Optimizing Quality of Information in Sensor Networks

Towards:

Control of Information Systems

Tarek Abdelzaher

University of Illinois at Urbana Champaign

Disclaimers

- I am not a control person…

Disclaimers

- I am not a control person…
- I am not a machine learning person…

Disclaimers

- I am not a control person…
- I am not a machine learning person…
- I *apply* control to solve computer-science problems such as:
  - Software performance problems
  - Real-time scheduling problems
  - Information networking problems

Disclaimers

- I am not a control person…
- I am not a machine learning person…
- I *apply* control to solve computer-science problems such as:
  - Software performance problems
  - Real-time scheduling problems
  - Information networking problems
- This talk: emerging application domains from CS

(10 year old) Example: Control of Computing Systems

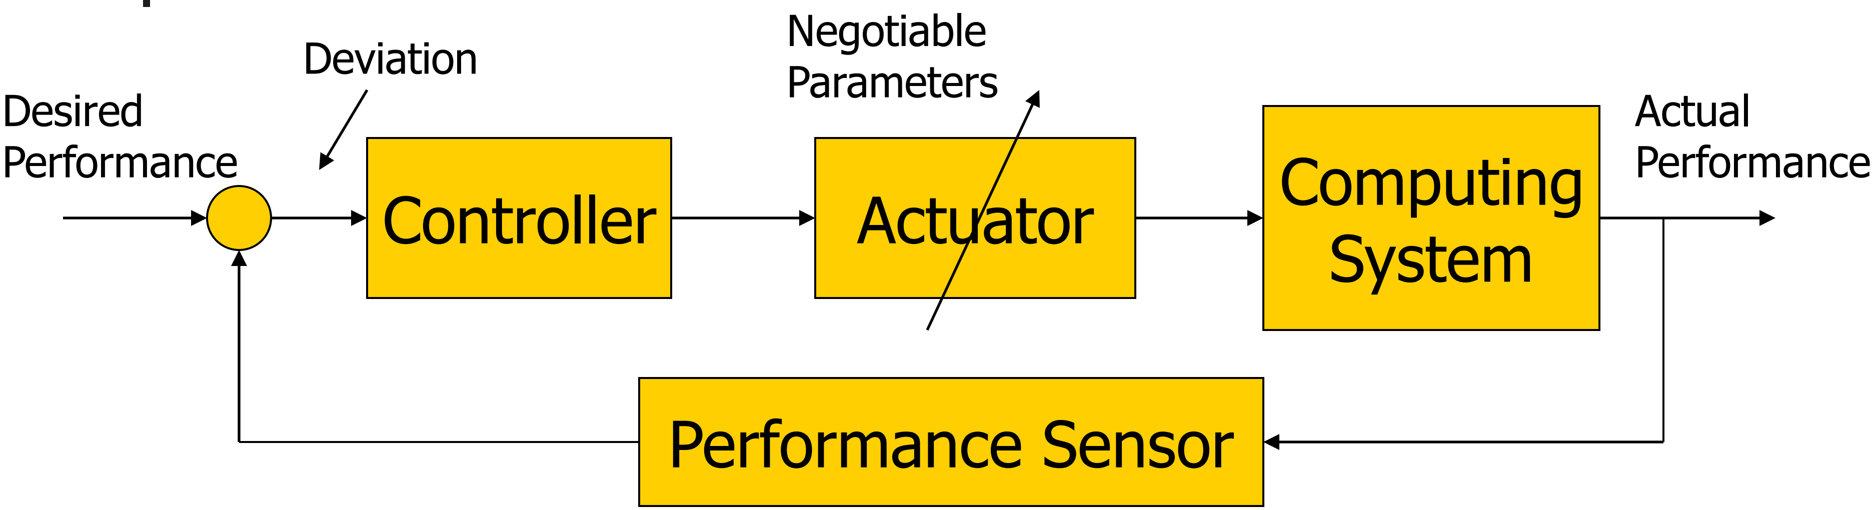

- Software is modeled by *differential equations* relating performance and resource allocation
- Sensors measure performance
- Actuators reallocate resources

Example: Web Server Modeling

Web servers can be modeled by difference equations!

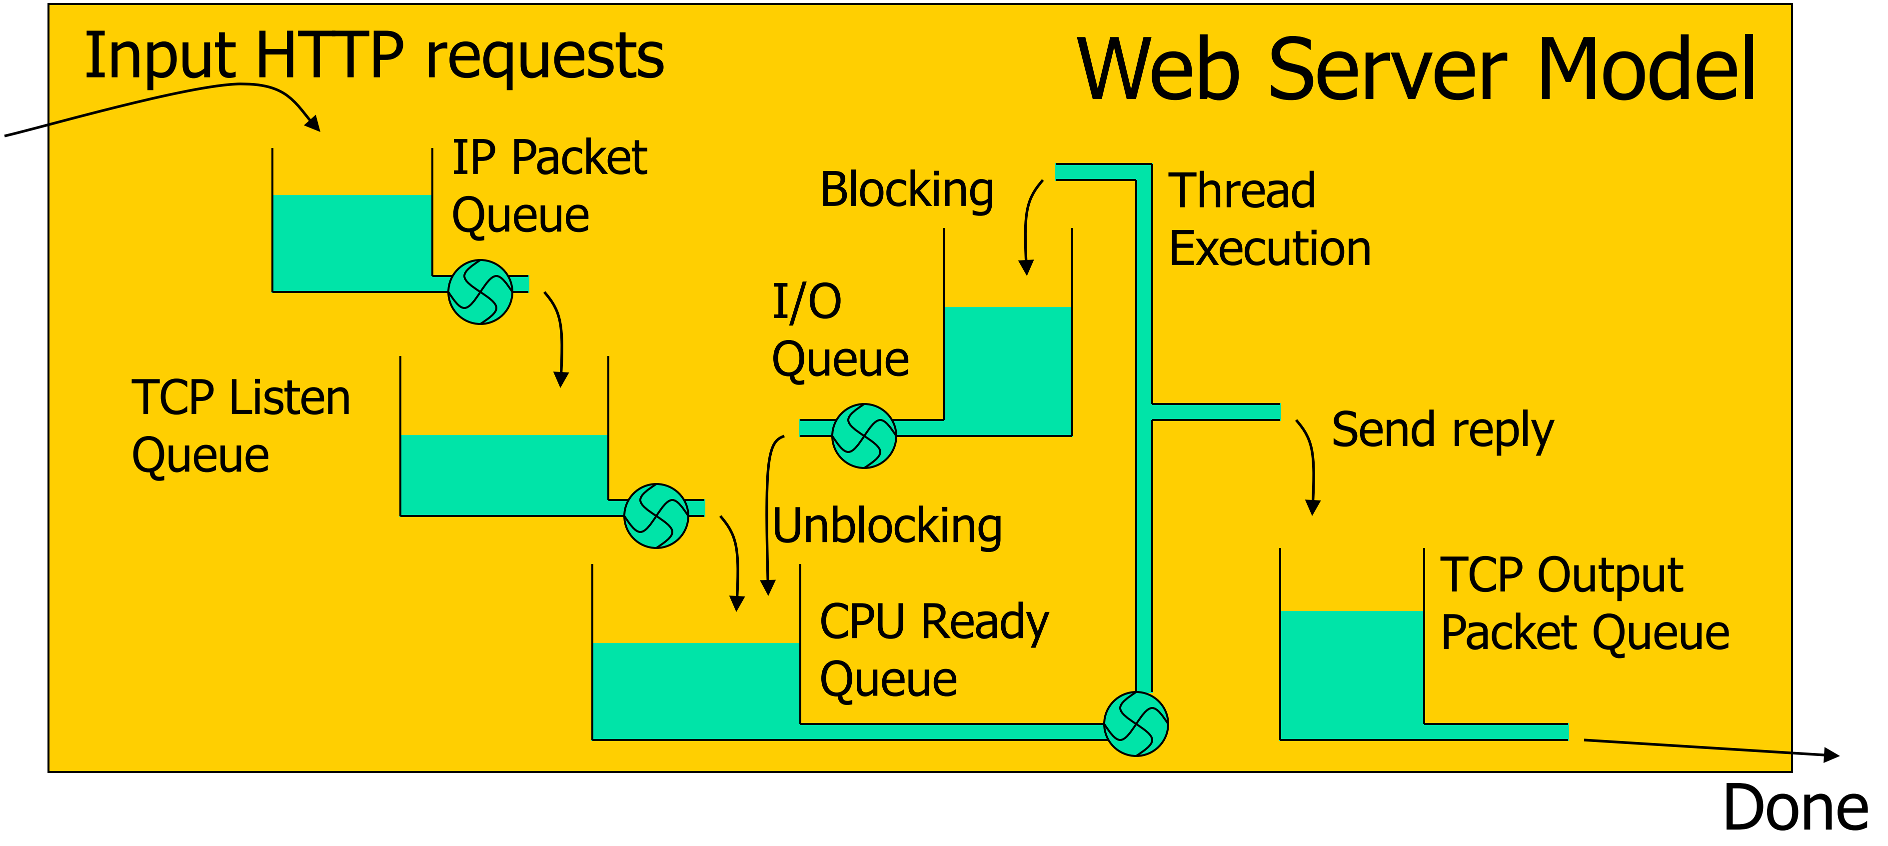

Example: Web Server Utilization Control Loop

- Load Monitor (measures utilization)
- Utilization Controller (determines degradation)
- Content Adaptor (serves appropriate content version)

