Real-Time Sensor Networks: Paradigms, Challenges, and Open Issues

Tarek Abdelzaher
Department of Computer Science
University of Virginia

Where is Computer Science Research Going?

The beginning:

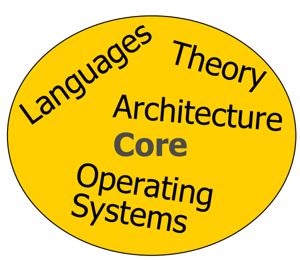

Where is Computer Science Research Going?

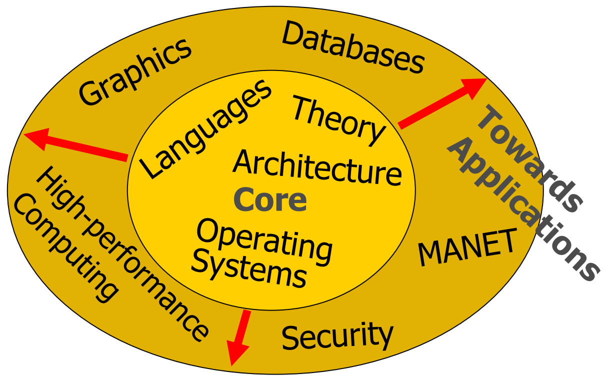

Where is Computer Science Research Going?

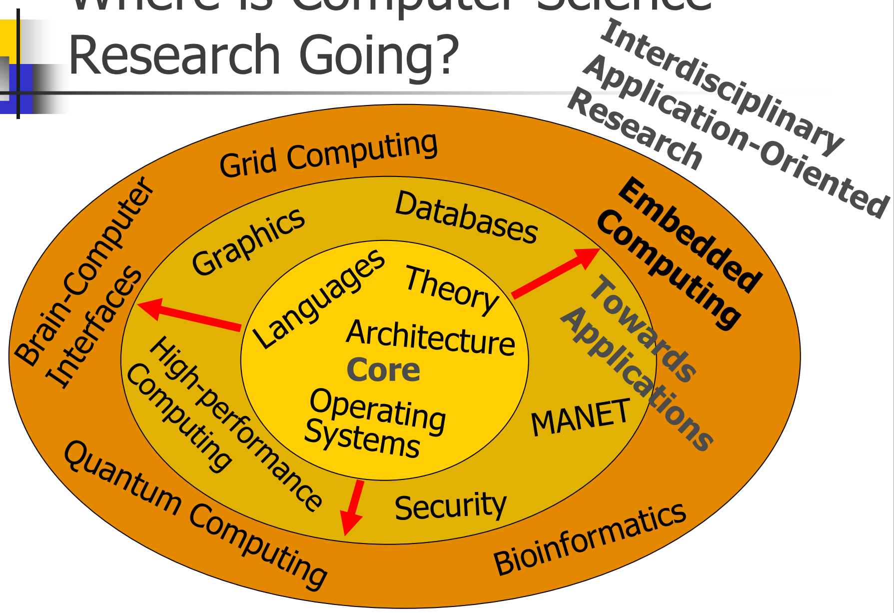

Real-Time and Embedded Systems

The Next Frontier

Why embedded computing?

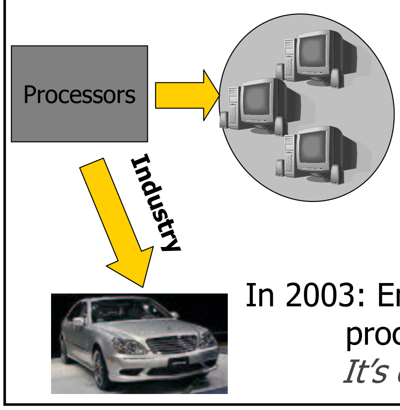

In 2003: Embedded processors are 98% of all processors manufactured

It's only the beginning…

Real-Time and Embedded Systems

The Next Frontier

Trend 1:
- Invisible (embedded) computing, implicit interfaces (users need only 1 mobile device – rest should be non-intrusive)
- Context-aware computing (new sensors, new effectors)
- Ubiquitous – instrument what we use most (attire, personal effects, …)

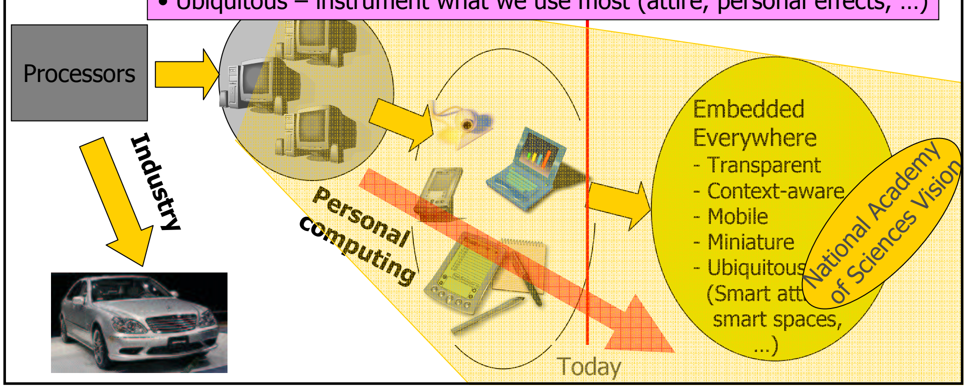

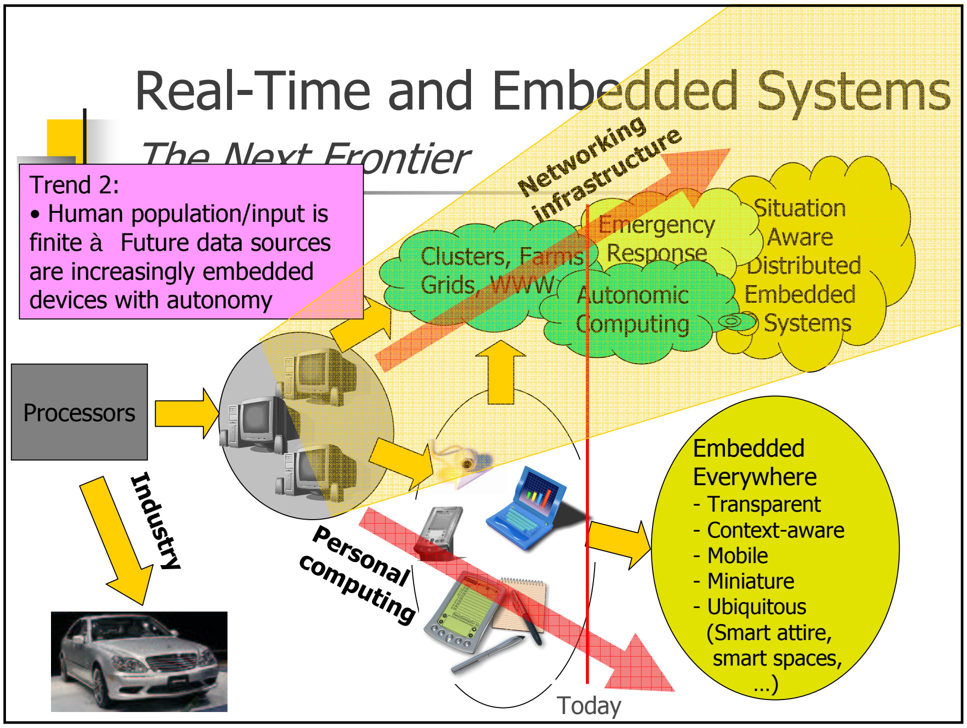
Real-Time and Embedded Systems
The Next Frontier
Trend 2:
• Human population/input is finite à Future data sources are increasingly embedded devices with autonomy
Networking infrastructure
Clusters, Farms Grids, WWW
Emergency Response
Autonomic Computing
Situation Aware Distributed Embedded Systems
Processors
Industry
Personal computing
Embedded Everywhere
- Transparent
- Context-aware
- Mobile
- Miniature
- Ubiquitous
(Smart attire, smart spaces, ...)
Today

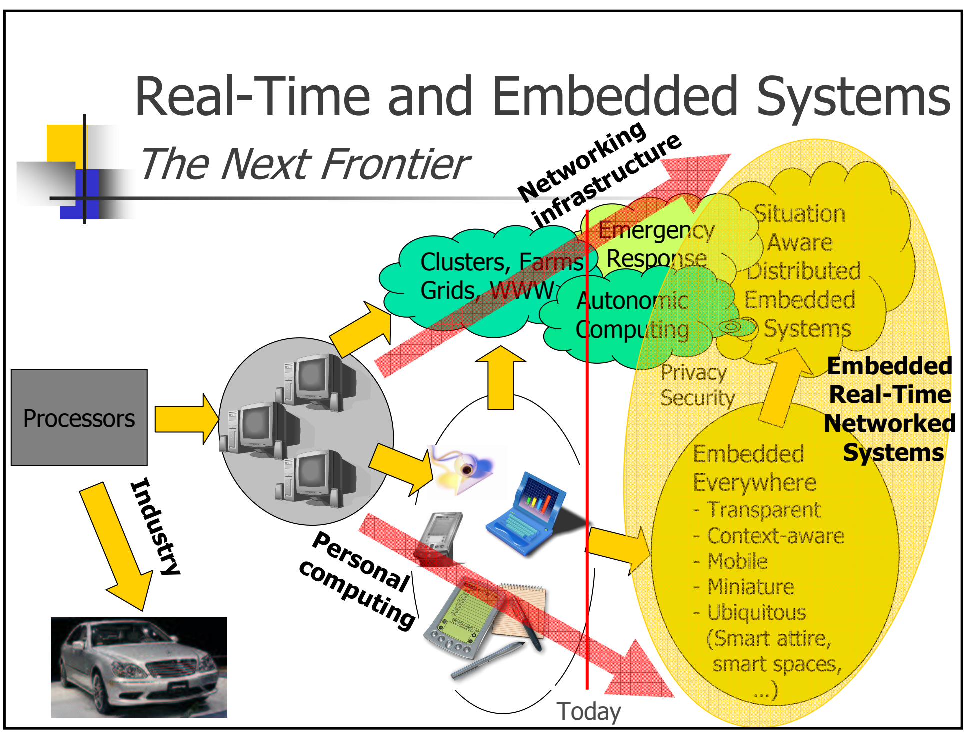
Real-Time and Embedded Systems
The Next Frontier
Networking infrastructure
Clusters, Farms Grids, WWW
Emergency Response
Autonomic Computing
Situation Aware Distributed Embedded Systems
Privacy Security
Embedded Real-Time Networked Systems
Processors
Industry
Personal computing
Embedded Everywhere
- Transparent
- Context-aware
- Mobile
- Miniature
- Ubiquitous
(Smart attire, smart spaces, ...)
Today

