Networking....

research trends and challenges for the coming decade

UNIVERSITY OF Massachusetts Amherst

Jim Kurose
Department of Computer Science
University of Massachusetts

Overview

- where are we now and how did we get here?
- challenges on beyond today's Internet
  - driven by applications
  - driven by fundamentals: network science
- closing thoughts
- discussion

Networking: expanding visions

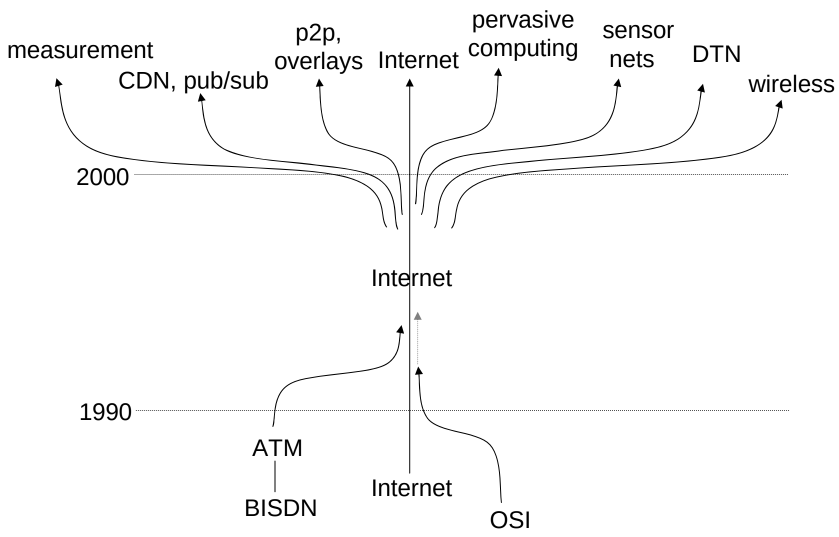

Networking: expanding visions

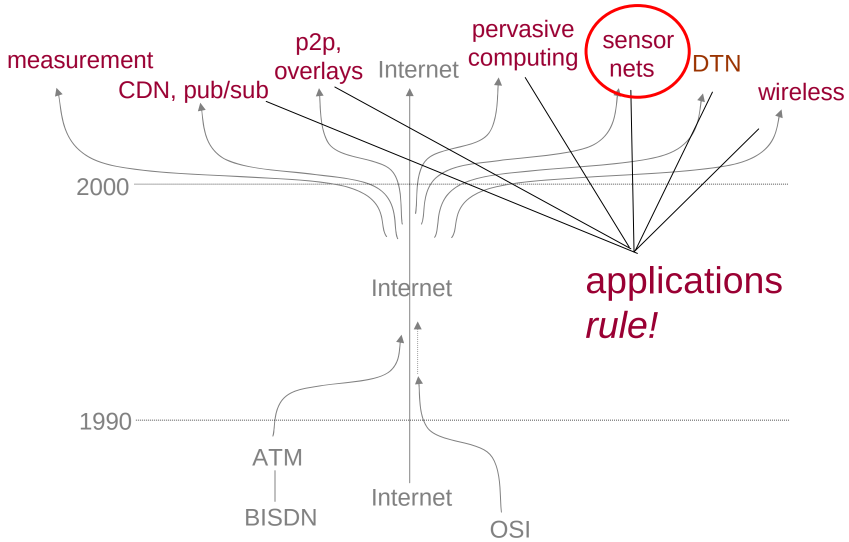

Sensor nets: wide range of characteristics

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

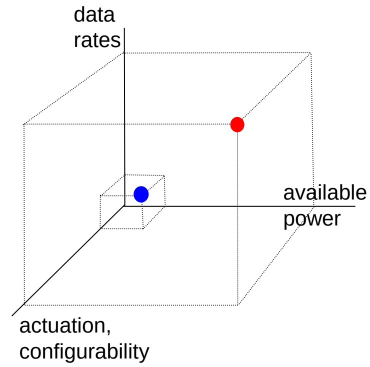

Wide range of sensor nets: embedded

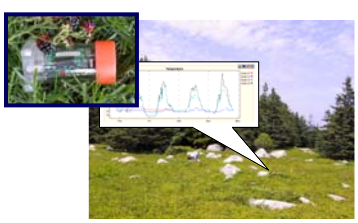

- power constrained sensors
- mostly data "push," re-tasking possible
- low bit rate data
- network design: from scratch

