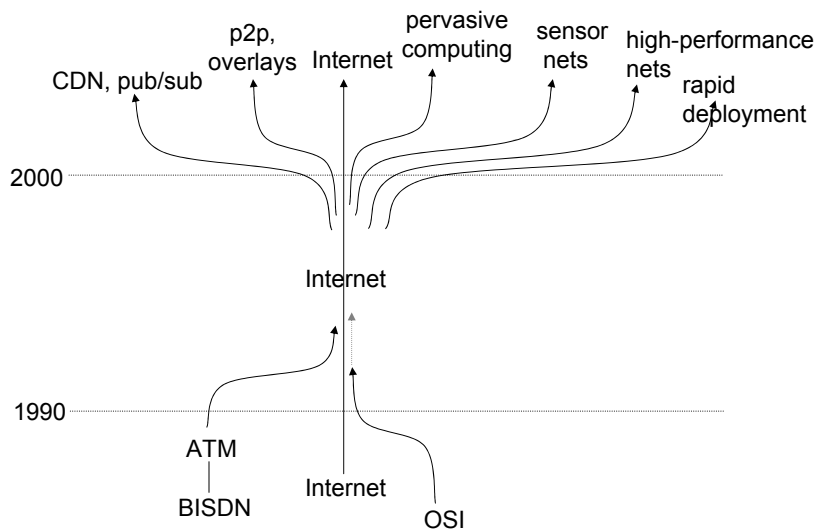
- QOS
- network calculi
- intserv diffserv
- signaling: RSVP
- high speed routers
- WWW protocols, caching
- CDNs
- streaming
- LRD traffic
- power laws
- TCP, closed loop control
- multicast
- measurement
- active networking
- traceback, DDoS
- BGP pathologies
- VoIP, SIP
- IP-over-*
-

many, many **successes!**

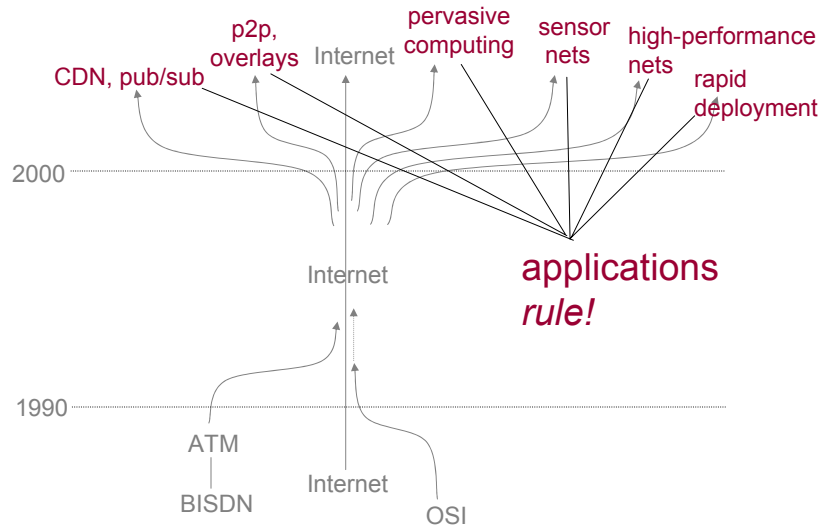
Overview

- where are we now?
- challenges: beyond today's network
 - ❖ driven by applications
 - ❖ driven by fundamentals: network science
- lessons learned

Networking: expanding visions




Networking: expanding visions




Sensor nets: wide range of characteristics

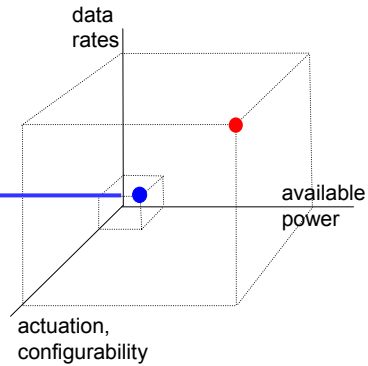
- ❑ **power**: constrained or “plugged in”?
- ❑ **data rate**: bit rate, duty cycle?
- ❑ **reconfigurability**: retasking, retargeting how often?
- ❑ **users**: single-purpose, many?
- ❑ in-situ or remote