Expanding Interest in Embedded Computing

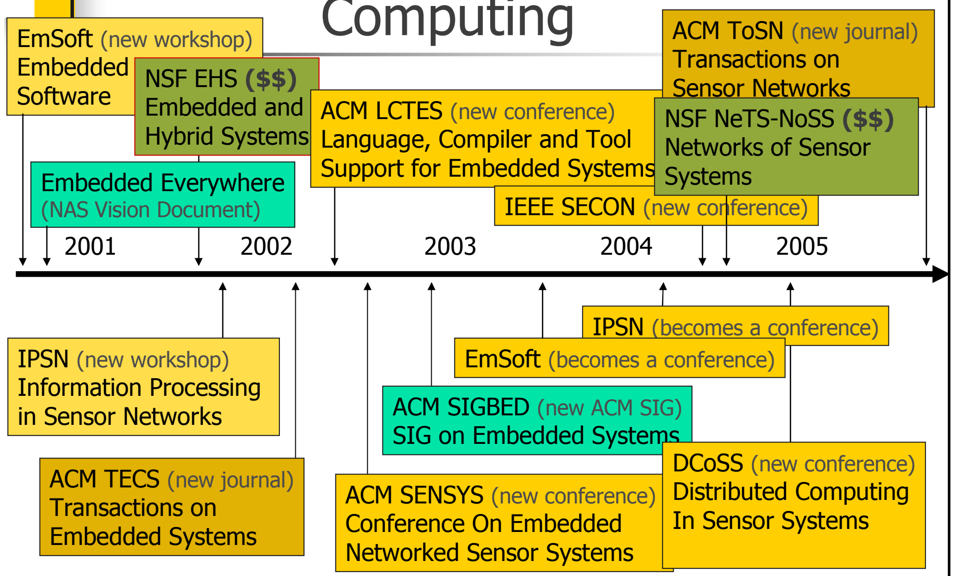

Overarching Challenge: A Science of *Design* and *Performance Analysis* for Embedded Networks

- An NSF initiative: Science of Design
  - "How can we explain, ***predict and control emergent properties*** of software-intensive systems?"
  - "To what extent can we systematize guidance that leads to systems that ***satisfy requirements***?"
  - "How can we develop theories that rely on aggregate ***reasoning about overall behavior*** rather than exact reasoning about all the details?"
- A sensor networks perspective
  - A science of design ***within constraints of time, space, and laws-of-physics***

Sensor Network Design

Three Software Challenges

Vision: Ubiquitous Sensor Networks

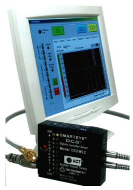

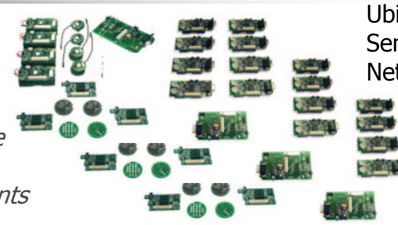

Trend: *Moving from single device to networks of thousands of small components*

Software Challenges:

- Programming Paradigms and Protocols
- Analysis of aggregate properties (lifetime, capacity, …)
- Controllable emergent behavior

Sensor Network Design

Three Software Challenges

Vision: Ubiquitous Sensor Networks

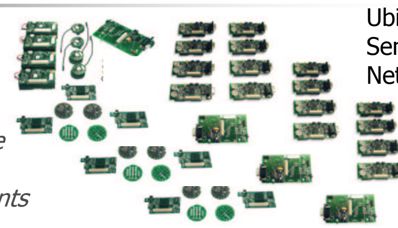

Trend: Moving from single device to networks of thousands of small components

Software Challenges:

- Programming Paradigms and Protocols
- Analysis of aggregate properties (lifetime, capacity, …)
- Controllable emergent behavior

Application Domain

Habitat Monitoring

Precision Agriculture

Sensing Applications

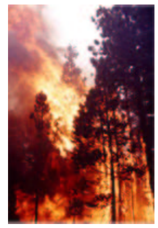

Disaster Response

Features
- Ad hoc deployment
- Massive distribution
- Interaction with a physical environment
- Unattended operation