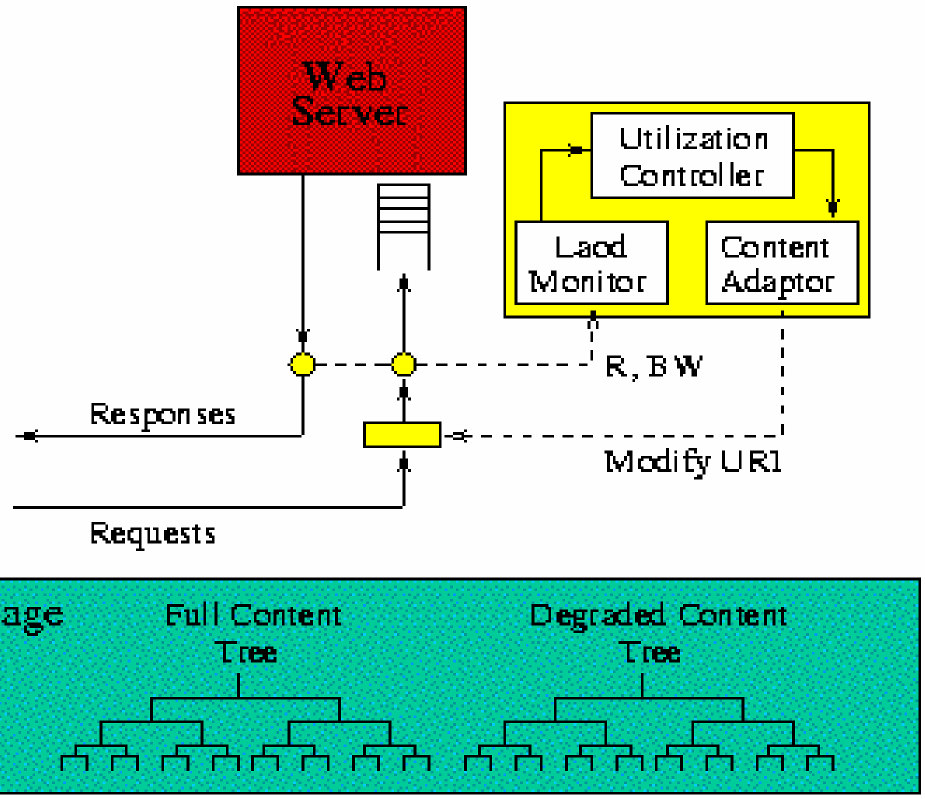

Actuation Example:
Choice of Content Version

<img src="room.gif" type=DEGRADABLE>

74KB GIF

High QoS

8.4 KB JPEG

Low QoS

Content Adaptor

- Degrades or rejects a fraction of requests depending on controller output, N.

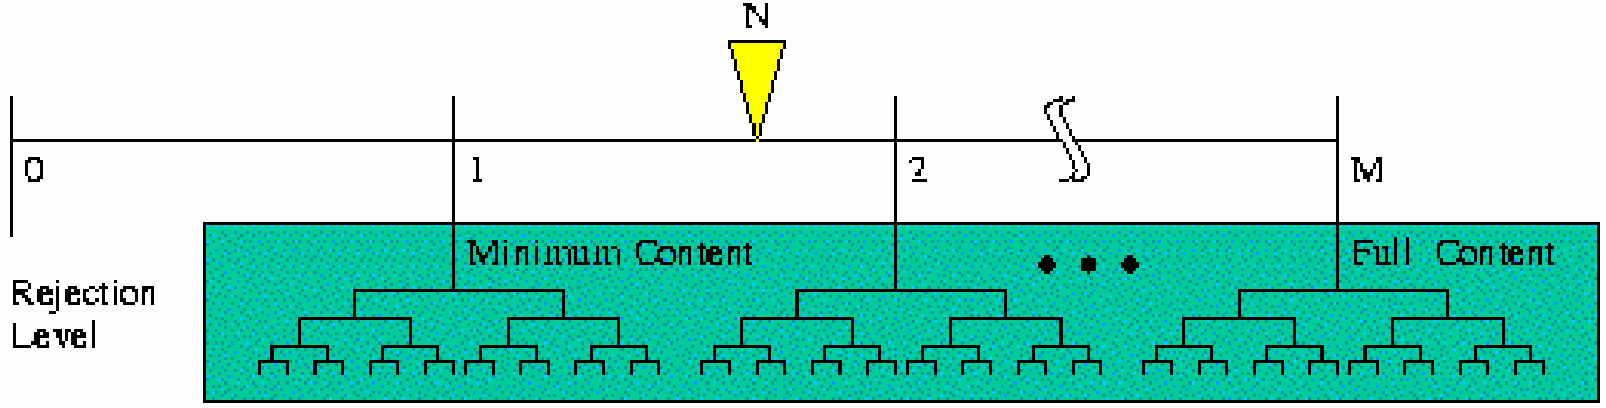

Experiment:
Efficacy of Utilization Control

- Offered Load = 300% (at t=13)
- Desired utilization = 85%

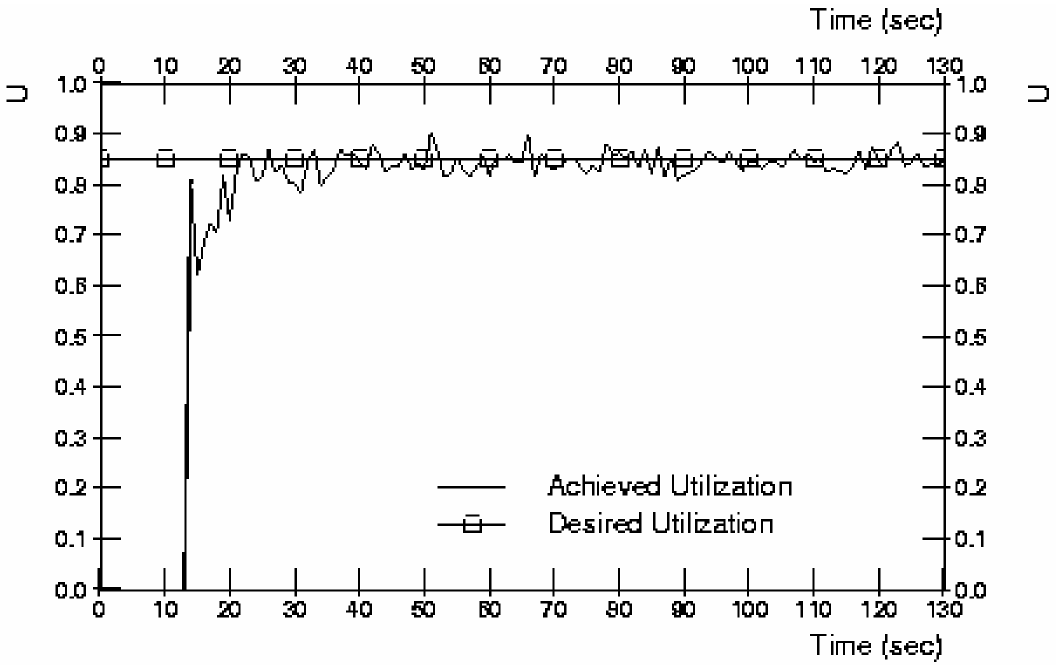

Controlling Server Utilization

A Case Study on Delay Control: Delay Control in Web Servers*

- Average delay is inversely proportional to average flow (nonlinear control)
  - From queueing theory, delay = $\frac{1}{\mu - \lambda}$