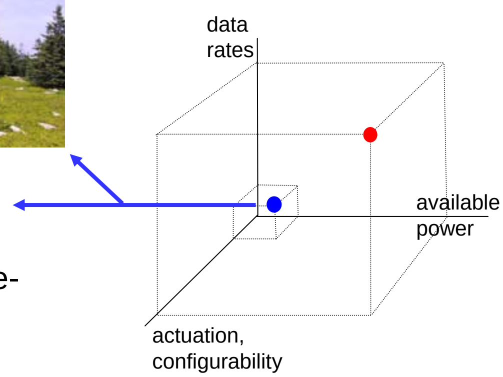

Embedded Networked Sensing Apps

- *embedded* micro-sensors, on-board processing, wireless interfaces at very small scale
  - in-situ sensing: need to "be there," monitor "up close"
- spatially, temporally dense environmental monitoring

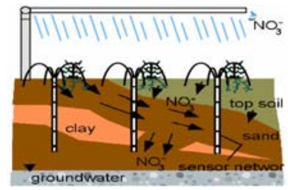

contaminant transport

seismic structure response

ecosystems, biocomplexity

marine microorganisms

Slide courtesy of D. Estrin

Wide range of sensor nets: "plugged in"

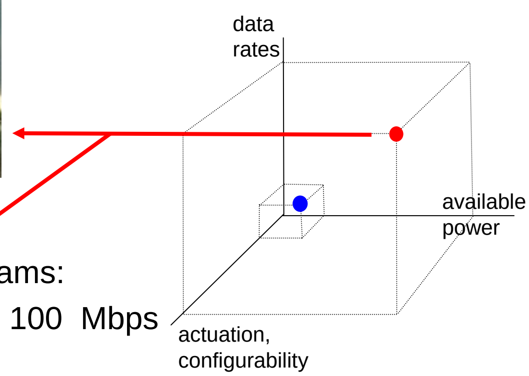

- powered radars
- rapidly steerable beams:
- data rates: 2 Mbps - 100 Mbps per radar
- multiple data consumers
- network design space: above IP?

Wide range of "sensor nets": sense-and-response systems

- habitat monitoring
- microclimate monitoring
- animal tracking
- vehicle tracking in sensor field
- auto traffic monitoring
- video surveillance
- radar/weather
- satellite observation (EODIS)
- network traffic monitoring

Atmospheric sensing: application driver

*Today: S*parse, high-power radar

- sensing gap: Earth curvature effects prevent 72% of the troposphere below 1 km from being observed
- coarse resolution

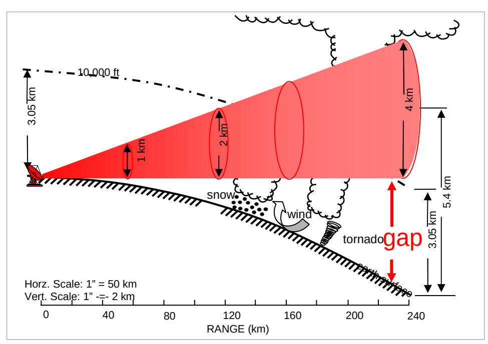

"There is insufficient knowledge about what is actually happening (or is likely to happen) at the Earth's surface where people live." [NRC 1998]

CASA: collaborative adaptive sensing of the atmosphere

CASA: dense network of low power radars:

- sense lower 3 km of earth's atmosphere
- *collaborating* radars:
  - improved sensing
  - improved detection, prediction
- responsive to multiple *end-user needs: sample atmosphere when and where end user needs are greatest*

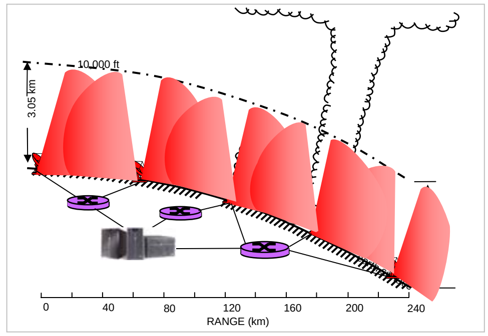

A multidisciplinary approac is required!