Border Control

Target Tracking

Infrastructure Protection

Programming Paradigms for Sensor Networks

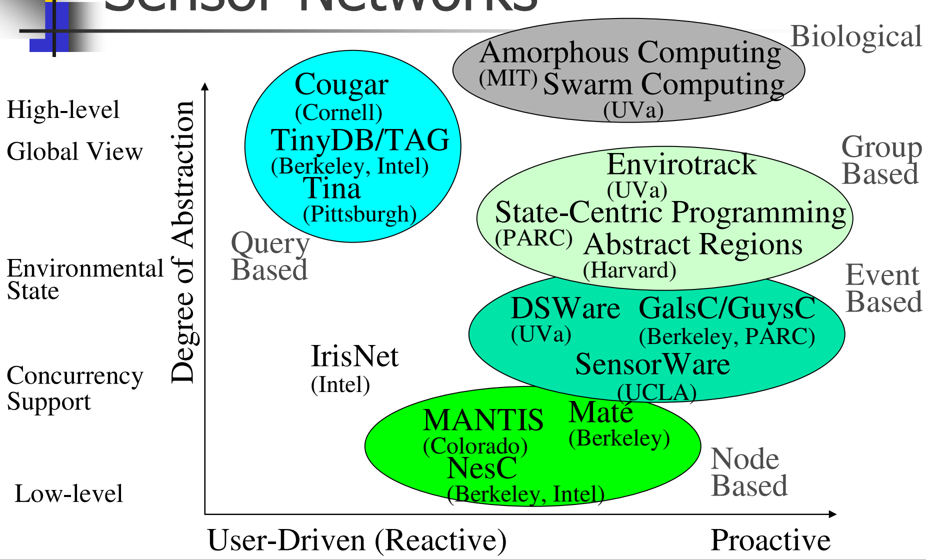

Programming Paradigms for Sensor Networks

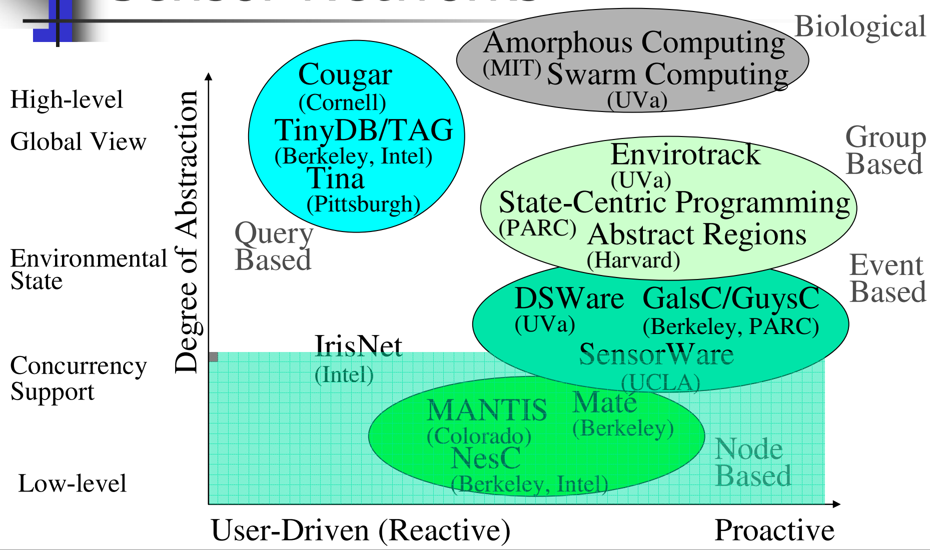

Challenge 1: Energy Balancing

- Problem: data collection by centralized entities

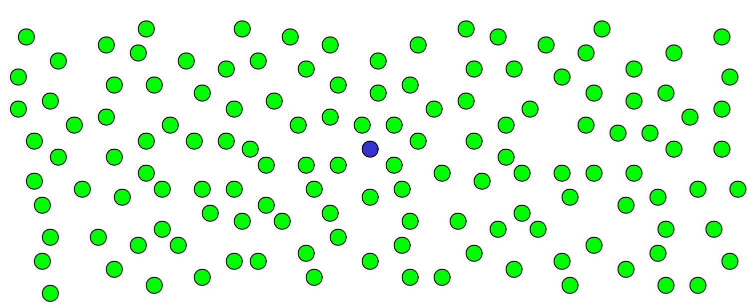

Challenge 1: Energy Balancing

- Problem: data collection by centralized entities

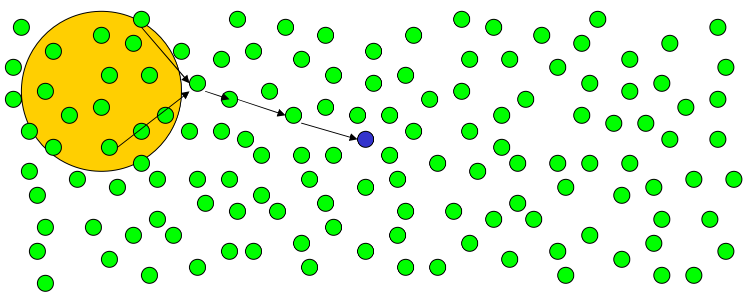

Challenge 1: Energy Balancing

- Problem: data collection by centralized entities

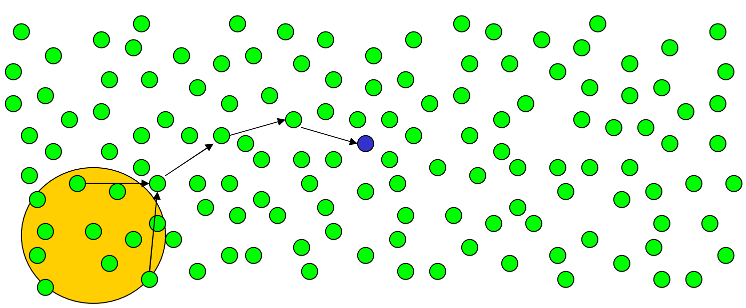

Challenge 1: Energy Balancing

- Problem: data collection by centralized entities

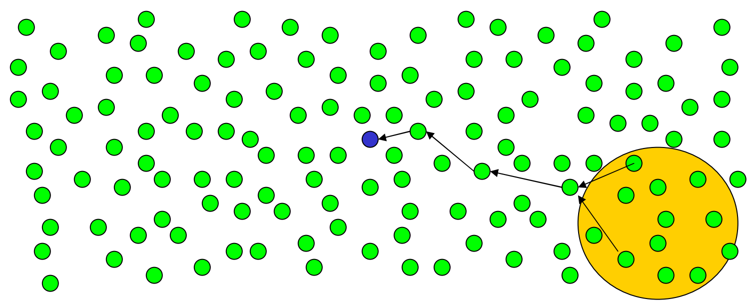

Challenge 1: Energy Balancing

- Problem: data collection by centralized entities

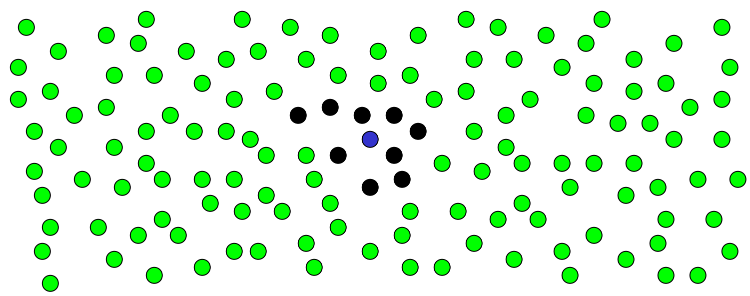

Challenge 1: Energy Balancing

- Problem: data collection by centralized entities

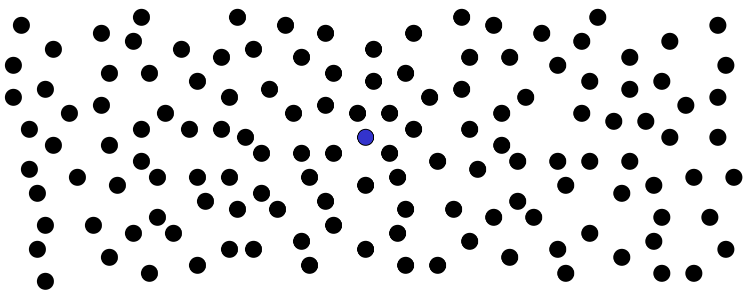

Challenge 1: Energy Balancing

- Problem: data collection by centralized entities
- Implications:
  - Programming paradigms should support autonomous operation
  - Push computation into the network
  - Take the user out of the loop
  - Query based? Event-based? Hierarchical? ...

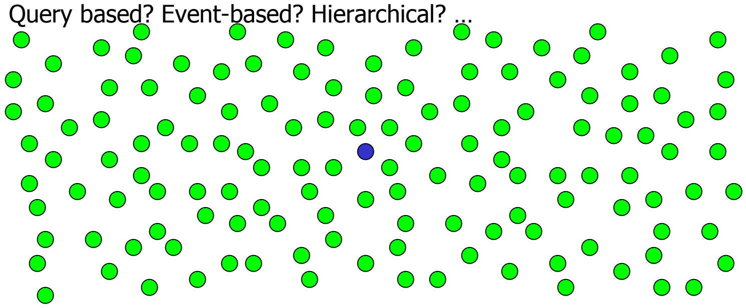

Programming Systems for Sensor Networks

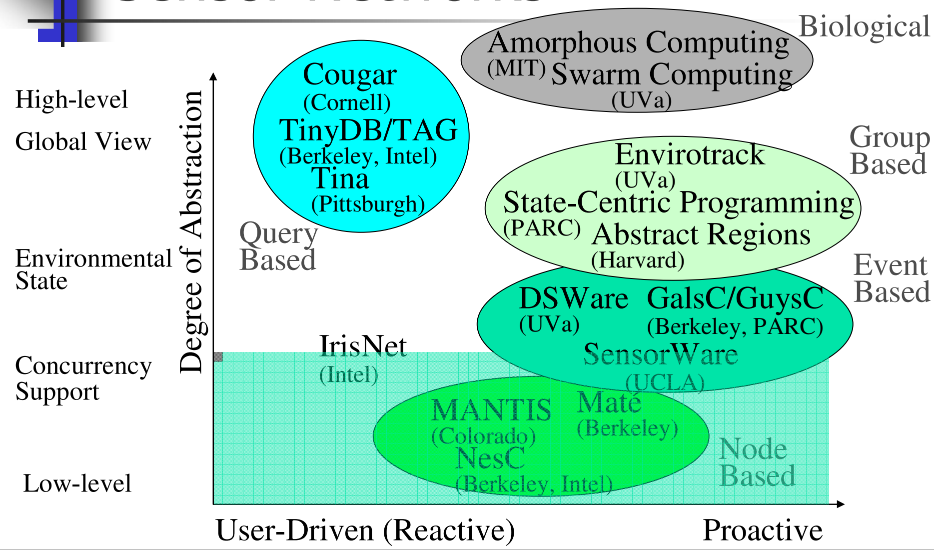

Programming Systems for Sensor Networks

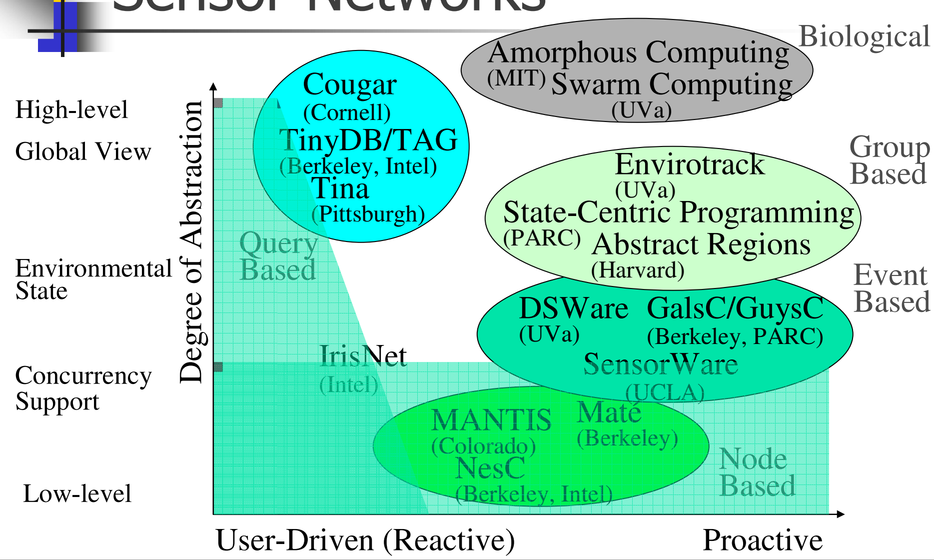

Challenge 2: Environmental Abstractions

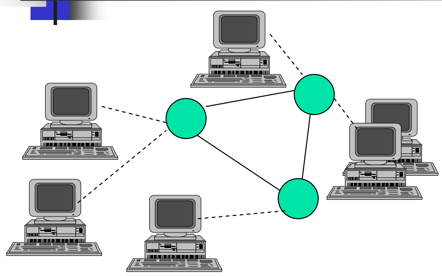

Challenge 2: Environmental Abstractions

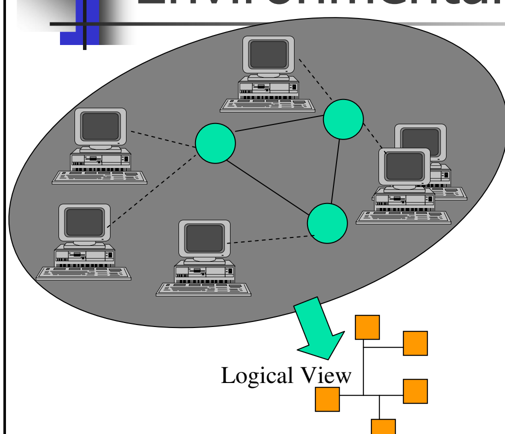

- Distributed programming paradigms
  - Abstract distributed communication
  - Provide location transparency