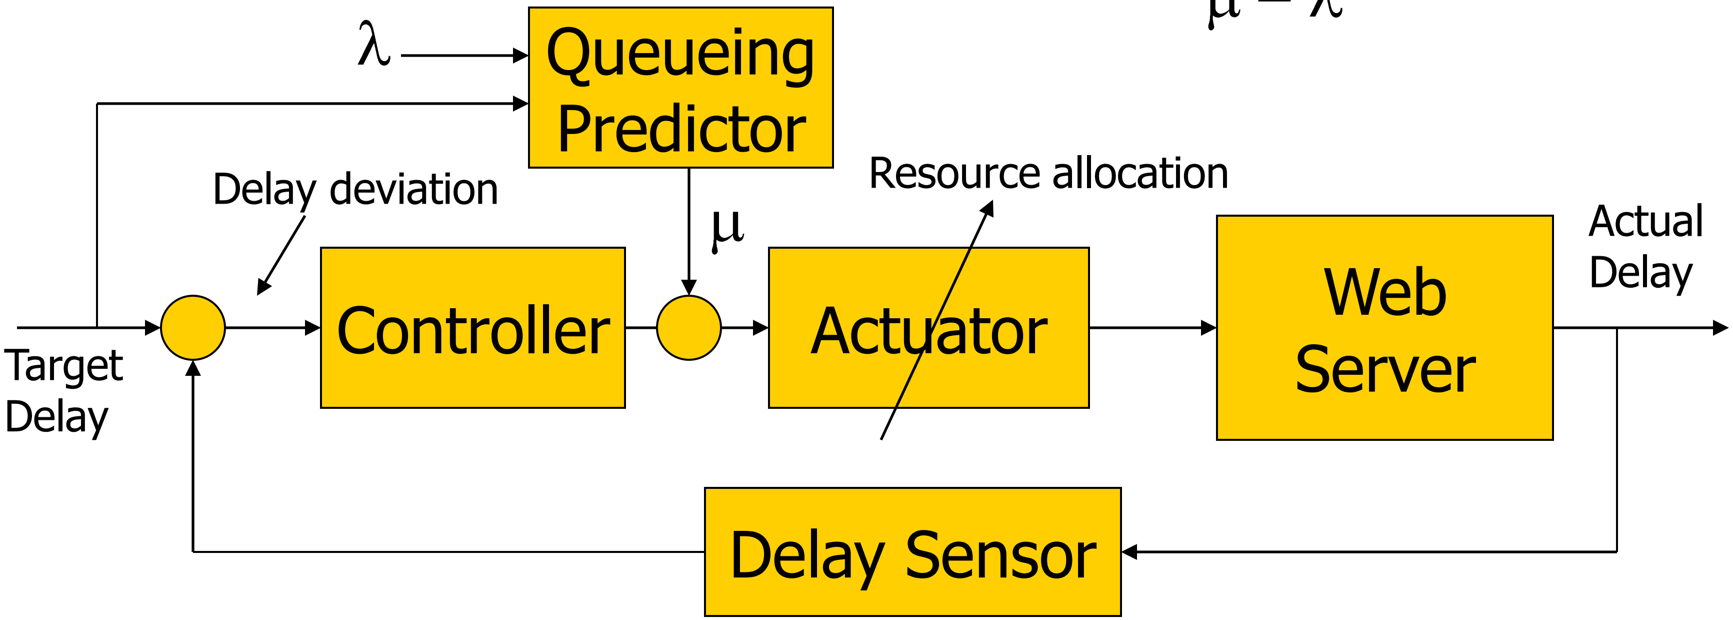

*This work is in collaboration with Lui Sha and Xue Liu from UIUC

Web Server Delay Control Experiment

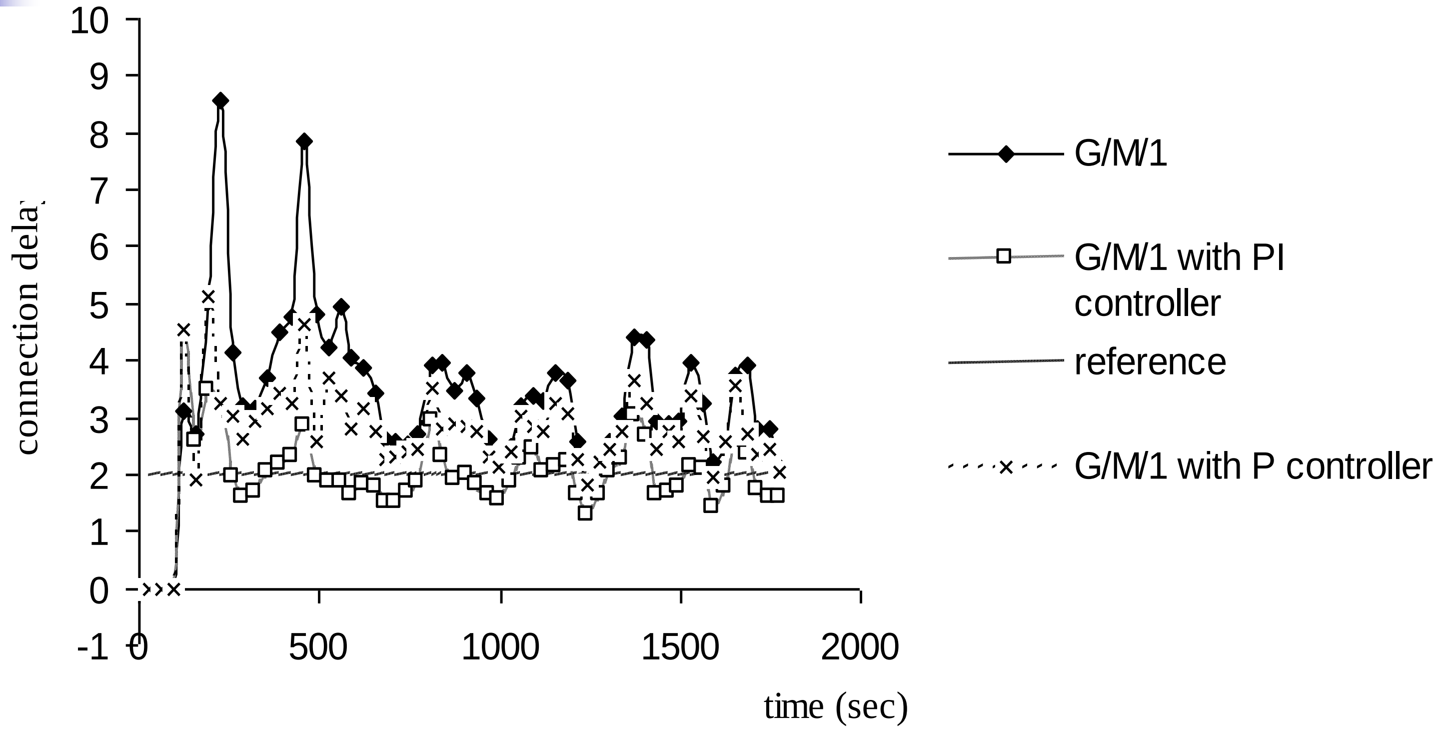

Graphs shown are a subset from Queueing Model Based Network Server Performance Control, by L. Sha, X. Liu (UIUC), Y. Lu, T. Abdelzaher (UVA)

*This work is in collaboration with Lui Sha and Xue Liu from UIUC

Adaptive Control of Web Cache Hit Ratio

- Internet performance is improved by caching frequently requested content closer to clients
- Caching needs storage space allocation to different content