CASA
Collaborative
Adaptive
Sensing
of the Atmosphere

NSF
Engineering Research Center
Sept. 2003

Remote sensing
Microwave engineering
Networking
Real-time systems
Storage
Numerical prediction
Emergency management
Radar meteorology
Quantitative inversion
Climate studies
Social impact
Antenna design

core partners

expertise

Oklahoma 4-Node Test Bed

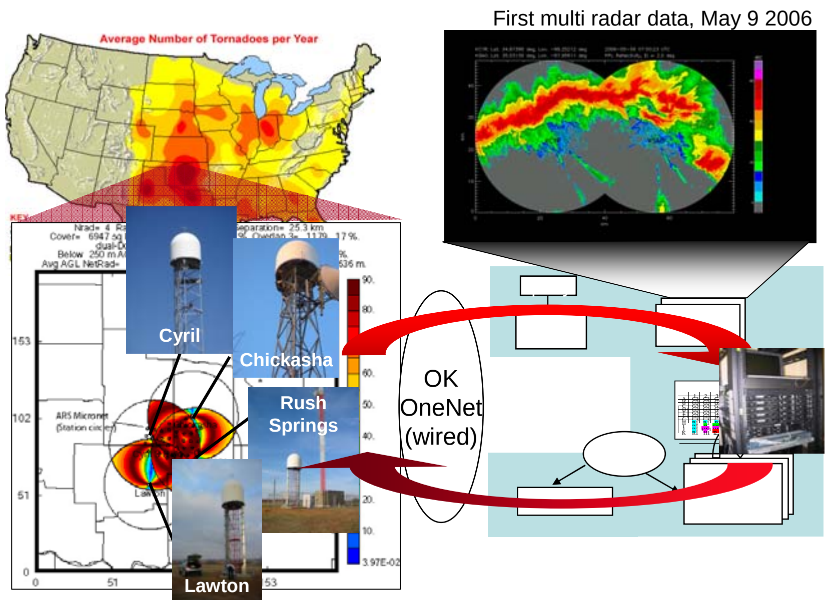

Puerto Rico, Off-the-Grid Test Bed

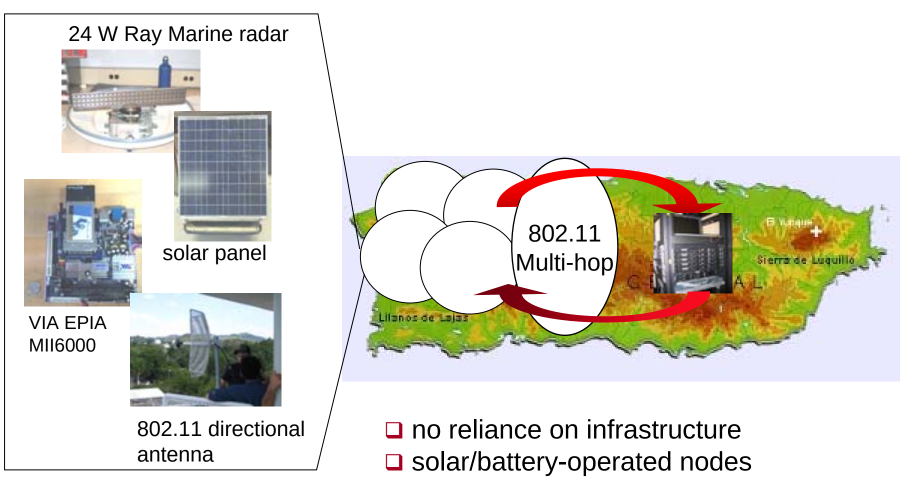

- no reliance on infrastructure
- solar/battery-operated nodes
- multi-antenna multi-hop 802.11 directional antenna