Challenge 2: Environmental Abstractions

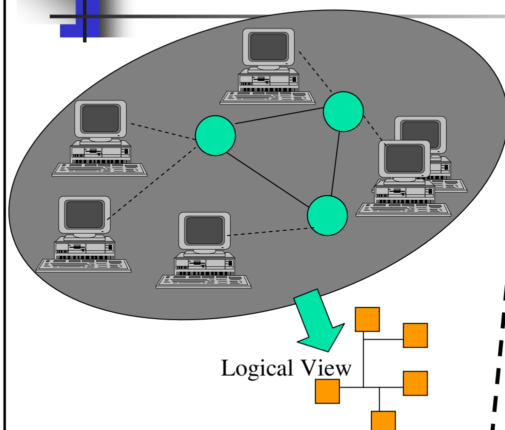

- Distributed programming paradigms
  - Abstract distributed communication
  - Provide location transparency
- Sensor Network Programming Abstractions
  - Represent the physical world to the programmer
  - Abstract distributed interaction with the physical environment

Challenge 2: Environmental Abstractions

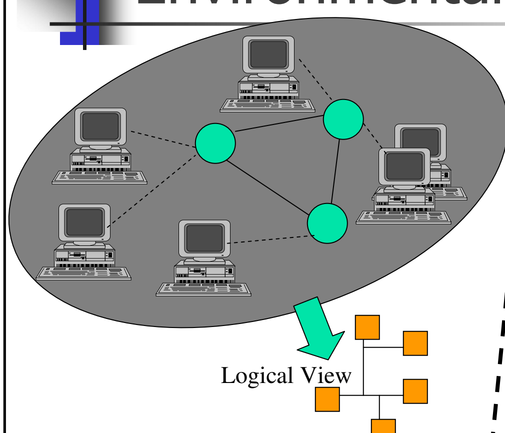

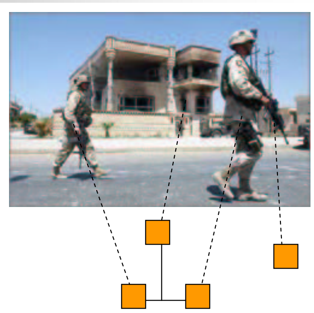

- Distributed programming paradigms
  - Abstract distributed communication
  - Provide location transparency
- Sensor Network Programming Abstractions
  - Represent the physical world to the programmer
  - Abstract distributed interaction with the physical environment

Programming Systems for Sensor Networks

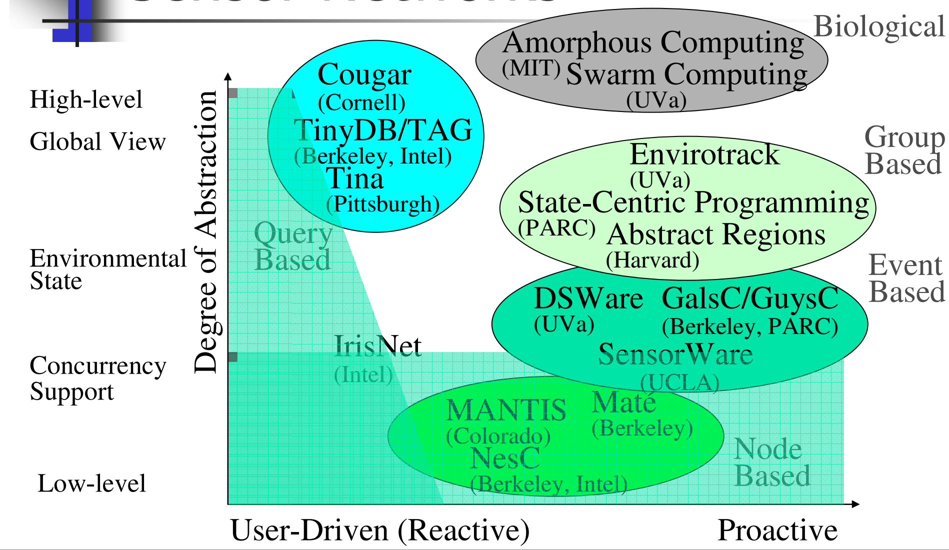

Programming Systems for Sensor Networks

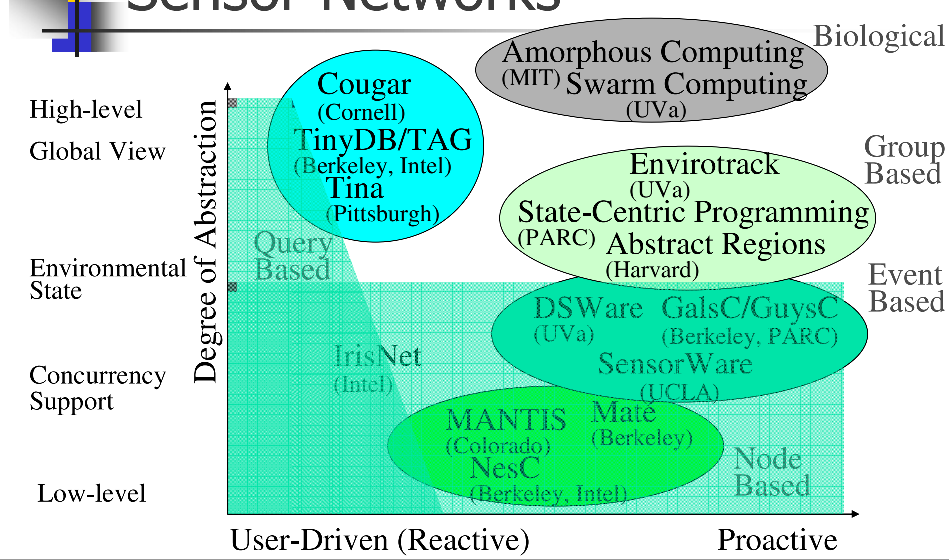

Challenge 3: Distributed Group Management Theory

- Individual nodes or events are insignificant – groups and aggregate state are good abstractions
- Leverage group communication (a mature topic)
  - Real-time
  - Temporal constraints à spatiotemporal constraints
- New group types/semantics:
  - Groups may be very dynamic, high failure-rate, variable topology, migration/mobility, relaxed semantics
  - Semantics cover interaction with the environment
    - Different interactions give rise to different group types
    - New semantics, group properties and invariants
    - Type-specific group communication protocols
  - Programming interface and primitives
- Example: Environmentally Immersive Programming

Environmentally Immersive Programming

- Exports a new address space in which the addressed entities (called *contexts*) are representations of physical entities in the external environment
- *Contexts*:
  - Logical representations of entities in the external world
  - Have unique names (context labels) – same as IP hosts
  - Instantiated when the corresponding external entities are observed in the environment – follow these entities around
  - Tracking objects (tasks) can be attached to contexts to execute in the vicinity of the corresponding real-world entity
- Tasks (attached to contexts) can communicate and invoke each other’s methods remotely

Programming Model

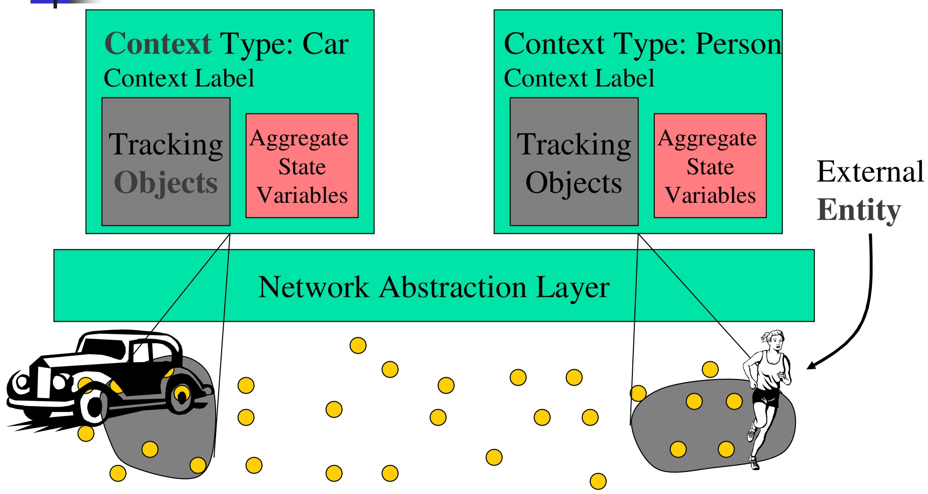

Contexts and Objects

- Contexts: Encapsulate entity state and tracking objects
- Tracking objects: Perform entity-specific computation, communication and sensing