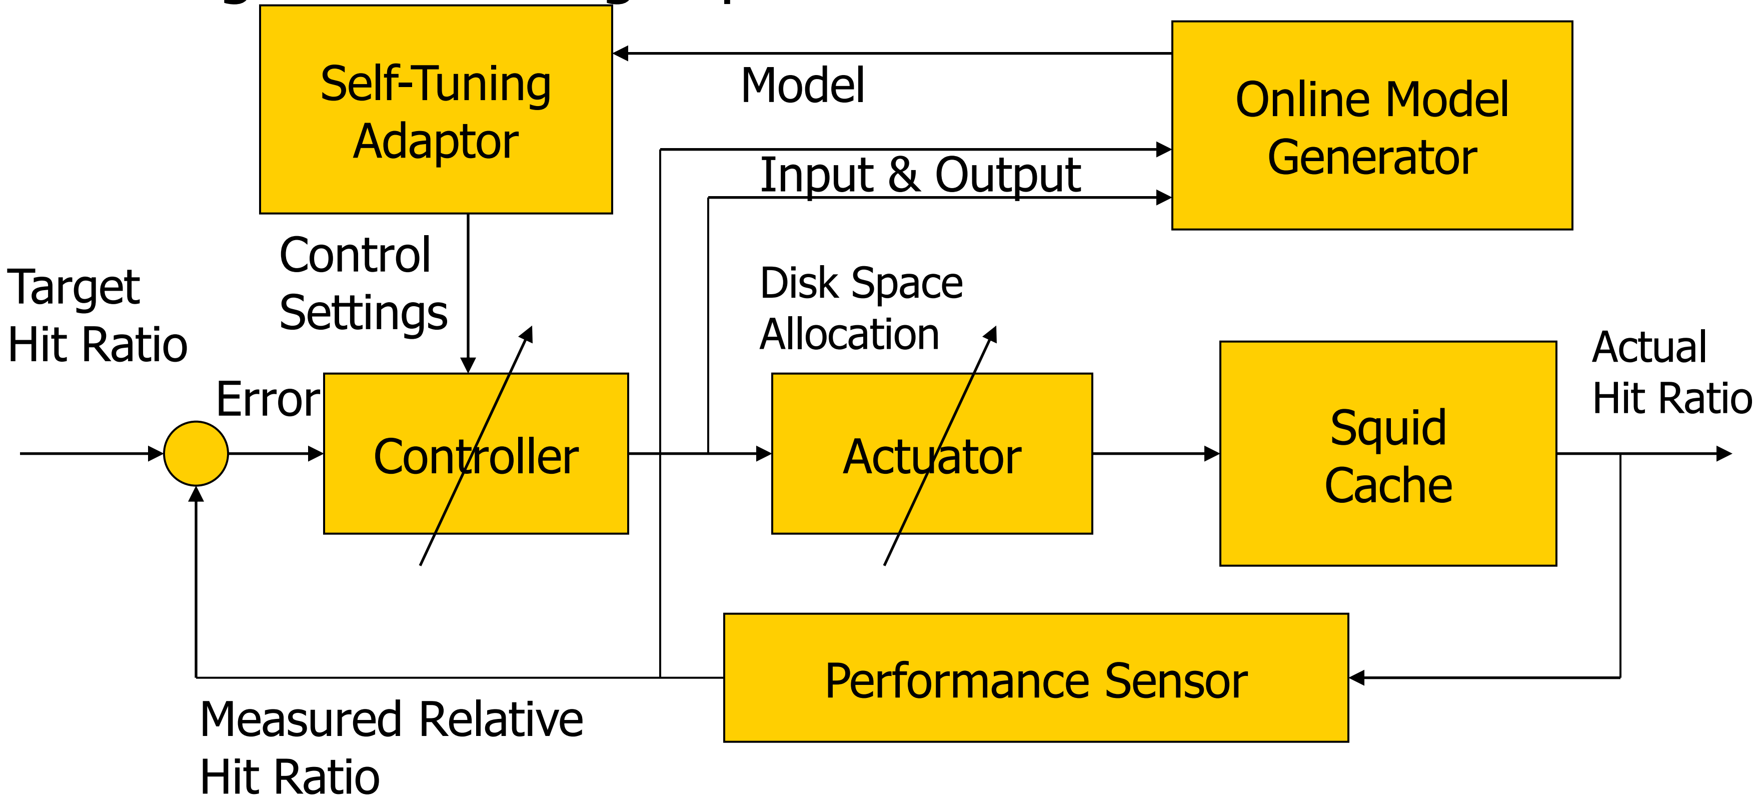

Courtesy of Ying Lu

Adaptive Controller Performance for Synthetic Log

Experiment with Adaptive Controller

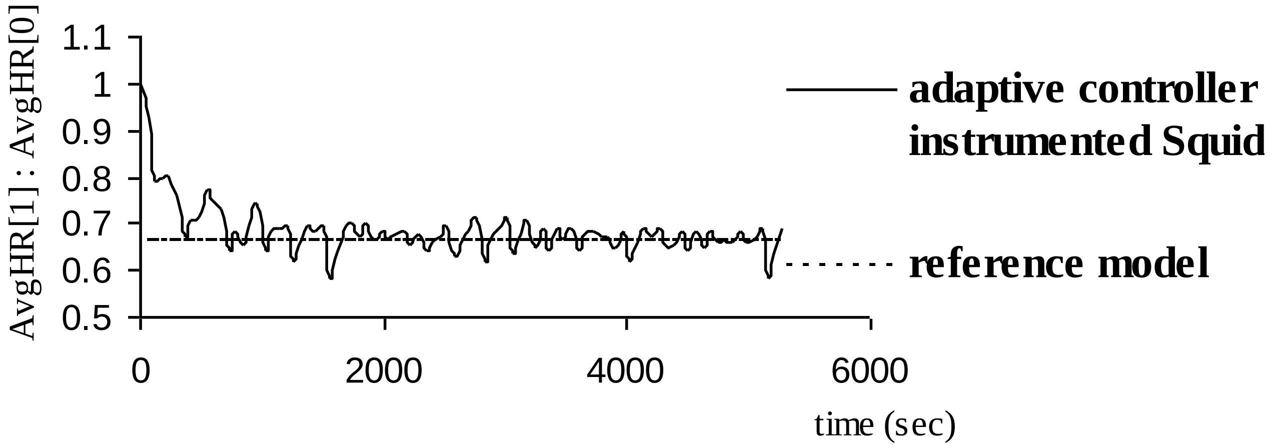

Experiment with Fixed Controller

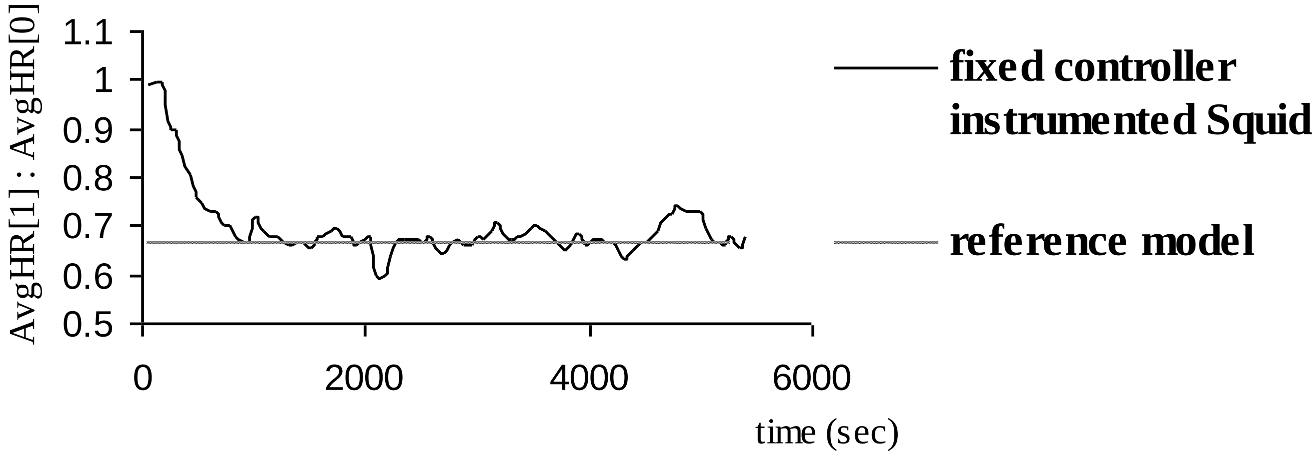

Trends in Computing

- Where is the next big application domain of control theory in computing?

Trend 1: Computing Device Proliferation (By Moore's Law)

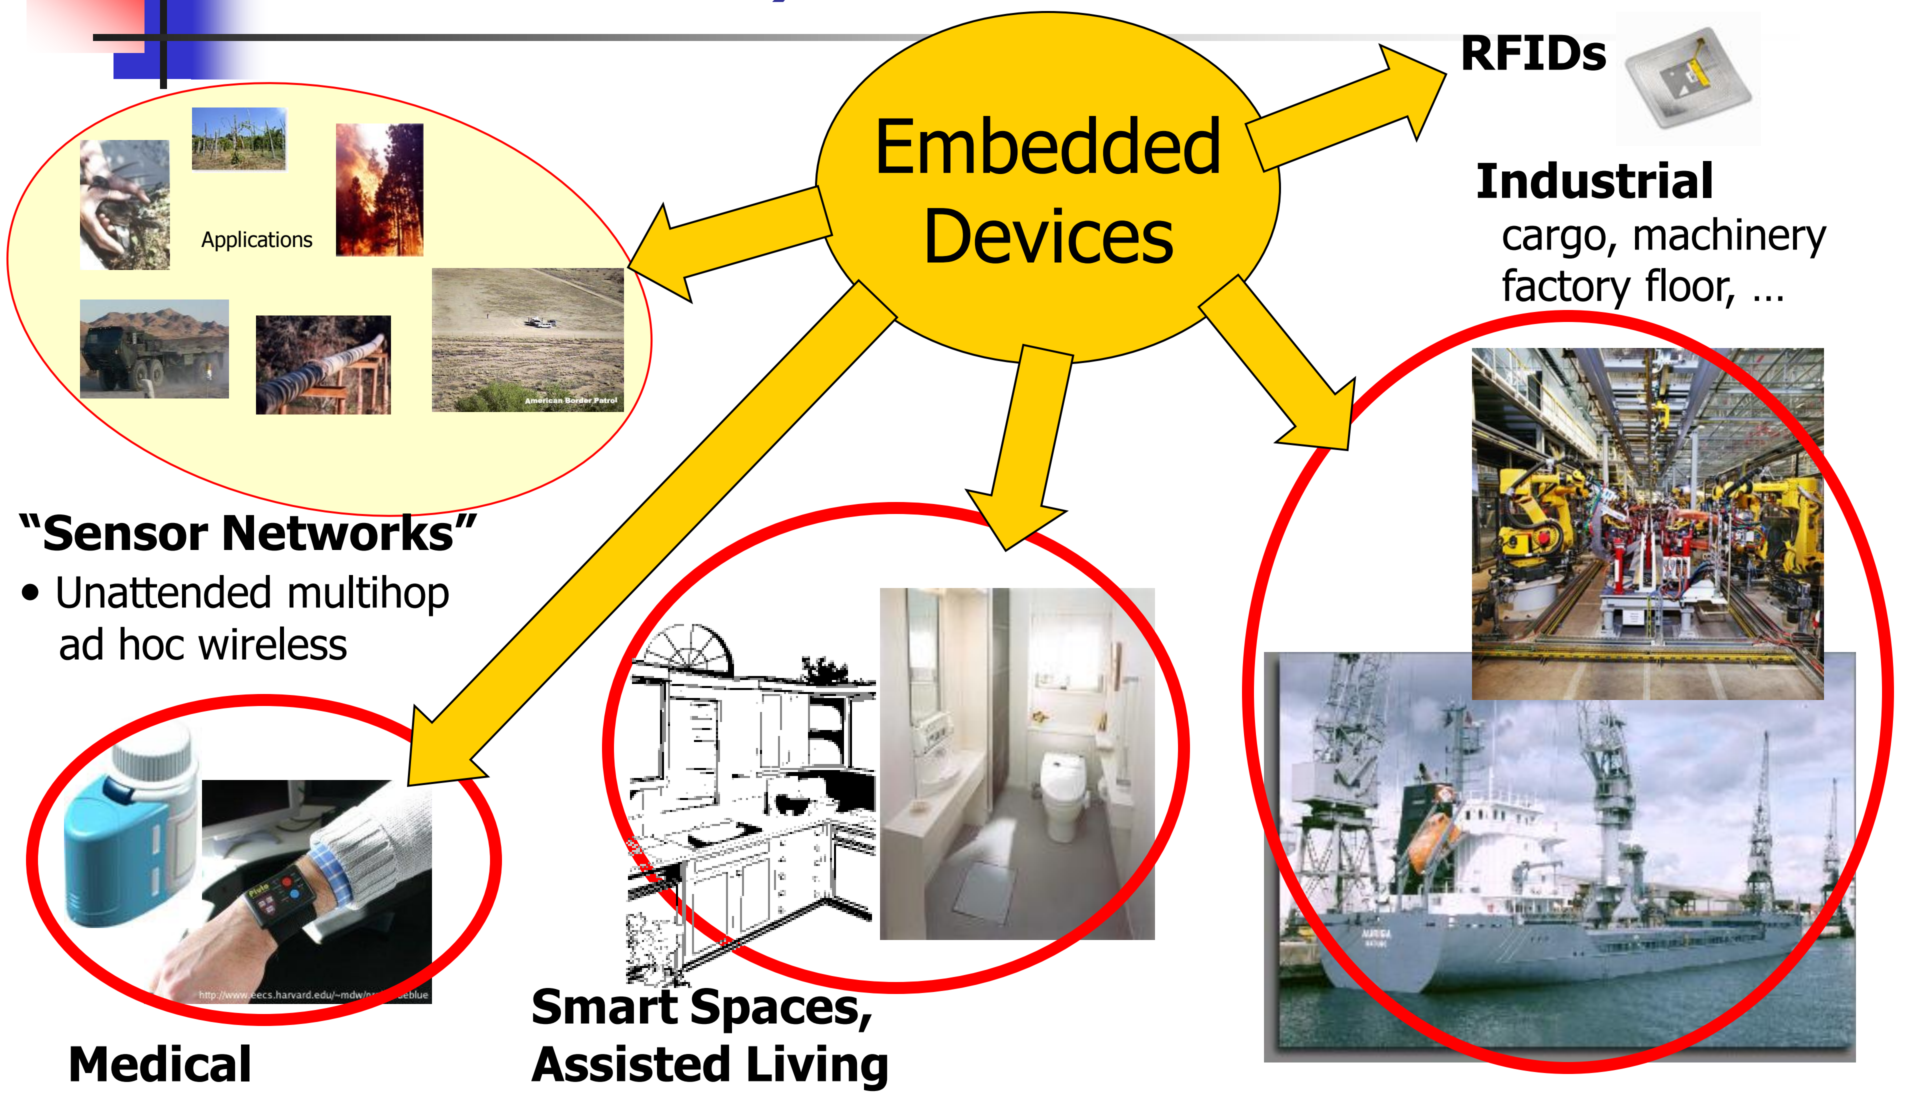

Trend 2:
Integration at Scale (Isolation has cost!)