Research efforts

- application-level congestion control
- joint sensing/routing optimization in remote sensing networks
- architecture: separation of control and data
- background data transfer
  - TCP-nice
  - distributed, application routed, data-file-transfer (many-to-1)
- optimization: planning, prediction (Kalman filtering) for radar targeting

Architecture overview

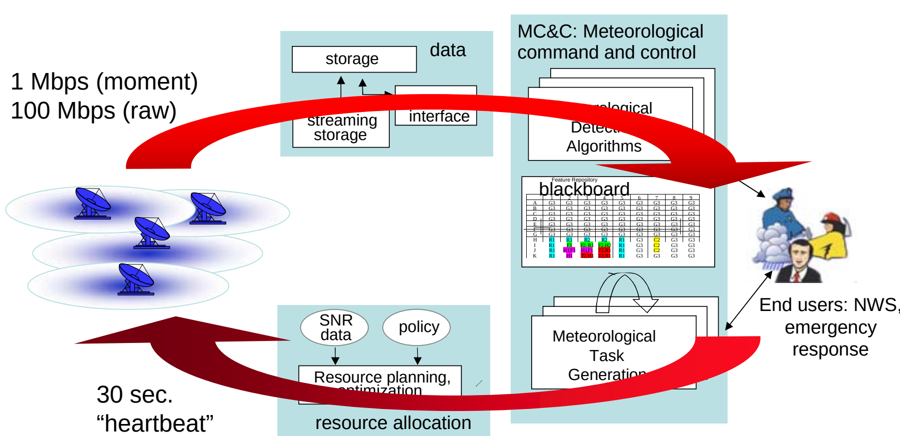

Optimized radar targeting

- where to point the radars?: “where user needs are greatest”
- multiple “users”
  - different sensing needs (e.g., beam targeting)
  - different utility
- policy
- environment (SNR) impacts sensing ability

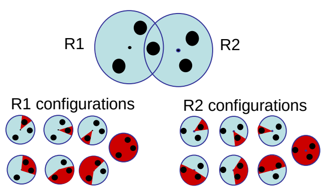

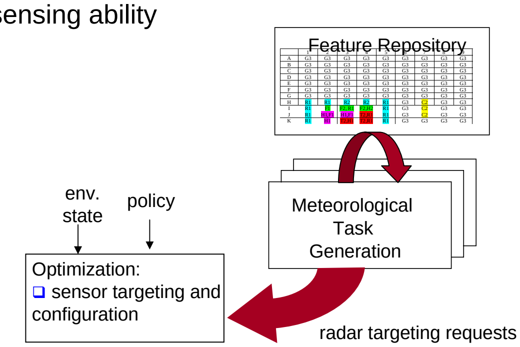

Optimizing Radar Targeting: incorporating end user considerations

Where to point?

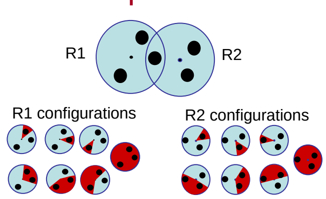

What to optimize?

Find ***configuration*** that gives the best quality data for the highest utility tasks (phenomena) at time step k::

$$J = \max_{configurations, C} \sum_{tasks, t} U(t,k) Q(t,C)$$

Utility – **“how important”** is task *t* to the users at time k?

$$U(t,k) = \sum_{groups, g} w_g U_g(t,k)$$

defining overall utility in terms of individual group utility

Quality – **“how good”** is scanning configuration *C (distance, coverage, # radars) for task t?*

Optimizing Radar Targeting: architecture!

What to optimize?