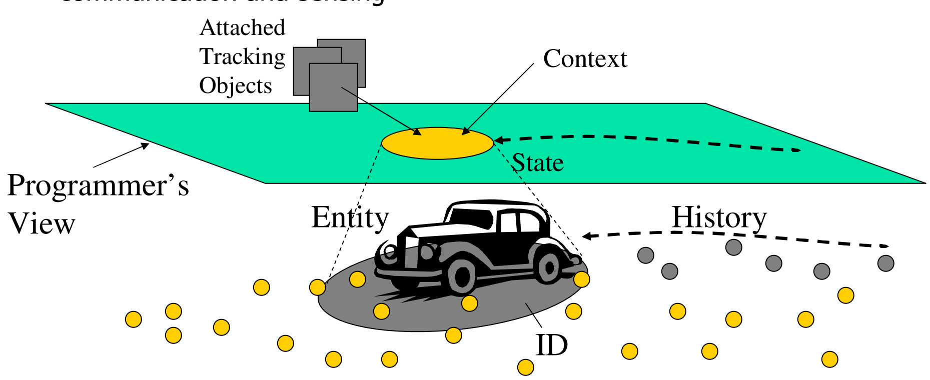

Communication

- Objects may export methods for remote invocation

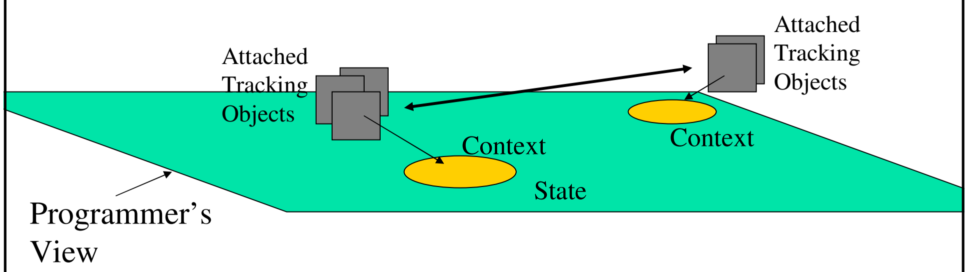

Context Example

```
begin context tracker
  sense: magnetic() + motion();
  state: location = avg (position);
end
```

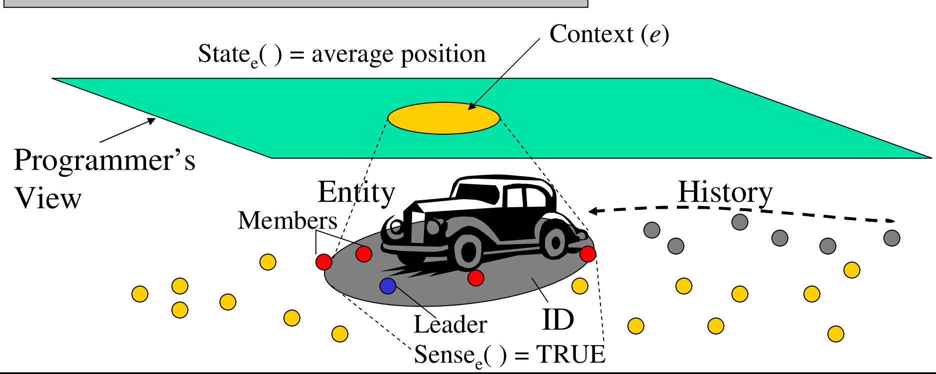

Attaching Objects

begin context *tracker*
  sense: magnetic() + motion();
  state: location = avg (position);
end

begin object *reporter*
  send (state, home);
end

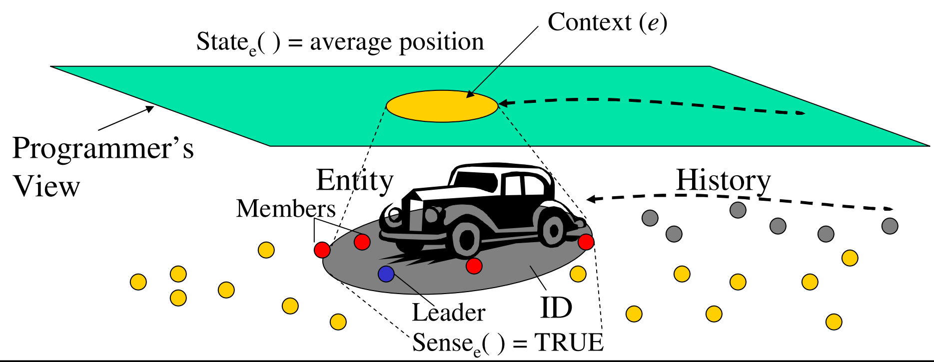

Attaching Objects

begin context *tracker*
  sense: magnetic() + motion();
  state: location = avg (position);
end

begin object *reporter*
  send (state, home);
end

Attach

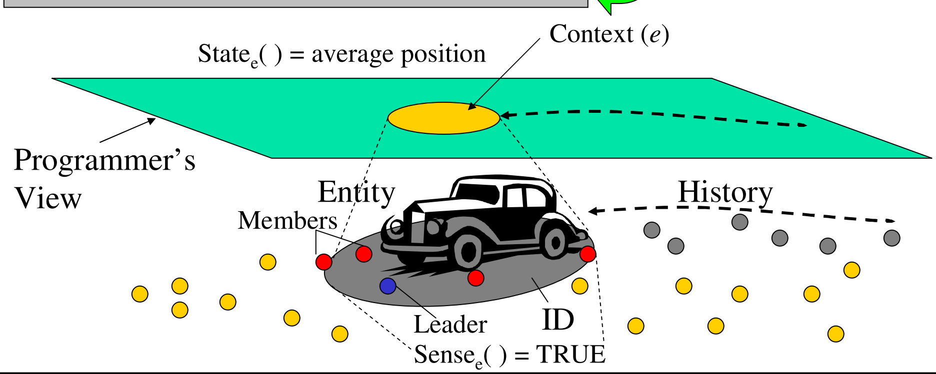

Attaching Objects

```
begin context tracker
  sense: magnetic() + motion();
  state: location = avg (position);
end
```

```
begin object reporter
  send (state, home);
end
```

```
begin object mic
  turn-on microphone
  send (sound, home);
end
```

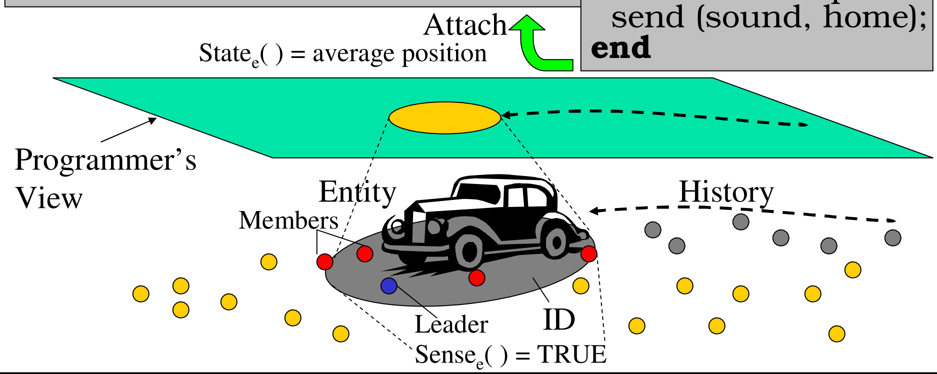

Challenges

- Context maintenance
  - Context definition and instantiation (discovery of external entities)
  - Aggregate state management (abstract state of a dynamically changing group of sensor nodes)
- Unique context representation
  - Uniqueness
  - Context migration and handoff (track locales of mobile physical entities in the environment)
- Communication, entity location and remote invocation

Group Management and Context Formation

- Nodes sensing a given target form a single context
- Context has leader, members, **and followers**

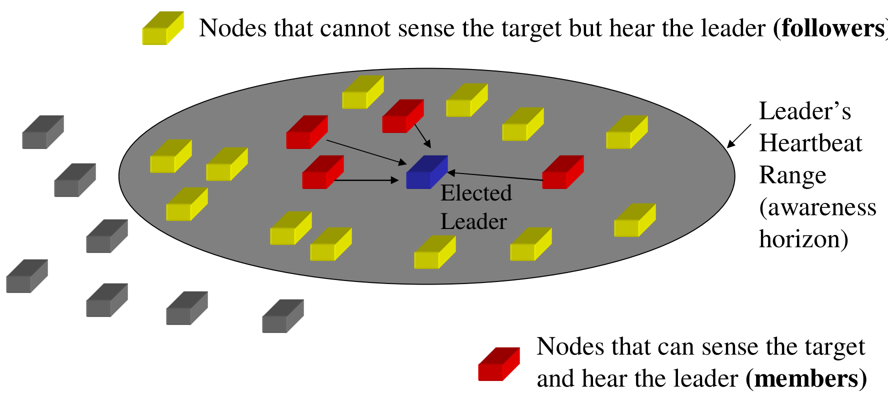

A New Spatiotemporal Problem: Unique Representation

- Target can't move fast enough to be sensed by a node that is outside leader heartbeat range:
- Communication range > 2 Sensing Range + Target speed x group migration time