Wide range of sensor nets: embedded



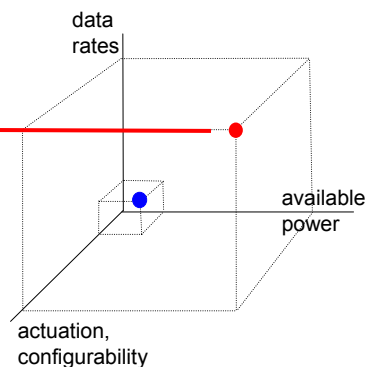
- ❑ power constrained sensors
- ❑ mostly data “push,” re-tasking possible
- ❑ low bit rate data
- ❑ network design: from scratch



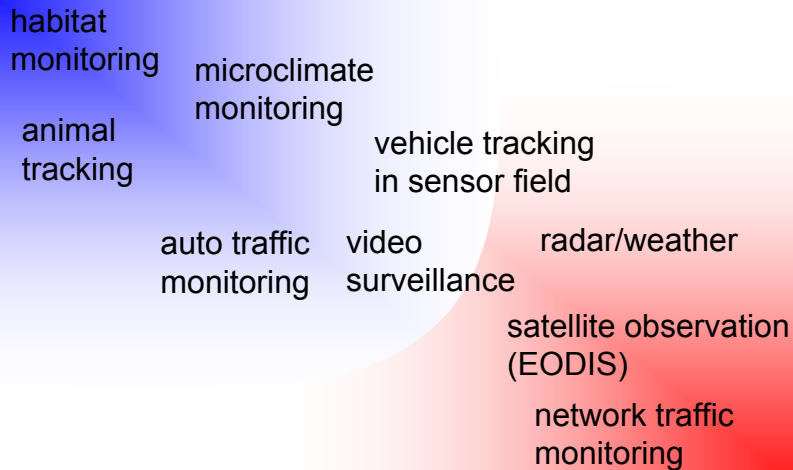
Wide range of “sensor nets”



- ❑ powered radars
- ❑ rapidly steerable beams:
- ❑ data rates: 2 Mbps - 100 Mbps per radar
- ❑ multiple data consumers
- ❑ network design space: above IP?



Wide range of “sensor nets”

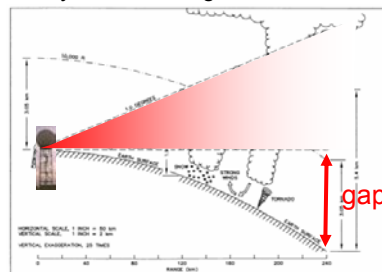


Atmospheric sensing: application driver

dense *network* of low power radars:


- ❑ *overcome blockage*: sense lower 3 km of earth's atmosphere
- ❑ *collaborating radars*:
 - ❖ improved sensing
 - ❖ improved detection, prediction
- ❑ responsive to multiple *end-user needs*

Today's meteorological radars:

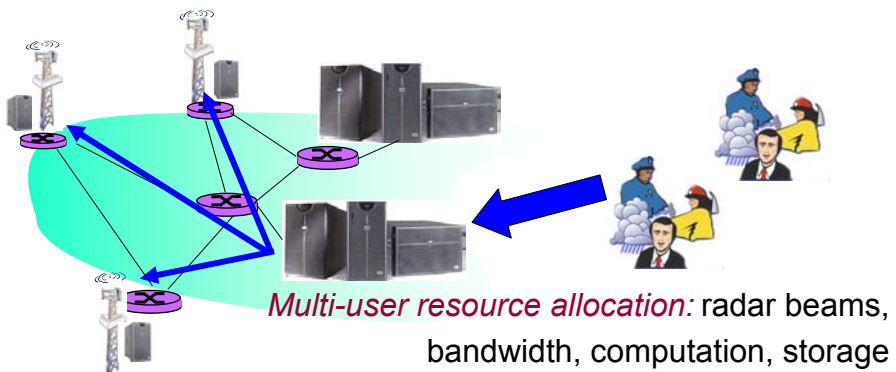


CASA: collaborative adaptive sensing of the atmosphere

NetRad: research challenges




NetRad: research challenges



NetRad: research challenges


networking:

- ❑ end-point congestion control,
- ❑ overlay routing
- ❑ monitoring management



NetRad: research challenges



data: application-level reliability semantics



Architecture: stovepipes

applications

habitat sensing geo sensing atmp. sensing



habitat sensing net



atmospheric sensing nets



physical



Architecture: stovepipes or layers?

applications

habitat sensing geo sensing atmp. sensing



habitat sensing net




atmospheric sensing nets

physical




Losing the hourglass: accommodating new applications

middle age: a narrowing mind, a widening waist




IP "hourglass"



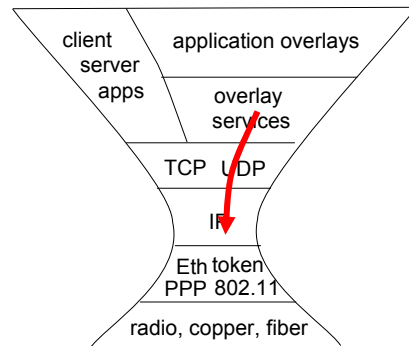
Middle-age IP "hourglass" ?

Losing the hourglass: accommodating new applications

middle age: a expanding mind, a slim waist



IP "hourglass"



Overview

- where are we now?
- challenges: beyond today's network
 - ❖ driven by applications
 - ❖ **driven by fundamentals: network science**
- lessons learned

Network theory and practice: ... many many challenges

- wireless networks: capacity, percolation
- resilience to faults, misconfiguration, bugs
- pricing, economics
 - ❖ utility-based view of protocols
- measurement:
 - ❖ techniques: sampling, inference, signal analysis
- management
 - ❖ auto-configuration, rapid deployment
- mobility
- security

Challenge: on beyond the data plane

Q: data plane performance really *the* major roadblock?

- ❖ “robustness”
- ❖ “complexity of control”
- ❖ maintainability
- ❖ evolvability
- ❖ adaptability
- ❖ reconfigurability
- ❖ security
- ❖ manageability

the “.....ities”

Fundamental advances here are hard!

- ❖ “efficiency” not always the most important measure
- ❖ little/no past work on the “...ities”
- ❖ goal: science of design

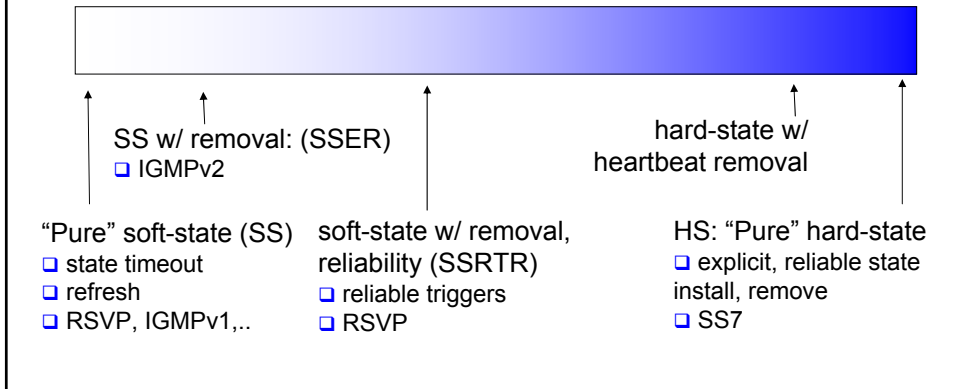
On beyond data plane: soft state control