Find ***configuration*** that gives the best quality data for the highest utility tasks (phenomena) at time step k::

$$J = \max_{configurations,C} \sum_{tasks,t} U(t,k)Q(t,C)$$

- separation of "how important," *U(t,k),* from how good ,*Q(t,C)*
- *U(t,k,Q(t,C))* would have been possible but:
  - complex to solve
  - complex to specify and update *U(t,k,Q(t,C))*
  - "stovepipe" design

Emulator: comparing scan strategies

- utility comparison: sit-and-spin versus sectored scan
  - utility gained according to MC&C-defined utility of scanning under configuration C

$$J = \max_{configurations,C} \sum_{tasks,t} U(t,k)Q(t,C)$$

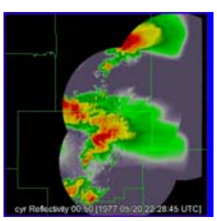

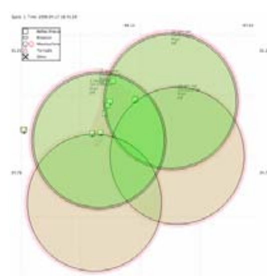

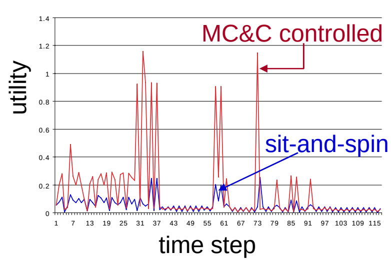

How to define "how important": $U_g(t,k)$

- "naturally": group-sensitive utility for each feature (tornado, wind shear, hail core) scanned
- … and the survey says….

User feedback:

- NWS: want "mental movie" scanning "areas of interest" at regular intervals
- Feature-based too "jumpy"
- Need context: scan areas around features

User Utility Rules (revised)

- *interval-based preferences:* "do X every Y time units"

Rules	Rule trigger	Sector Selection	Elevations	# radars	Contig.	Sampling interval
NWS						
N1	time	360	lowest	1	Yes	1 / min
N2	storm	task size	[illegible]	1	Yes	1 / 2.5 min
EMs						
E1	time	360	[illegible]	1	Yes	1 / min
E2	reflectivity over AOI	task size	lowest	1	Yes	1 / min
E3	velocity over AOI	task size	lowest	2+	Yes	1/ 2.5 min

The *really* big picture

- importance of user requirements

“It’s the end-user, stupid”

“It’s the application, stupid”

“It’s the network, stupid”

- how to embed user requirements into network?
  - sensor networks
  - content distribution
  - special-purpose overlays

The big picture: again

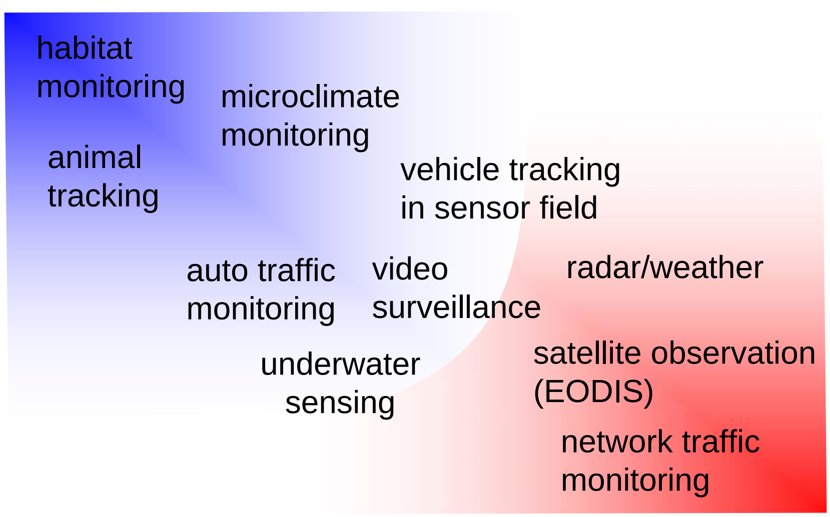

in spite of differences, commonalities as well!