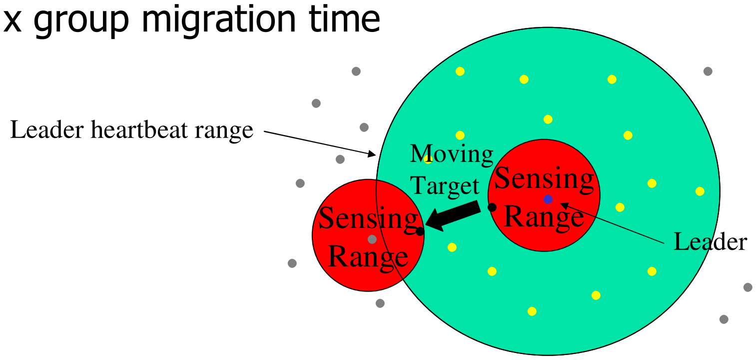

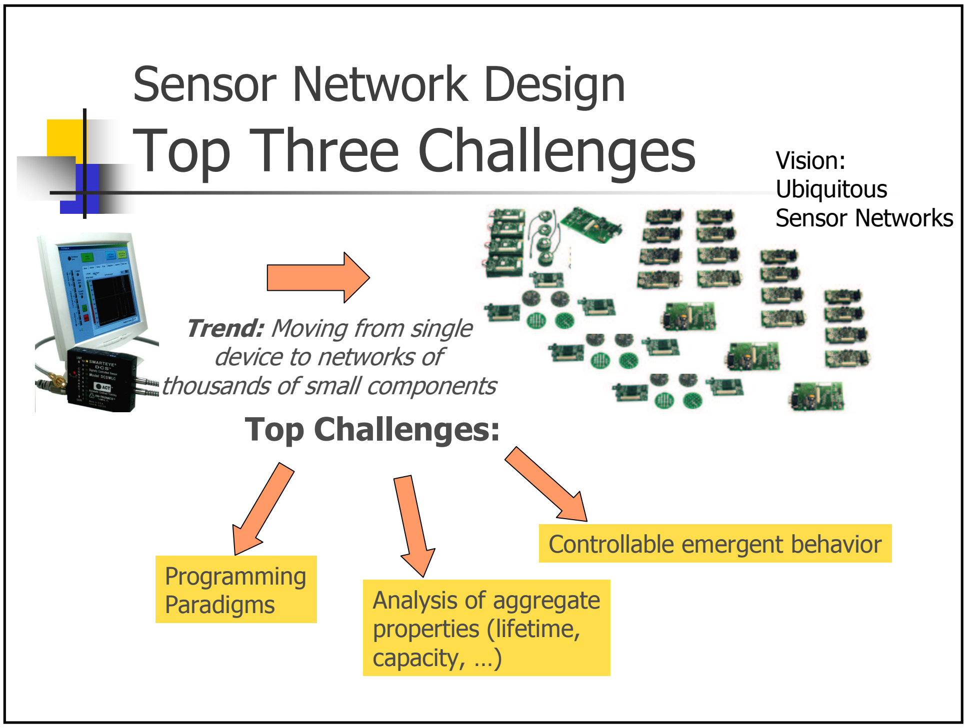
Sensor Network Design
Top Three Challenges
Vision:
Ubiquitous
Sensor Networks
Trend: Moving from single device to networks of thousands of small components
Top Challenges:
Programming Paradigms
Analysis of aggregate properties (lifetime, capacity, …)
Controllable emergent behavior

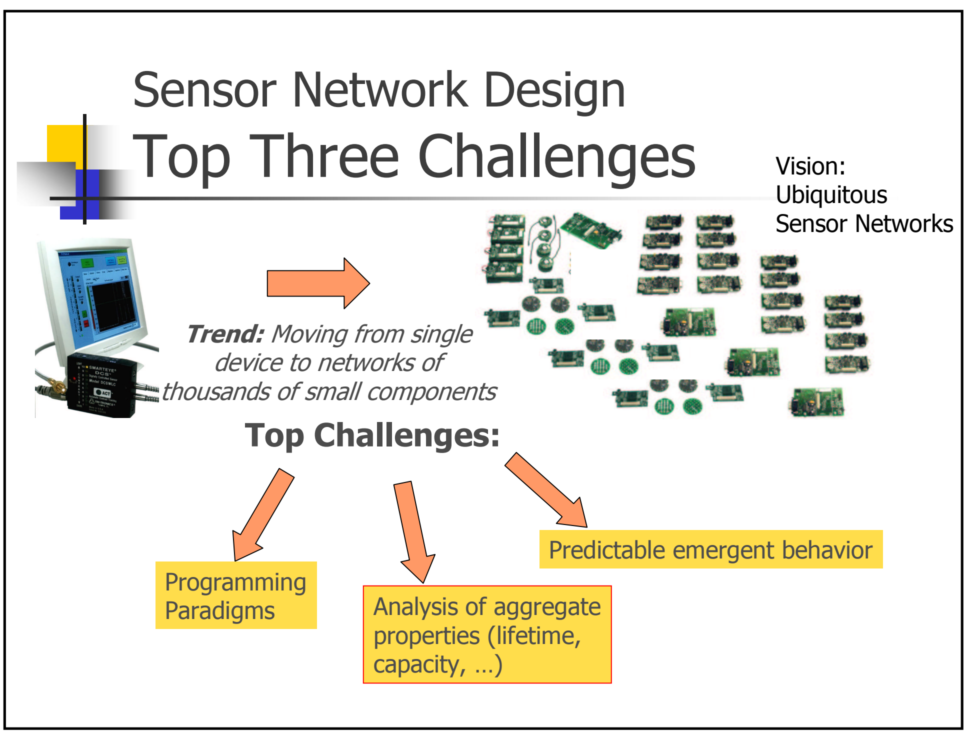
Sensor Network Design
Top Three Challenges
Vision:
Ubiquitous
Sensor Networks
Trend: Moving from single device to networks of thousands of small components
Top Challenges:
Programming Paradigms
Analysis of aggregate properties (lifetime, capacity, …)
Predictable emergent behavior

Real-Time Analysis of Sensor Networks

- What is the relation between radio range, network density, total number of nodes, number of sinks, packet length, packet scheduling policy, MAC-layer protocol, and end-to-end **packet ability to meet deadlines?**
- How to choose network parameters that satisfy delay bounds on communication?
- Optimality (of MAC scheduling policies), capacity proofs, convergence proofs, etc.

Real-Time Objects

```
begin context tracker
  sense: magnetic() + motion();
  state: location = avg (position,3,2);
end
```

```
begin RTobject reporter
  send (state, home,30s);
end
```

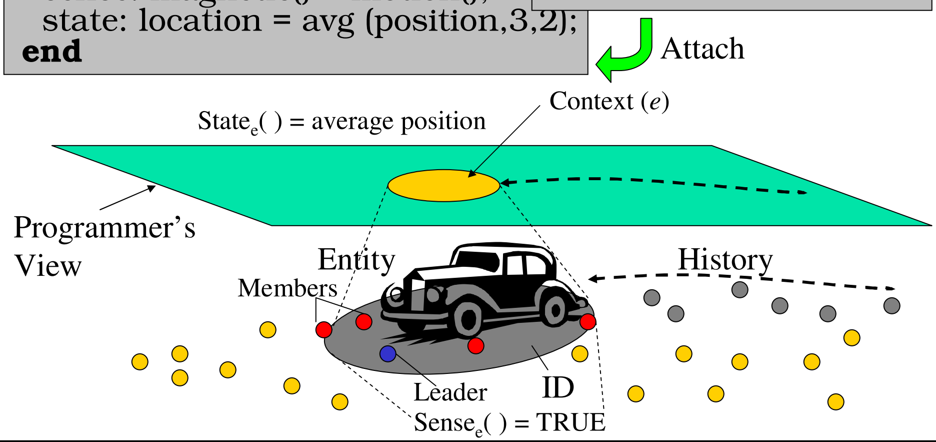

Capacity Planning for Real-Time Wireless Sensor Networks

- Seminal recent work established wireless network capacity bounds
- What if traffic has deadlines and only bits that make it by the latency constraint are counted towards throughput?
- Problem: express *real-time* network capacity that quantifies the throughput of timely bits only as a function of network parameters and time constraints

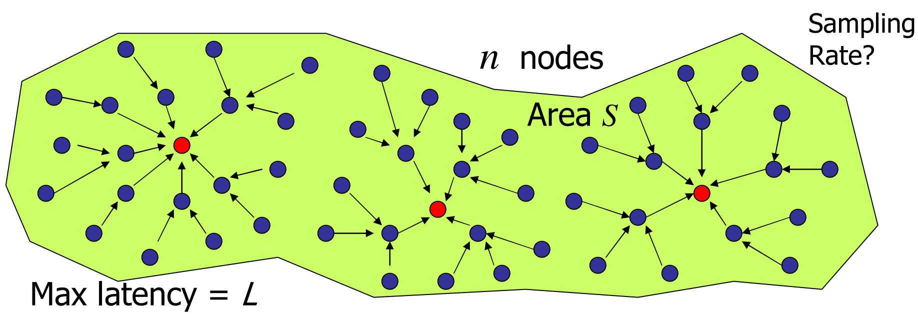

Network Real-time Capacity

- Network bandwidth is the bottleneck (communication scheduling problems)
  - Task processing time à packet transmission time
  - Scheduler queue à network queue
- Intuitively, network schedulability *decreases* with:
  - Increased packet size (task processing time), C
  - Increased distance between source and destination, L
  - Decreased end-to-end latency constraint, D
- Schedulability decreases with CL/D
- Is there a bound $Capacity_{RT}$, such that all packets, i, reach their destinations by their deadlines if:

$$\sum_i \frac{C_i L_i}{D_i} \leq Capacity_{RT}$$

The Stage Delay Theorem

- Let us define synthetic utilization at a resource as: $\Sigma\, C_i/D_i$ for all eligible tasks through that resource (those that have arrived to the system but whose end-to-end deadline has not expired)

The Stage Delay Theorem: If the synthetic utilization of resource j, does not exceed U_j, then no task is queued on resource j for more than a fraction β_j of its end-to-end deadline*, where:

$$\beta_j = U_j\,(1 - U_j/2)/(1 - U_j)$$

*This is under deadline monotonic scheduling. Similar results derived for other policies.