- Low end: ubiquitous embedded devices
  - Large-scale networked embedded systems
  - Seamless integration with a physical environment
- High end: complex systems with global integration
  - Examples: Global Information Grid, Total Ship Computing Environment

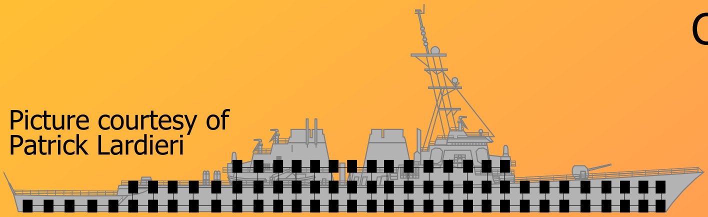

Picture courtesy of
Patrick Lardieri

Total Ship Computing Environment
(TSCE)

Global Information Grid

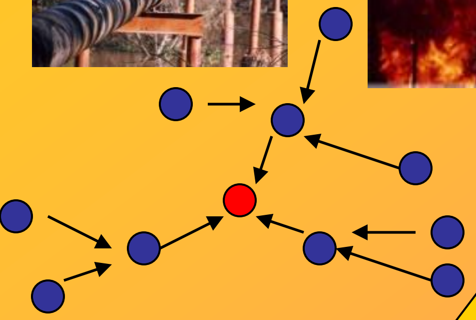

World Wide Sensor Web
(Feng Zhao)

Integration
and Scaling
Challenges

Future Combat System
(Rob Gold)

Low End — High End

Trend #3: Biological Evolution

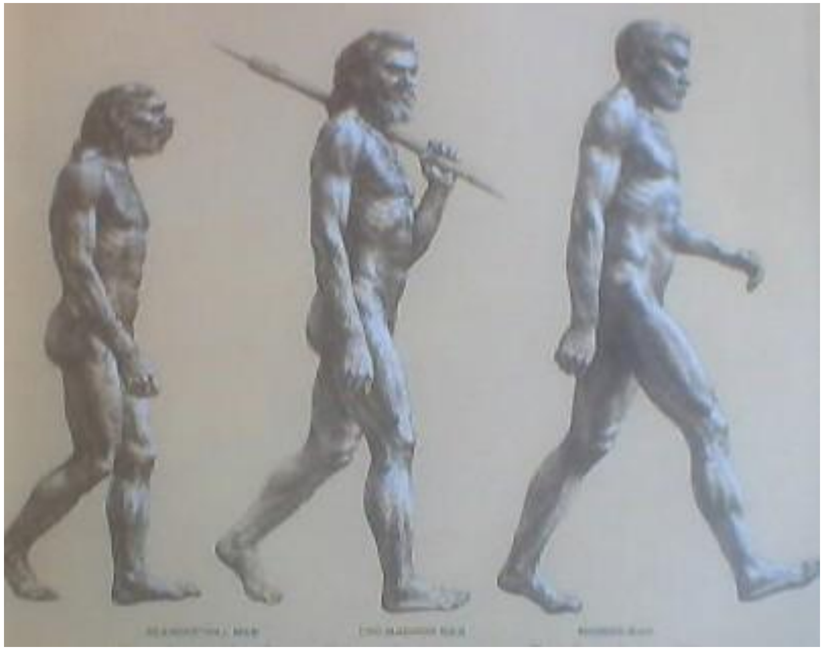

Trend #3: Biological Evolution

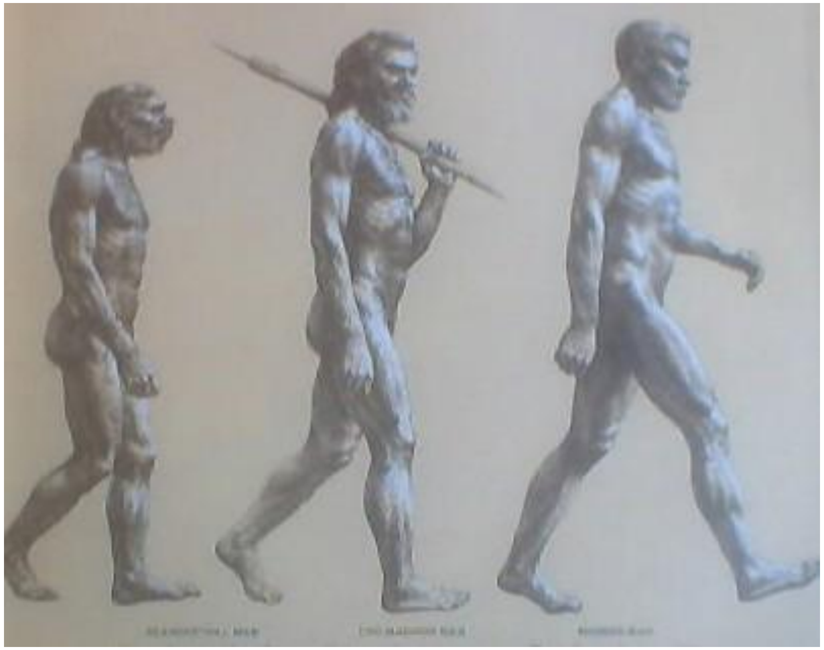

- **It's too slow!**
  - The exponential proliferation of networked embedded sensing devices (data sources), afforded by Moore's Law, is ***not*** matched by a corresponding increase in human ability to consume information!

Confluence of Trends

The Overarching Challenge

Trend2: Integration at Scale
(Isolation has cost)

Trend1: Device/Data Proliferation
(by Moore's Law)

Trend3: The capacity gap
(Humans are not getting faster)

Confluence of Trends

The Overarching Challenge

Trend2: Integration at Scale
(Isolation has cost)

An increasing category of networked systems to bridge the increasing capacity gap

Trend1: Device/Data Proliferation
(by Moore's Law)

Trend3: The capacity gap
(Humans are not getting faster)

Examples of “Bridging the Capacity Gap”

- Google
- Participatory sensing
  - Large populations of individuals measure phenomena of joint interest (e.g., traffic using GPS) to inform their decisions (e.g., avoid congestion).
- Military intelligence
  - Where to place which sensors (data sources) to collect the most relevant information to particular predictions?

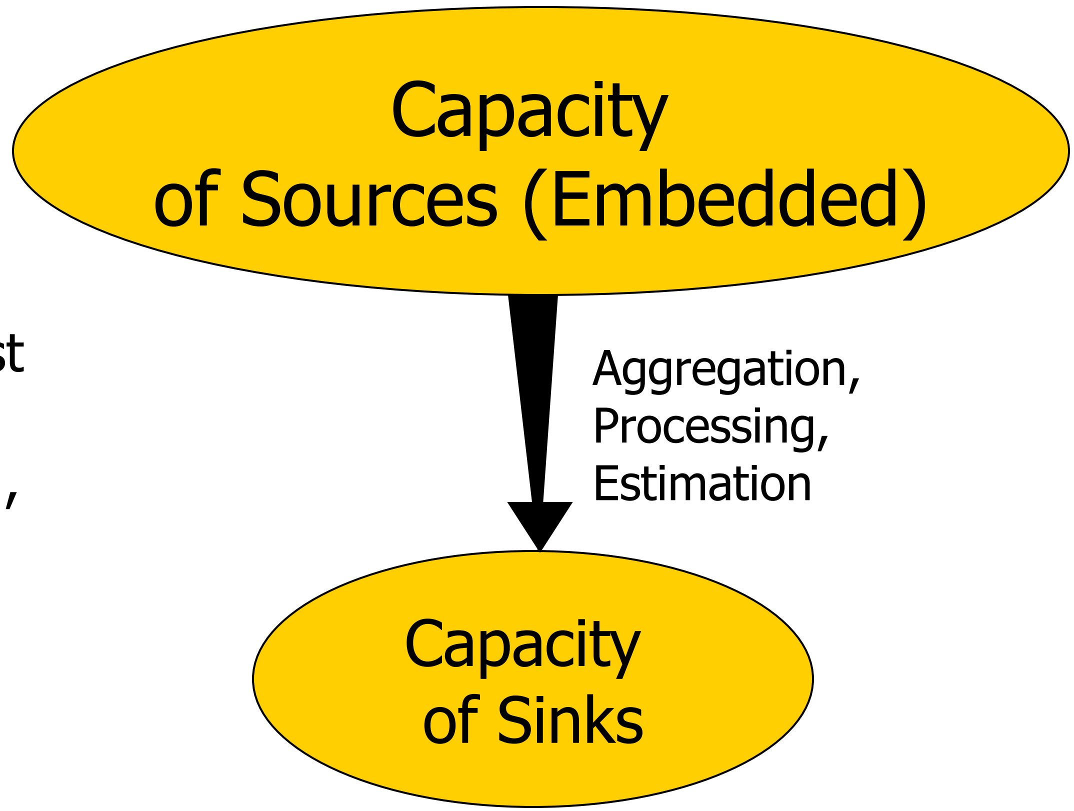

Examples of "Bridging the Capacity Gap"

- Google
- Participatory sensing
  - Large populations of individuals measure phenomena of joint interest (e.g., traffic using GPS) to inform their decisions (e.g., avoid congestion).
- Military intelligence
  - Where to place which sensors (data sources) to collect the most relevant information to particular predictions?