The big picture: stovepipes or layers

applications

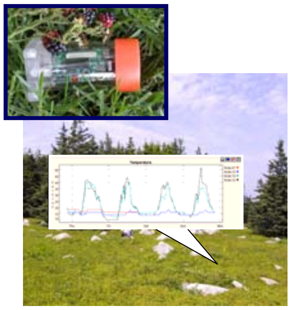

habitat sensing net

habitat sensing | geo sensing | atmosp. sensing

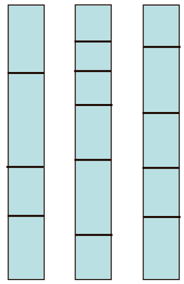

physical

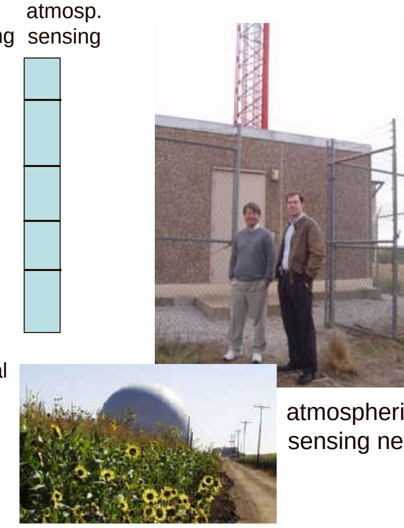

atmospheric sensing nets

The big picture: stovepipes or layers

applications

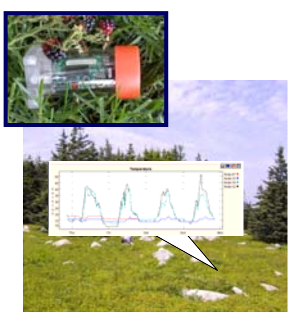

habitat sensing net

habitat sensing | geo sensing | atmosp. sensing

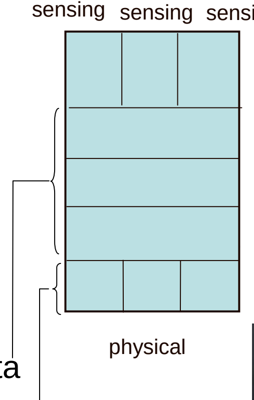

Data

Network

physical

atmospheric sensing nets

Big picture: supporting new applications – losing the IP hour glass figure?

middle age: a narrowing mind, a widening waist

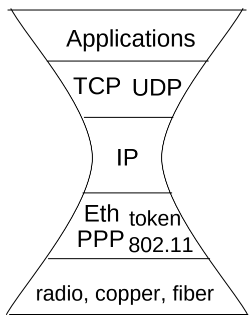

IP “hourglass”

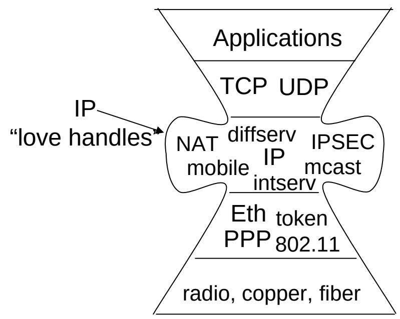

Middle-age IP “hourglass” ?

Big picture: supporting new applications – losing the IP hour glass figure?

middle age: a expanding mind, a slim waist

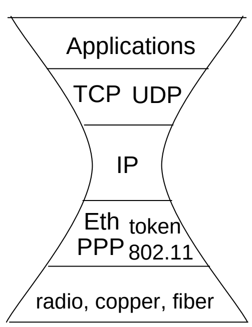

IP “hourglass”

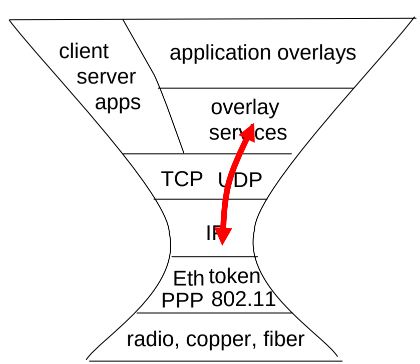

Overlays: bringing networking techniques up to the application

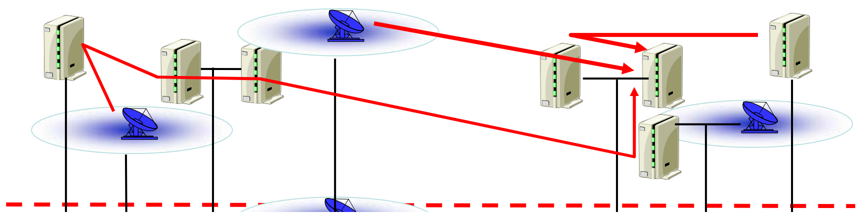

Overview

- 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, coverage
- pricing, economics
  - utility-based view of protocols
- measurement:
  - techniques: sampling, inference, signal analysis
- management
  - auto-configuration, rapid deployment, resilience to faults, misconfiguration, bugs
- overlay networks
  - competing levels of underlay/overlay control
- mobility
- security