The Single Hop Problem

- Only one node in the vicinity of a receiver can send at a time

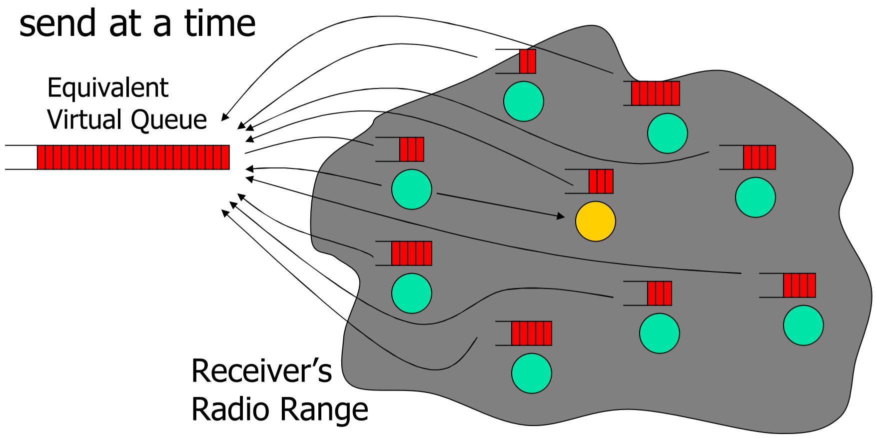

Throughput Optimization: Maximizing Real-Time Capacity

- Consider a localized communication pattern where each node communicates with nodes at most N hops away.
- On any path, maximize $\Sigma_j U_j$ subject to

$$\sum_{j=1}^{N} \frac{U_j(1-U_j/2)}{1-U_j} \le 1$$

- From symmetry, $U_j = U$

$$\frac{U(1-U/2)}{1-U} = 1/N$$

- Hence, $U = \left(1 + 1/N - \sqrt{1+(1/N)^2}\right)$

Real-Time Capacity

- ***The total capacity theorem:*** In a ***load-balanced*** network of n nodes, each with a radio of transmission speed W and m neighbors on average, if communication is localized within at most N hops:

$$Capacity_{Opt} = \frac{nW}{m}\left(1 + 1/N - \sqrt{1+(1/N)^2}\right)$$

- For large N:

$$Capacity_{Opt} \approx \frac{nW}{mN}$$

Real-time Capacity of Multi-hop Data-Collection in Sensor Networks

- In a ***data collection*** sensor network with K collection points, maximum path length N, and radio transmission speed W, what is a sufficient bound on real-time capacity?

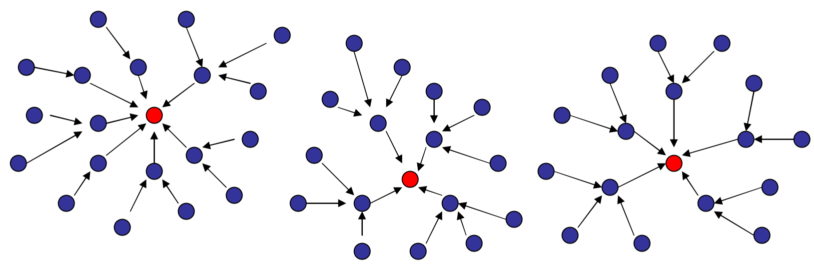

Real-time Capacity of Multi-hop Data-Collection in Sensor Networks

- In a ***data collection*** sensor network with K collection points, maximum path length N, and radio transmission speed W, what is a sufficient bound on real-time capacity?

$$\sum_{j \in path} \frac{U_j(1-U_j/2)}{1-U_j} \leq 1, \quad U_j \propto 1/j$$

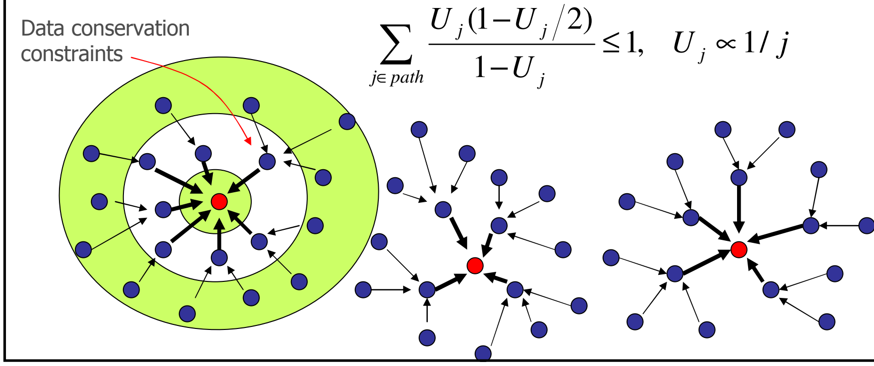

Real-time Capacity of Multi-hop Data-Collection in Sensor Networks

- In a ***data collection*** sensor network with K collection points, maximum path length N, and radio transmission speed W, a sufficient bound on real-time capacity is:

$$Capacity_{DC} \approx \frac{KNW}{1 + 0.5 \ln N}$$

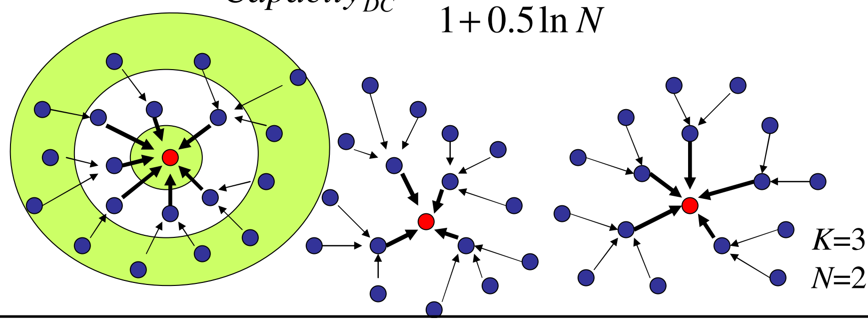

Evaluation: How Pessimistic is Real-time Capacity?

- Simulation versus analytic prediction of the onset of deadline misses in a 1600 node network

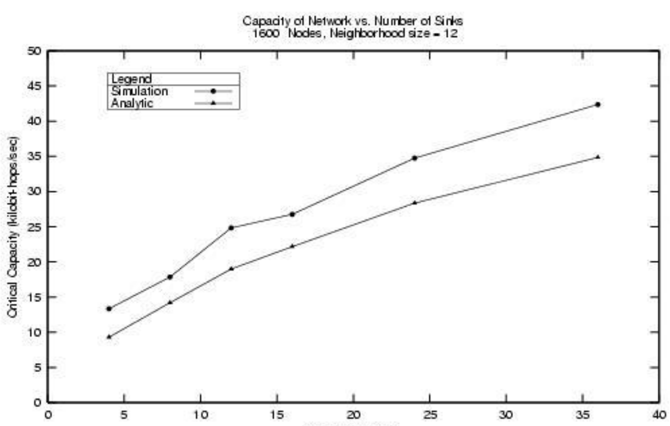

Sensor Network Design
Top Three Challenges

Vision:
Ubiquitous
Sensor Networks

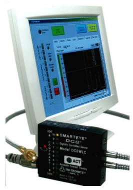

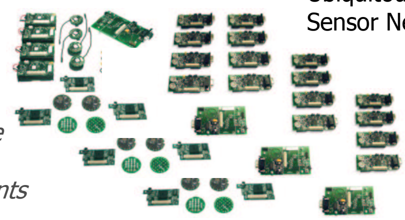

Trend: *Moving from single device to networks of thousands of small components*

Top Challenges:

- Programming Paradigms
- Analysis of aggregate properties (timing, capacity, …)
- Controllable emergent behavior

Sensor Network Design
Top Three Challenges

Vision:
Ubiquitous
Sensor Networks

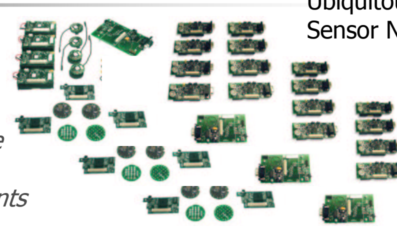

Trend: *Moving from single device to networks of thousands of small components*

Top Challenges:

- Programming Paradigms
- Analysis of aggregate properties (timing, capacity, …)
- Controllable emergent behavior