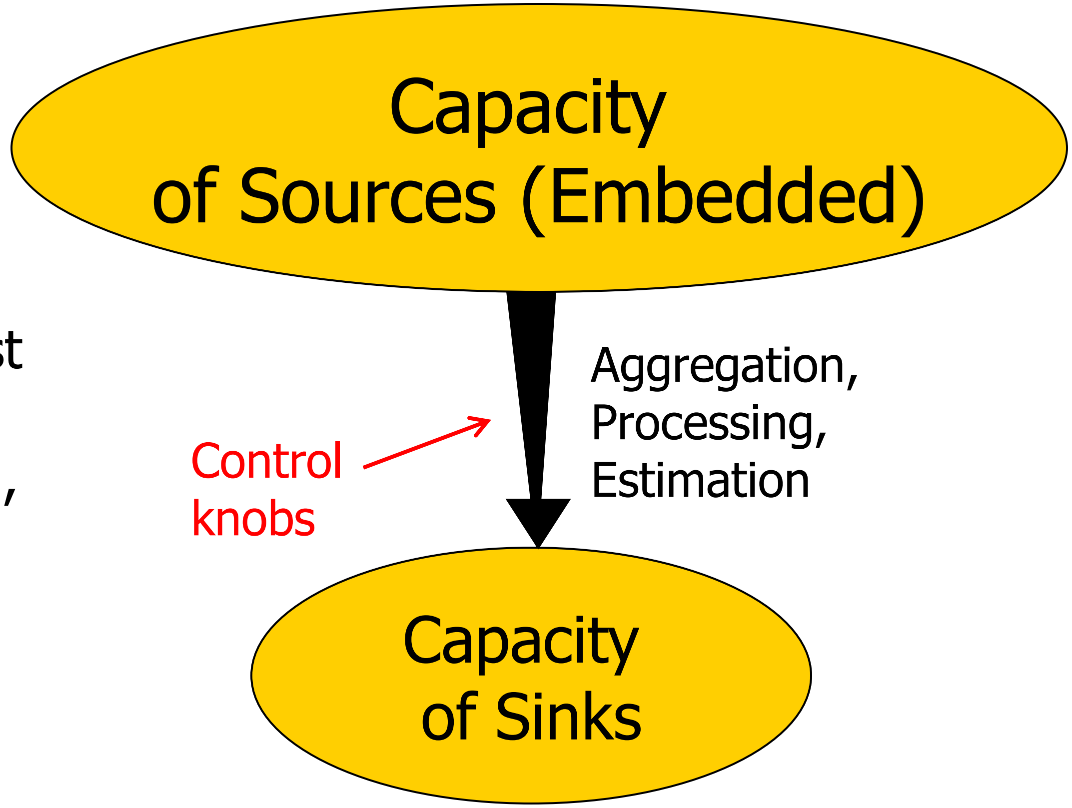

A Control Problem

A Control Problem

High quality information

Data fusion (data filtering and aggregation control knobs)

- Information Distillation

Raw data

ARL: A UIUC Center on Information Networks

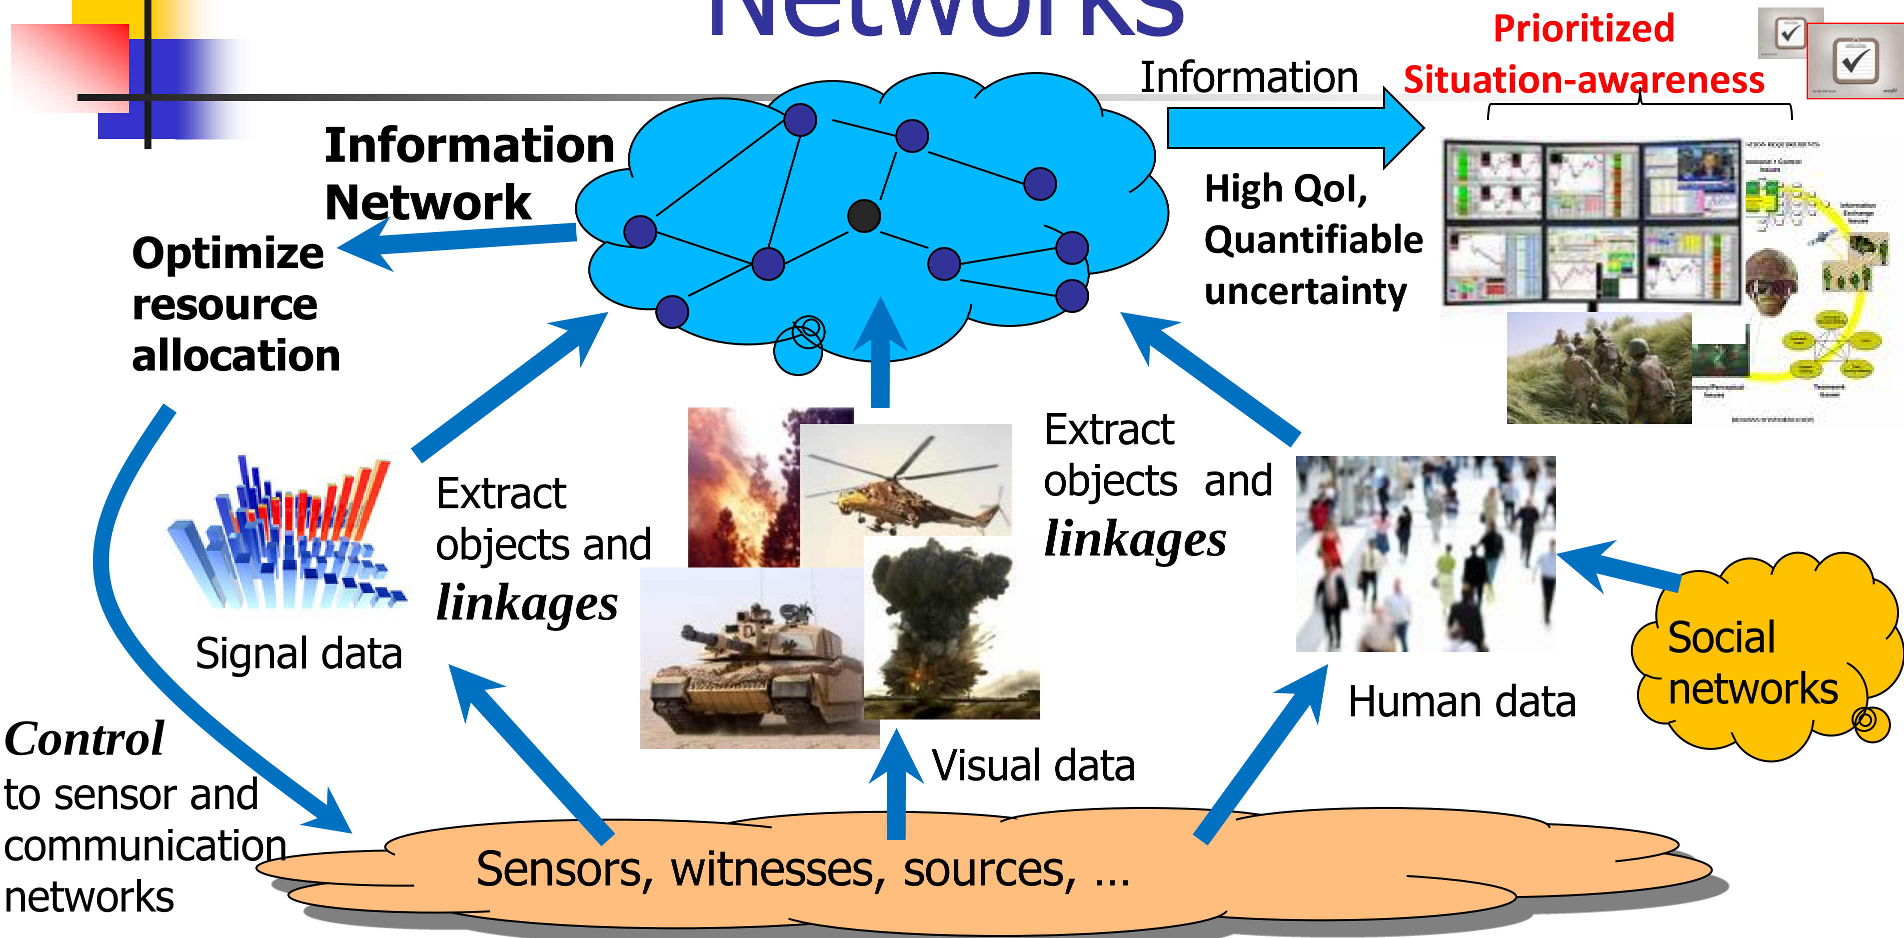

- **Two general themes**
  - Optimizing the quality of information
  - Quantifying uncertainty as a first class attribute

An Example "Deliverable": Rethink Communication Networks

- Problem:
  - Current networks optimize *communication objectives* (delay, throughput, etc)
  - But, they are being used increasingly as *information sources*.
- Challenge: Redesign networks around the notion of information quality optimization
  - Information quality is a *physical* concept (akin to modeling accuracy)
- A *cyber-physical systems* problem: Designing cyber-components with an understanding of physical models (and vice-versa)

A Simple Example: Sensor Networks that Measure Physical Phenomena

- Sensors (e.g., temperature) periodically measure temperature and send it to a base-station.
- Data travels through multiple hops
- Congestion occurs when a node cannot transmit (due to energy or bandwidth constraints) the data it generates/forwards
  - How should the sensor network react?

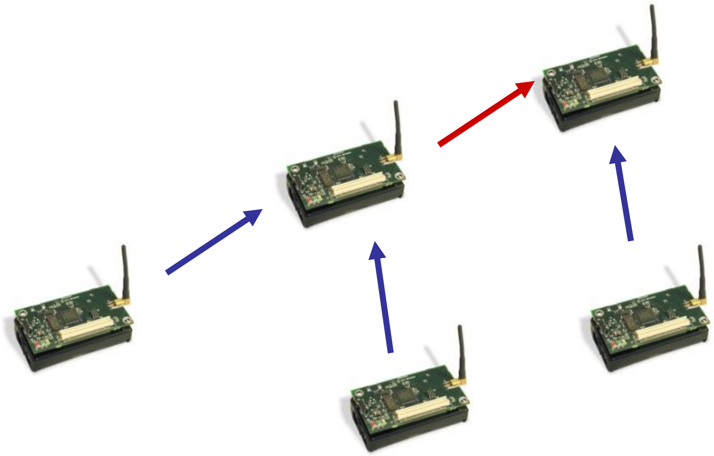

A Simple Example: Sensor Networks that Measure Physical Phenomena

- How should the sensor network react?
  - Classical response: Do backpressure on sources to reduce network load
  - Desired response: Throttle flows that do not contribute substantially to the quality of information at the receiver.