- ❑ 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

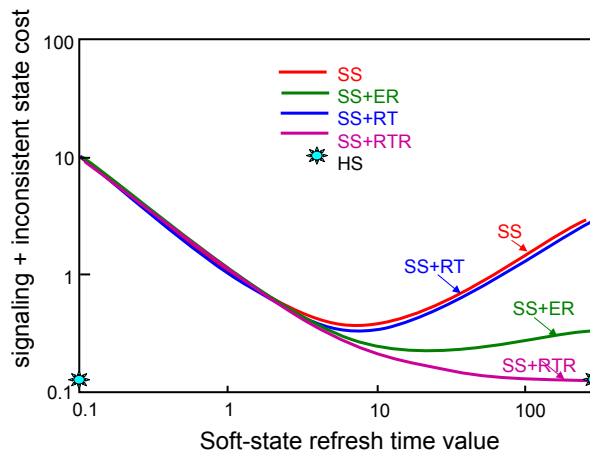
Soft-state versus hard-state

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



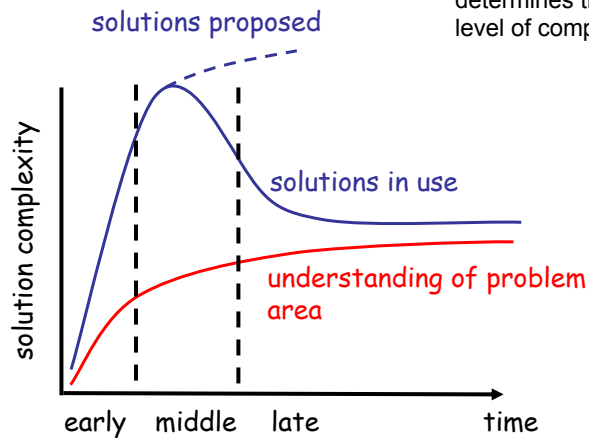
J. Ping, G. Zihui, J. Kurose, D. Towsley
 “A comparison of hard-state and soft-state signaling protocols,”
 ACM Sigcomm 2003

Soft-state versus hard-state

- 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

The right level of complexity


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



[adapted from Hluchyj 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...

On beyond data plane: traffic, routing

overheard a few years ago 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

□ SD traffic matrix estimation from link counts

□ topology inference
□ topology generators

□ robust 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!"

Overview

- where are we now?
- challenges on beyond today's Internet
 - ❖ driven by applications
 - ❖ driven by fundamentals: network science
- closing thoughts

The constancy of change

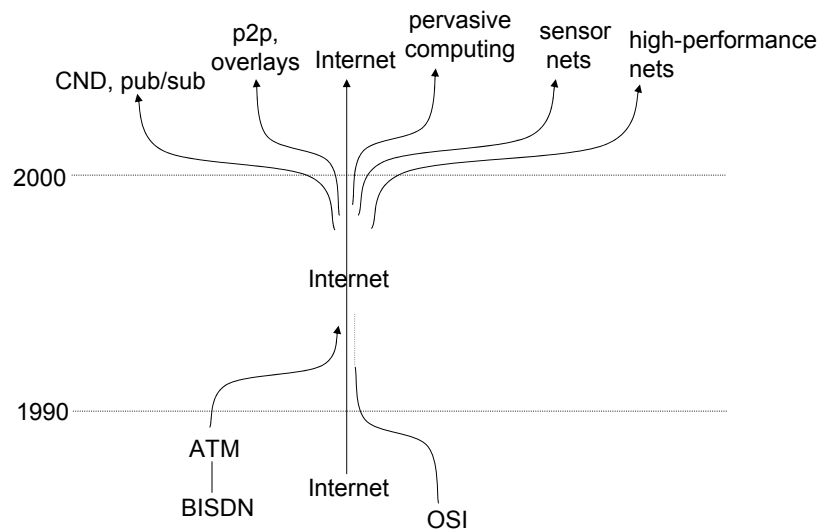
Q: What will do to today's (Internet) network what the Internet did to the telephone network?

- ❖ “Bellheads”: today's network of the past
- ❖ “IP hourglass heads:” designing tomorrow's 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
 - ❖ applications, applications, applications
 - ❖ fundamental questions with large half life (thinking outside the box)
 - ❖ “on beyond data plane”: the “...ities”
 - ❖ disruptive technology push

Thanks!

- *acknowledgements*: D. Towsley, M. Ammar, C. Diot, H. Schulzrinne, G. Parulkar, L. Zhang, several generations of students
- *slides available*: <http://gaia.cs.umass.edu/kurose/talks/>