Challenge: on beyond the data plane

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

- "robustness" (non-fragility)
- complexity of control
- maintainability
- evolvability
- adaptability
- reconfigurability
- security
- manageability

the "X-ities"

Fundamental advances here are hard!

- "efficiency" not always the most important measure
  - tradeoff between x-ity, expected performance
- little/no past work on the "X-ities"

Theme: performance "fragility"

- use protocols, parameter settings that work well for wide range of scenarios
  - avoid "fragile" solutions for "robust" ones
- *robust optimization:* optimize expected performance plus variability measure
- *routing:* choose routes that work well for *wide range* of traffic matrices
  - "On optimal routing with multiple traffic matrices," C. Zhang, Y Liu, W. Gong, J. Kurose, R. Moll, D. Towsley, *IEEE Infocom 2005*
- *overlays:* ability to optimize routing in overlay compensates for "poor" underlay routing

(Non) fragility: routing

- accommodate wide range of conditions, not so highly optimized to fail catastrophically when operating beyond "normal" operating regime.
- traditionally: optimize for known TM, but….
  - uncertainty in TMs
  - TMs change (e.g. time of day)
- "*robust" routing:* performs "well" over range of TMs, even if non-optimally for a given TM
  - oblivious routing (AT&T)

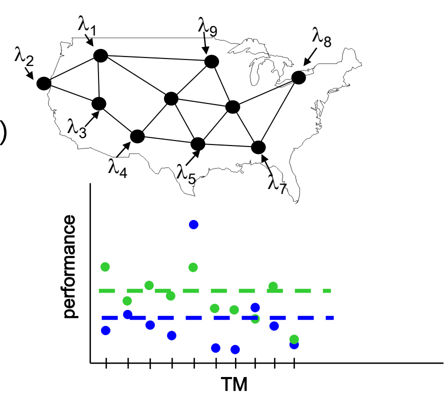

Routing: avoiding fragile solutions

- avoiding "fragility" (achieving robustness): routing with
  - good expected case performance
  - avoiding "terrible" worst case performance
- tradeoff between average and worst case performance
  - optimizing weighed sum of expected, worst case

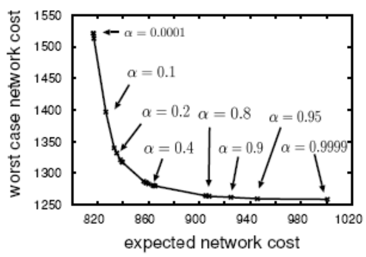

AT&T PoP-level topology
(400 OD pairs)
AT&T inferred hourly TMs

- other changes: link/node failures, impairments, compromise

Robust routing in wireless networks

- "robust" with respect to changes in link connectivity (mobility)
  - similarity with failover/backup routing
- Idea: use multiple paths at transport layer
  - Can move flow rate among different transport-layer paths as underlying network paths change
- what is right *timescale* for optimization?
  - joint optimization: each/every topology change?
  - decoupled/layered:
    - optimize (network layer) routing at slow time scale
    - multipath endpoint rate adaption (transport layer) at finer time scale