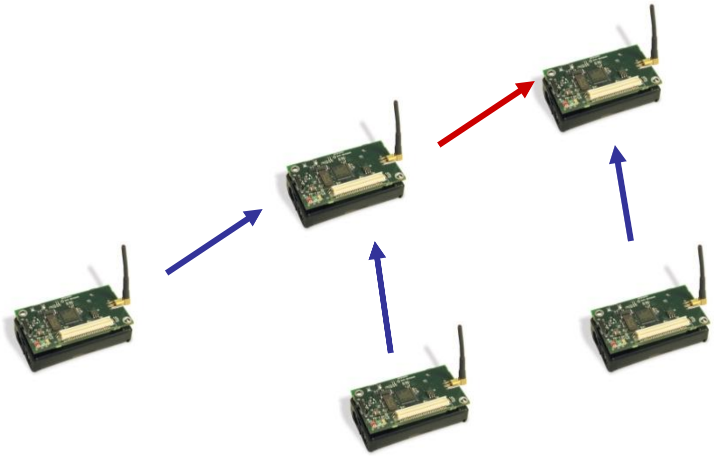

A “Cyber-physical” Approach

- Nodes indirectly affect the flow rate by specifying their *allowable estimation error* based on observed congestion
  - Adjust the allowable estimation error experienced on a bottleneck link and inform the upstream nodes (control knob to exert backpressure).
  - Upstream nodes summarize their data locally as much as possible while meeting the set value of allowable estimation error

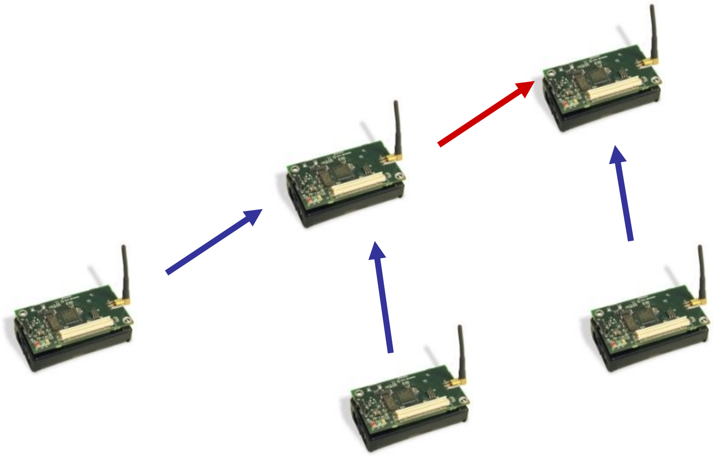

A Cyber-physical Approach

- Nodes indirectly affect the flow rate by specifying their *allowable estimation error* based on observed congestion
  - Adjust the allowable estimation error experienced on a bottleneck link and inform the upstream nodes (control knob to exert backpressure).
  - Upstream nodes summarize their data locally as much as possible while meeting the set value of allowable estimation error

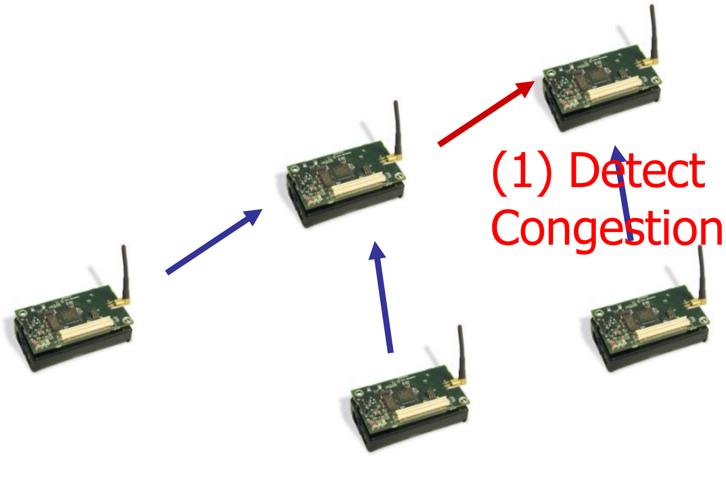

A Cyber-physical Approach

- Nodes indirectly affect the flow rate by specifying their *allowable estimation error* based on observed congestion
  - Adjust the allowable estimation error experienced on a bottleneck link and inform the upstream nodes (control knob to exert backpressure).
  - Upstream nodes summarize their data locally as much as possible while meeting the set value of allowable estimation error

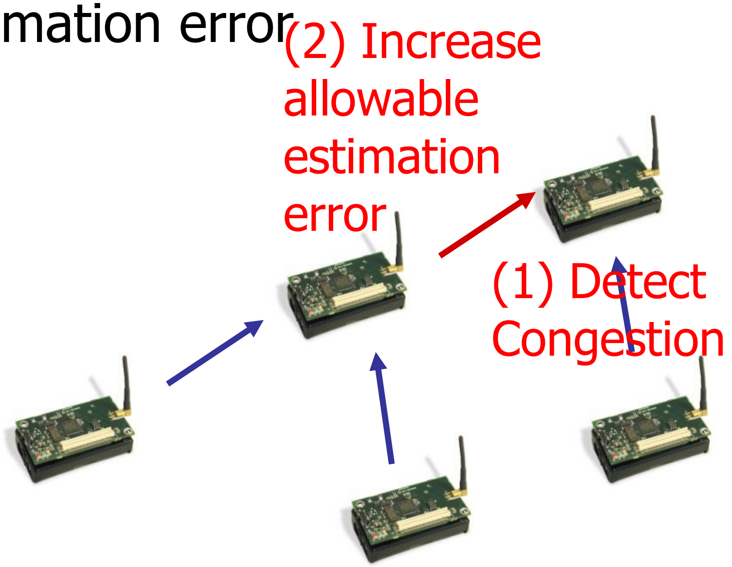

A Cyber-physical Approach

- Nodes indirectly affect the flow rate by specifying their *allowable estimation error* based on observed congestion
  - Adjust the allowable estimation error experienced on a bottleneck link and inform the upstream nodes (control knob to exert backpressure).
  - Upstream nodes summarize their data locally as much as possible while meeting the set value of allowable estimation error

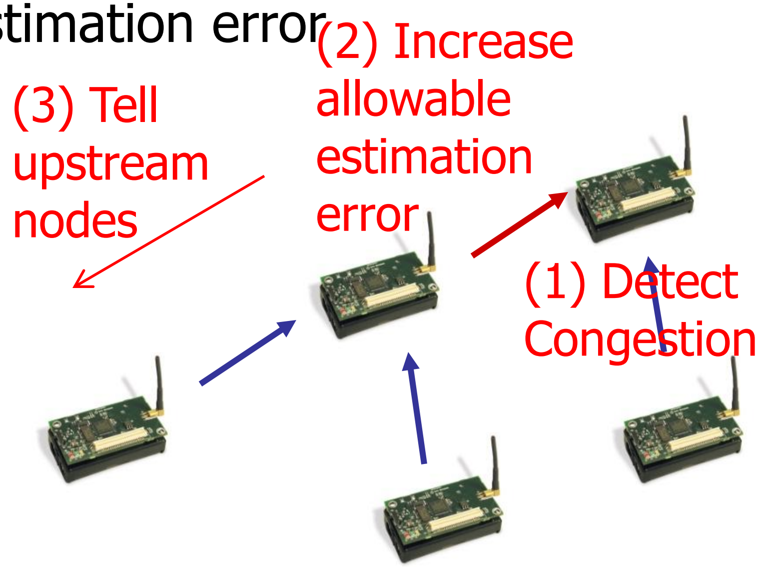

A Cyber-physical Approach

- Nodes indirectly affect the flow rate by specifying their *allowable estimation error* based on observed congestion
  - Adjust the allowable estimation error experienced on a bottleneck link and inform the upstream nodes (control knob to exert backpressure).
  - Upstream nodes summarize their data locally as much as possible while meeting the set value of allowable estimation error

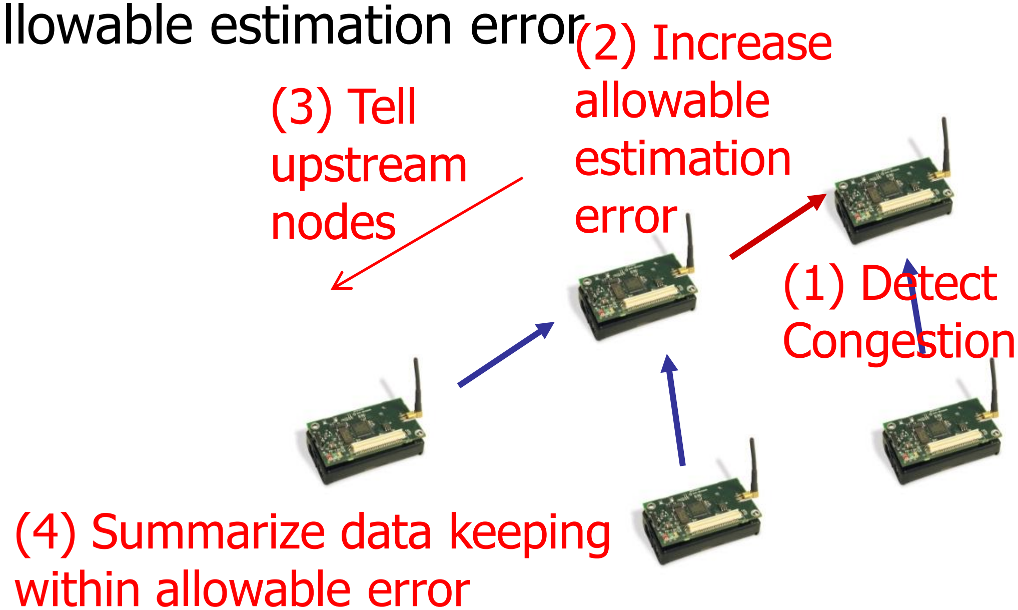

A Cyber-physical Approach

- Nodes indirectly affect the flow rate by specifying their *allowable estimation error* based on observed congestion
  - Adjust the allowable estimation error experienced on a bottleneck link and inform the upstream nodes (control knob to exert backpressure).
  - Upstream nodes summarize their data locally as much as possible while meeting the set value of allowable estimation error

Note: Relation between data summarization and model accuracy itself depends on model (a learning problem)