Fundamental Performance Tradeoffs in Sensor Networks

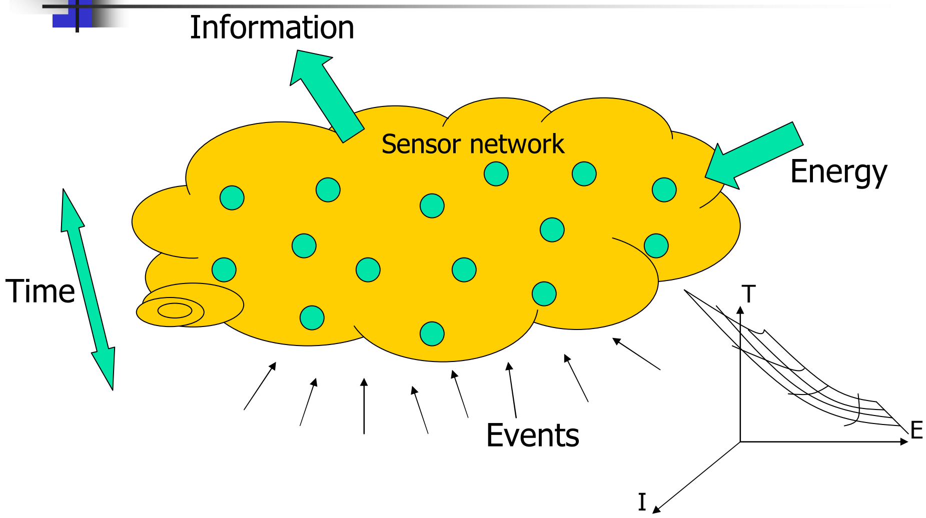

Fundamental Performance Tradeoffs in Sensor Networks

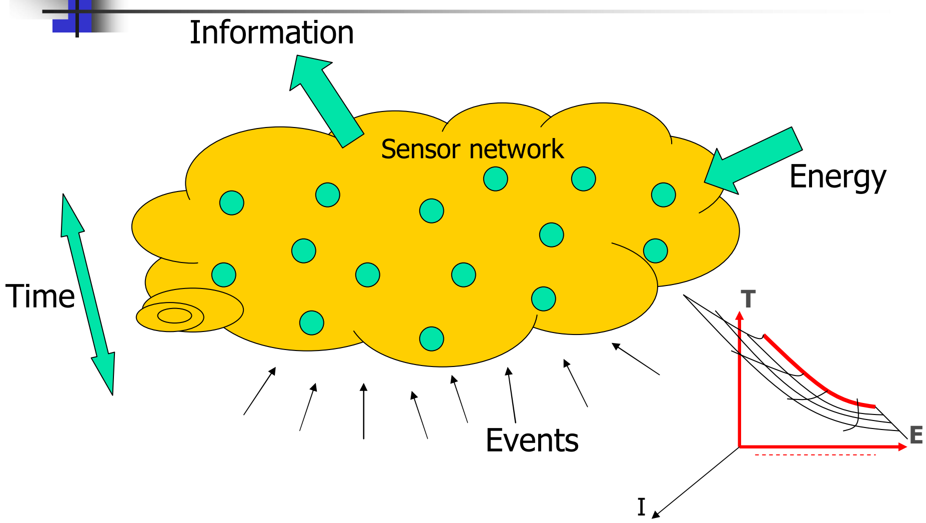

Fundamental Performance Tradeoffs in Sensor Networks

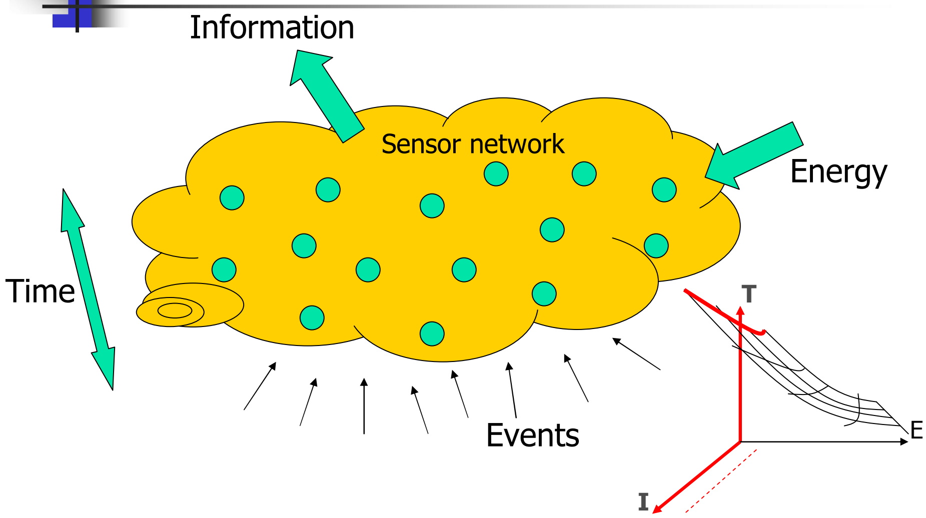

Fundamental Performance Tradeoffs in Sensor Networks

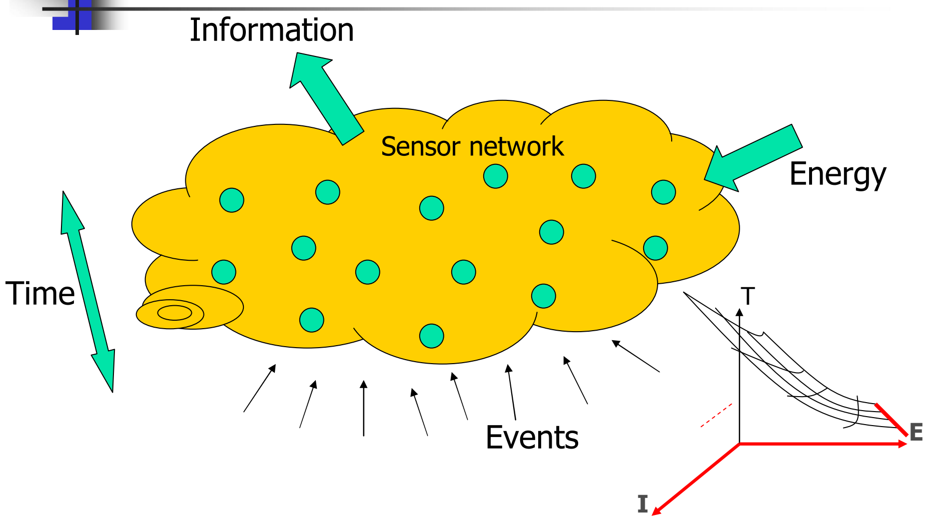

Constrained Optimization and Feedback Control in Sensor Networks

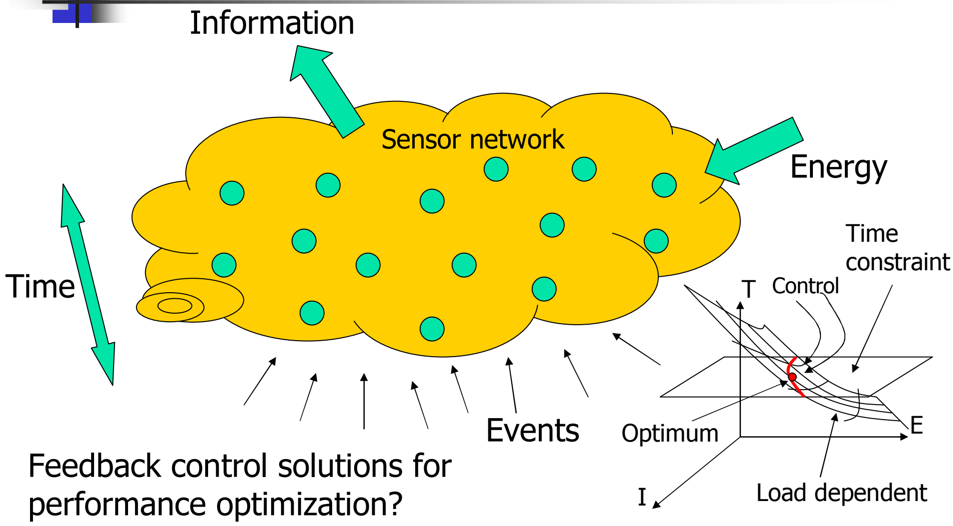

Feedback control solutions for performance optimization?

Convergence of Global Behavior in Sensor Networks

- Sensor network protocols are localized; nodes act independently, locally, in response to local stimuli.
- How to argue about the global effects of such localized protocols when performed by all nodes?
- How to induce and analyze convergence to desired global properties?
- Inspirations from control theory (convergence, stability), Markov decision theory (Markov chains, stochastic models), biology (behavior of social insects, swarm intelligence, bio-differentiation), physics (phase transitions, crystallization), …

Designing Adaptive Performance-Assured Sensor Networks

- Achieving desired temporal behavior without fine-grained knowledge
  - Compute feasibility boundaries
  - Control the system not to escape them

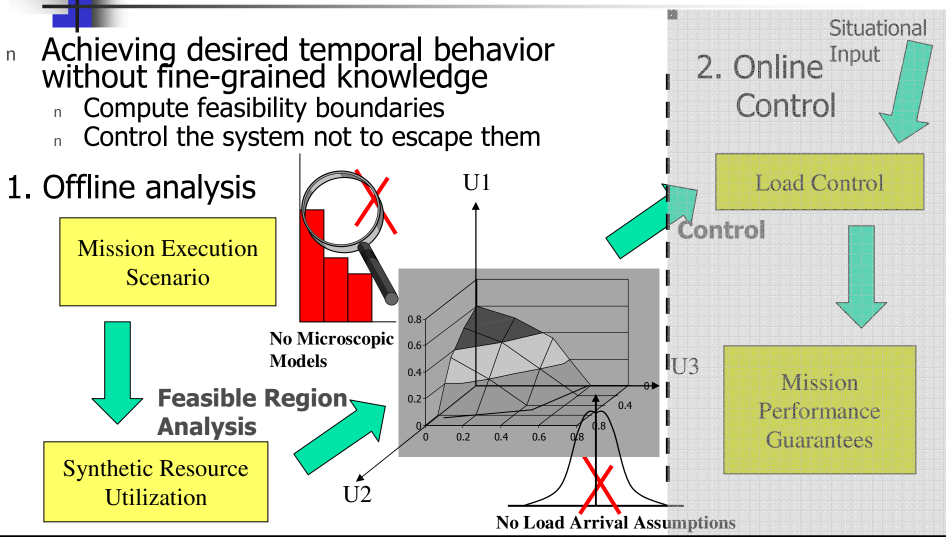

Final Word: **Sufficient Simplicity**
An Underlying Design Principle

Envisioned solution Techniques should be:

- **deterministic** analysis
- **aperiodic** tasks
- very **simple** analysis
- **high-level** models
- only **sufficient**

Less error prone (simplicity), more scalable (simplicity), and safe (sufficiency)

Trend

Increasing scale/complexity
Decreasing resource cost
Increasing degree of interaction
Less predictable dynamics

Current solution techniques are:

exact

complex

resource optimal

Conclusions

- Embedded computing is of growing importance in computer science as a discipline
- New computing paradigms and abstractions are needed
- Real-time analysis must catch up with sensor network realities
  - Fundamentally new problems with space and time constraints
  - Aggregate results as opposed to microscopic models
- Convergence and aggregate behavior analysis is to be developed
  - Real-time control of aggregate network properties
- The science of sensor network design is yet to emerge