Layering provides for timescale decomposition

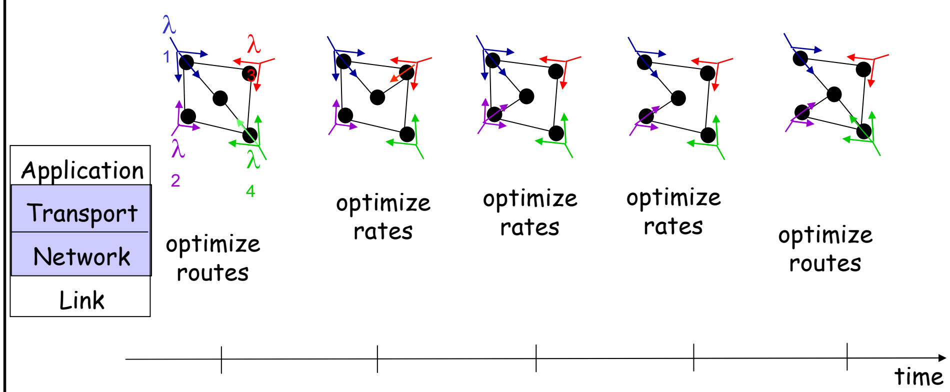

Research Questions:

- specific routing, multipath transport protocols
- performance penalty for decoupled optimizations?
  - scenarios (mobility, topology) when penalty is high?
  - function of timescale difference
  - similarities with underlay/overlay optimization
  - initially investigate via simulation
- performance gains for decoupled optimizations
  - quantify decrease in signaling

The “right” level of complexity

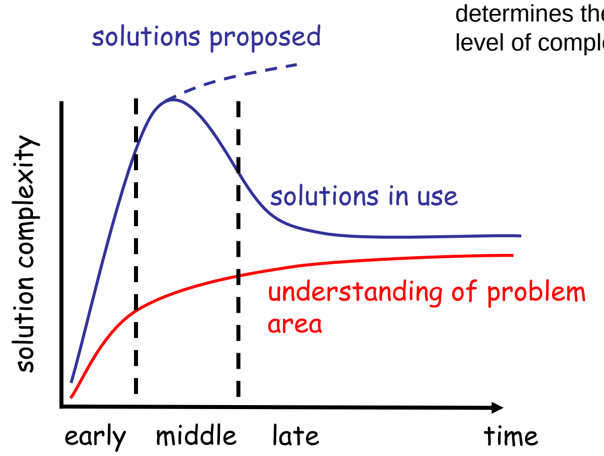

Q: What process determines the “right” level of complexity?

[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]

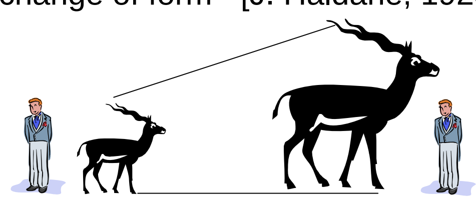

On being the right complexity?

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…

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 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!

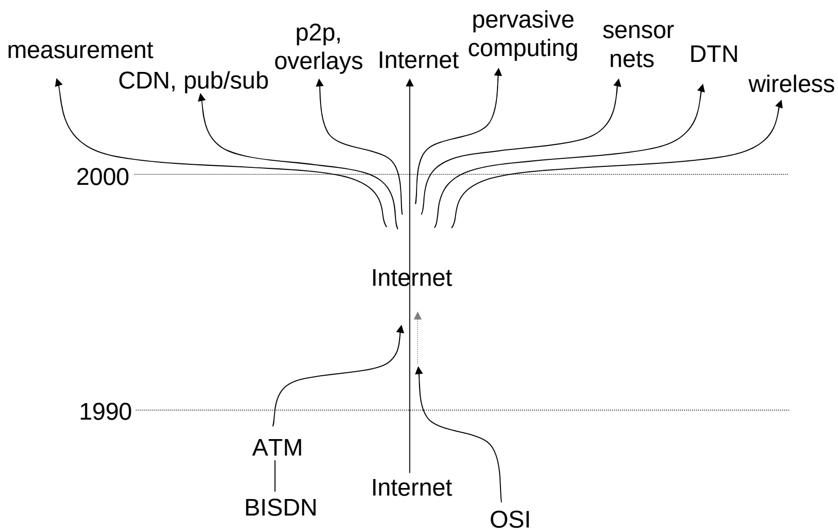

Summary: advice to students

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

Thanks!

- *acknowledgements*: D. Towsley, M. Ammar, C. Diot, H. Schulzrinne, G. Parulkar,J. Rexford, L. Zhang, several generations of students

- slides available: http://gaia.cs.umass.edu/kurose/talks/