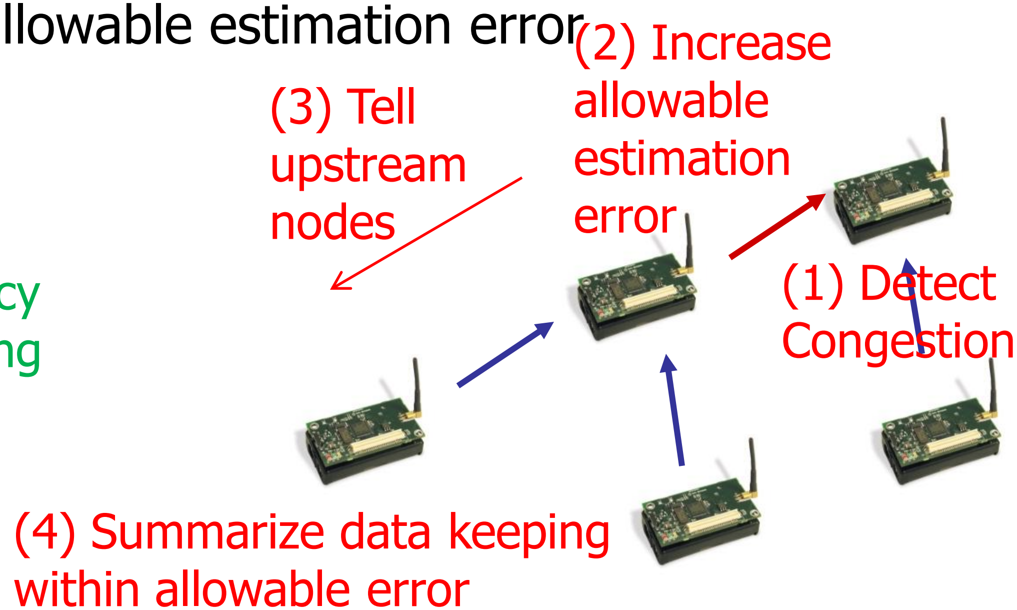

A Cyber-physical Approach

- Nodes indirectly affect the flow rate by specifying their *allowable estimation error* based on observed congestion
  - Adjust the allowable estimation error experienced on a bottleneck link and inform the upstream nodes (control knob to exert backpressure).
  - Upstream nodes summarize their data locally as much as possible while meeting the set value of allowable estimation error

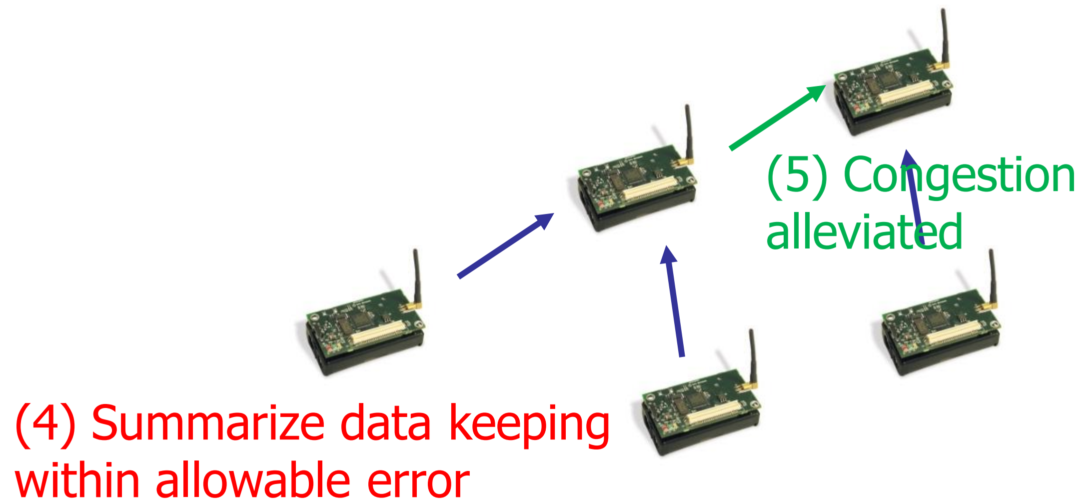

Experimental Testbed

20 nodes with sensing and communication capabilities forming multihop routes to basestation

Comparison:

- Previous congestion control protocols (ESRT, RMST)
- Unreliable transport
- CPS Approach

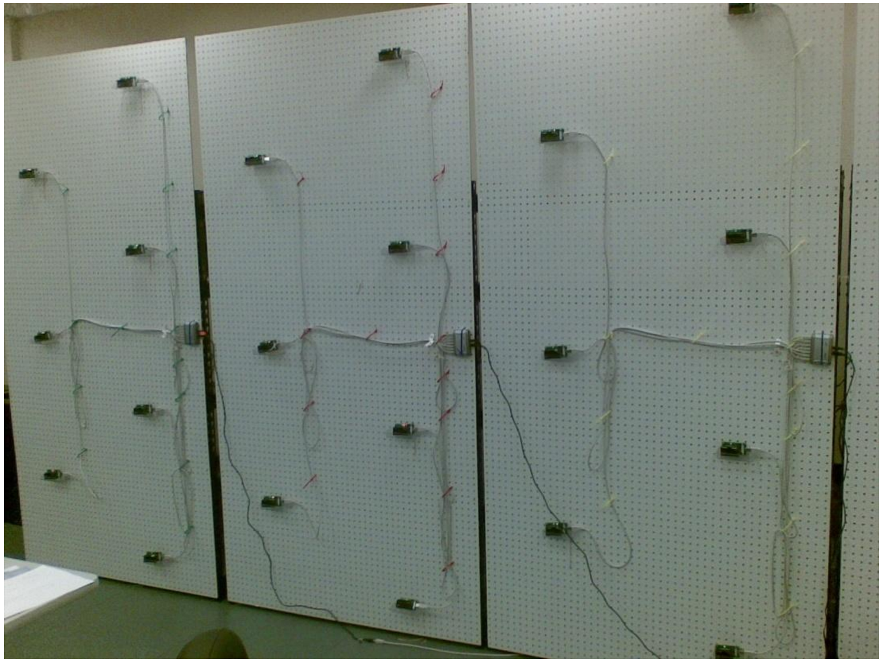

Evaluation: Error and Energy

- Our approach results in (up to) an order of magnitude less estimation error than a traditional congestion control

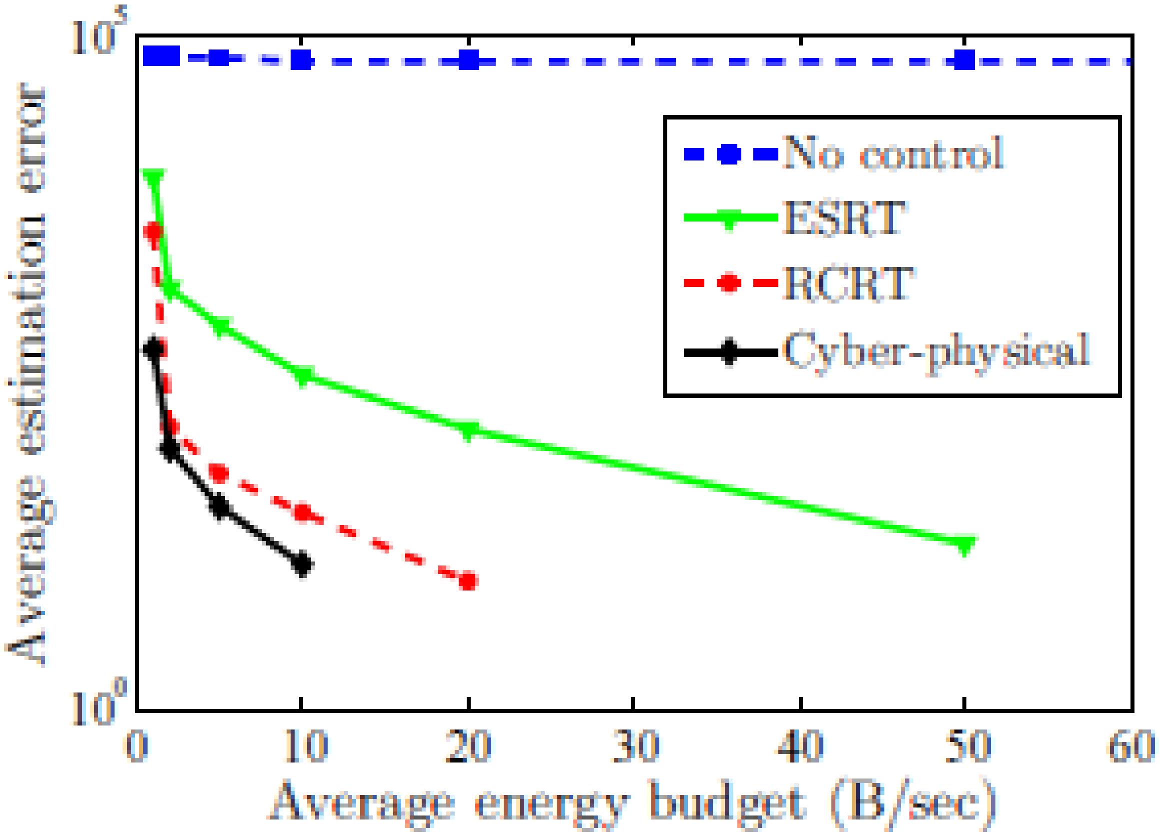

Summary

- Embedded devices generate "data overload"
- Future distributed computing services will be increasingly geared towards reducing raw data to actionable information
- The information processing system itself can be thought of as a controlled process in a loop that aims to achieve informational objectives
  - Understanding loop properties is a fertile area for collaboration between CS and Control researchers

Future Research Questions

Quoting the US Airforce Initiative: *Joint Control of Physical and Information State*

- Can we represent physical states and information states in a joint state space?
- Can we define control policies that achieve informational objectives?
- Can we develop theory that allows making trade-offs between quality of information and physical "control effort" (cost of information acquisition)?