

# Challenge Problems in Sensornet Research

Feng Zhao

Microsoft Research

<http://research.microsoft.com/zhao>


# Acknowledgment

- Benefited from discussions with my colleagues at MSR, Xerox PARC, Stanford
- Joint work with Jie Liu, Leo Guibas, Jeremy Elson, Alec Woo, Kamin Whitehouse, Elaine Cheong, Prabal Dutta, Zoe Abrams, Siddharth Seth, Ryan Newton, Andrew Parker

# Challenge for the Sensornet Community

- Great to see startups begin to tackle practical problems, collecting low-hanging fruits
- But, the research community must set the sight further
  - Every field needs grand challenges to focus energy and drive progress, so does sensornet
  - Here is my take on the grand challenges; substitute it with you favorite ...

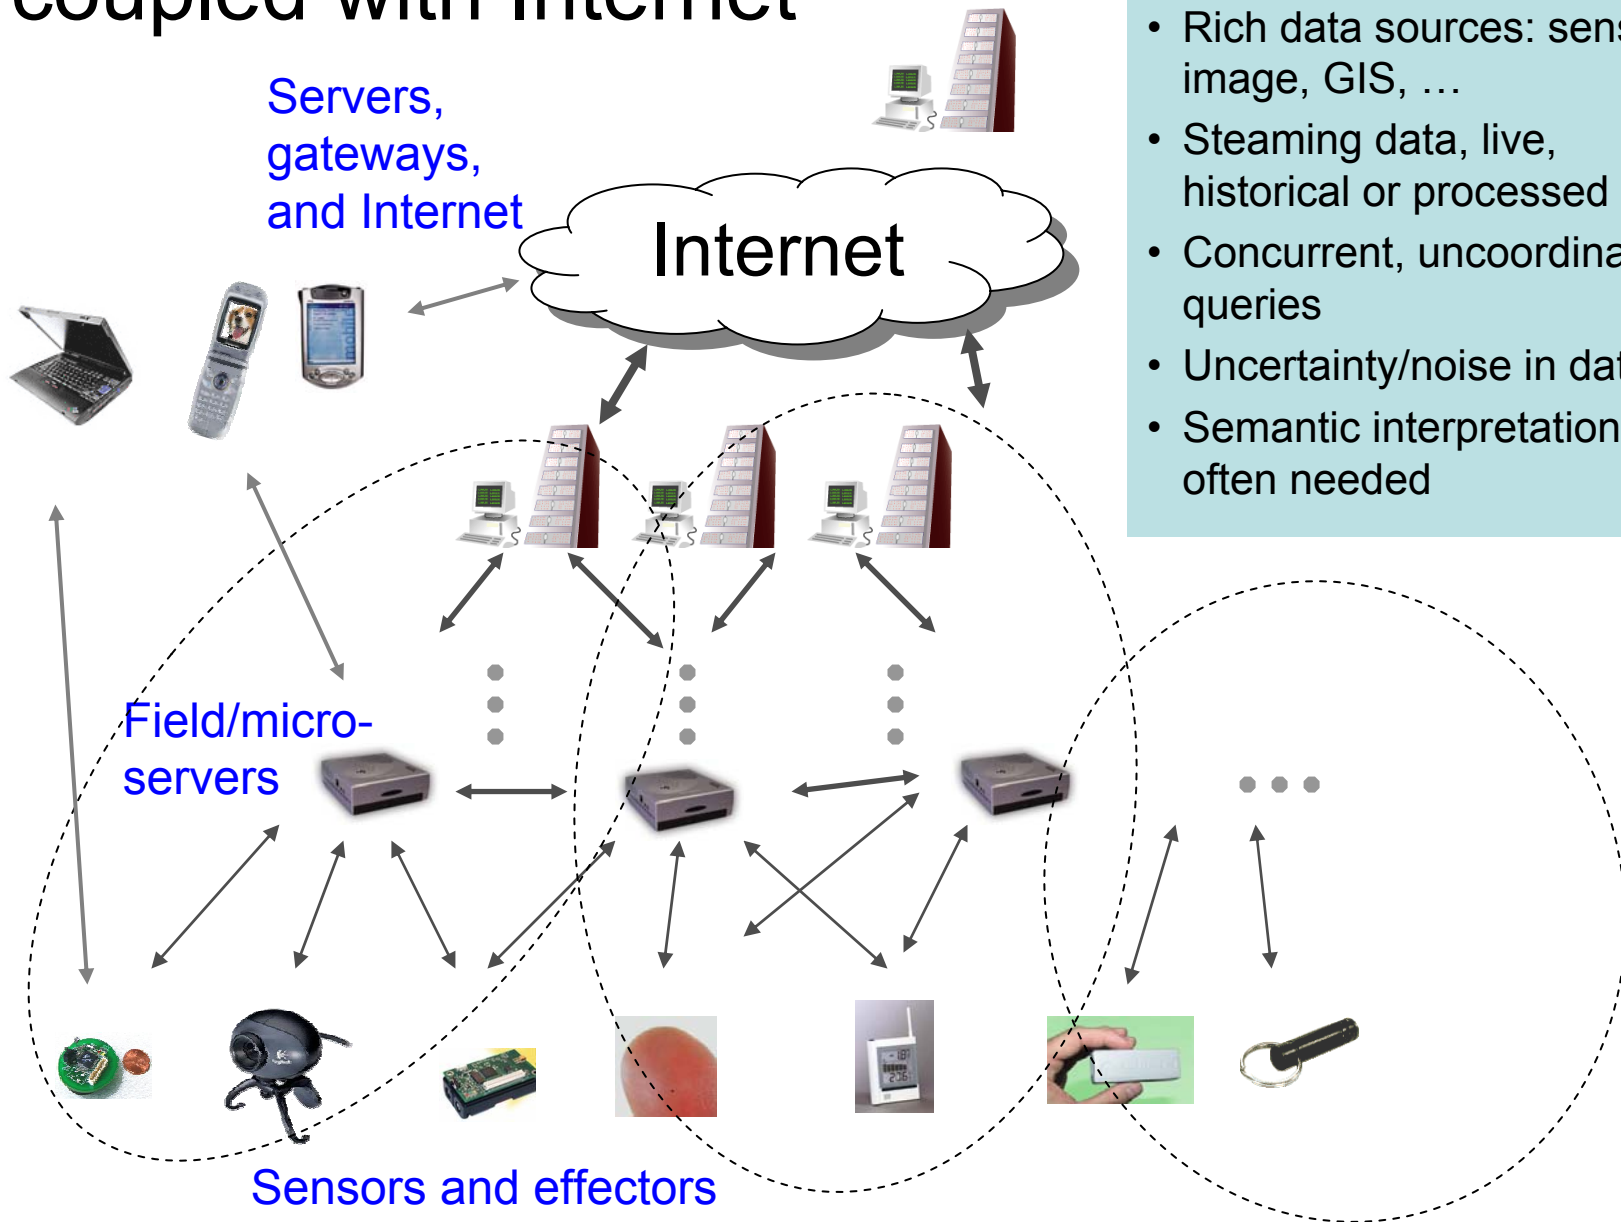
# Reality Browser: Query physical world, live and up close, from anywhere



# Potential Apps

- Environment
  - Volcano, underwater, rainforest
- Education
  - K-12 interactive learning, virtual laboratory
- Leisure
  - Virtual travel, sunny spot tracker (microclimate), what is the temperature at my favorite beach? what is the water algae level?
- Getting around
  - Where is the nearest available parking space? What is the traffic like on the bridge? How long is the queue at the gas station? Where is the bus?

# Multi-tier net of sensornets, coupled with Internet



## Characteristics:

- Rich data sources: sensor, image, GIS, ...
- Streaming data, live, historical or processed
- Concurrent, uncoordinated queries
- Uncertainty/noise in data
- Semantic interpretation often needed

# Three Challenge Problems

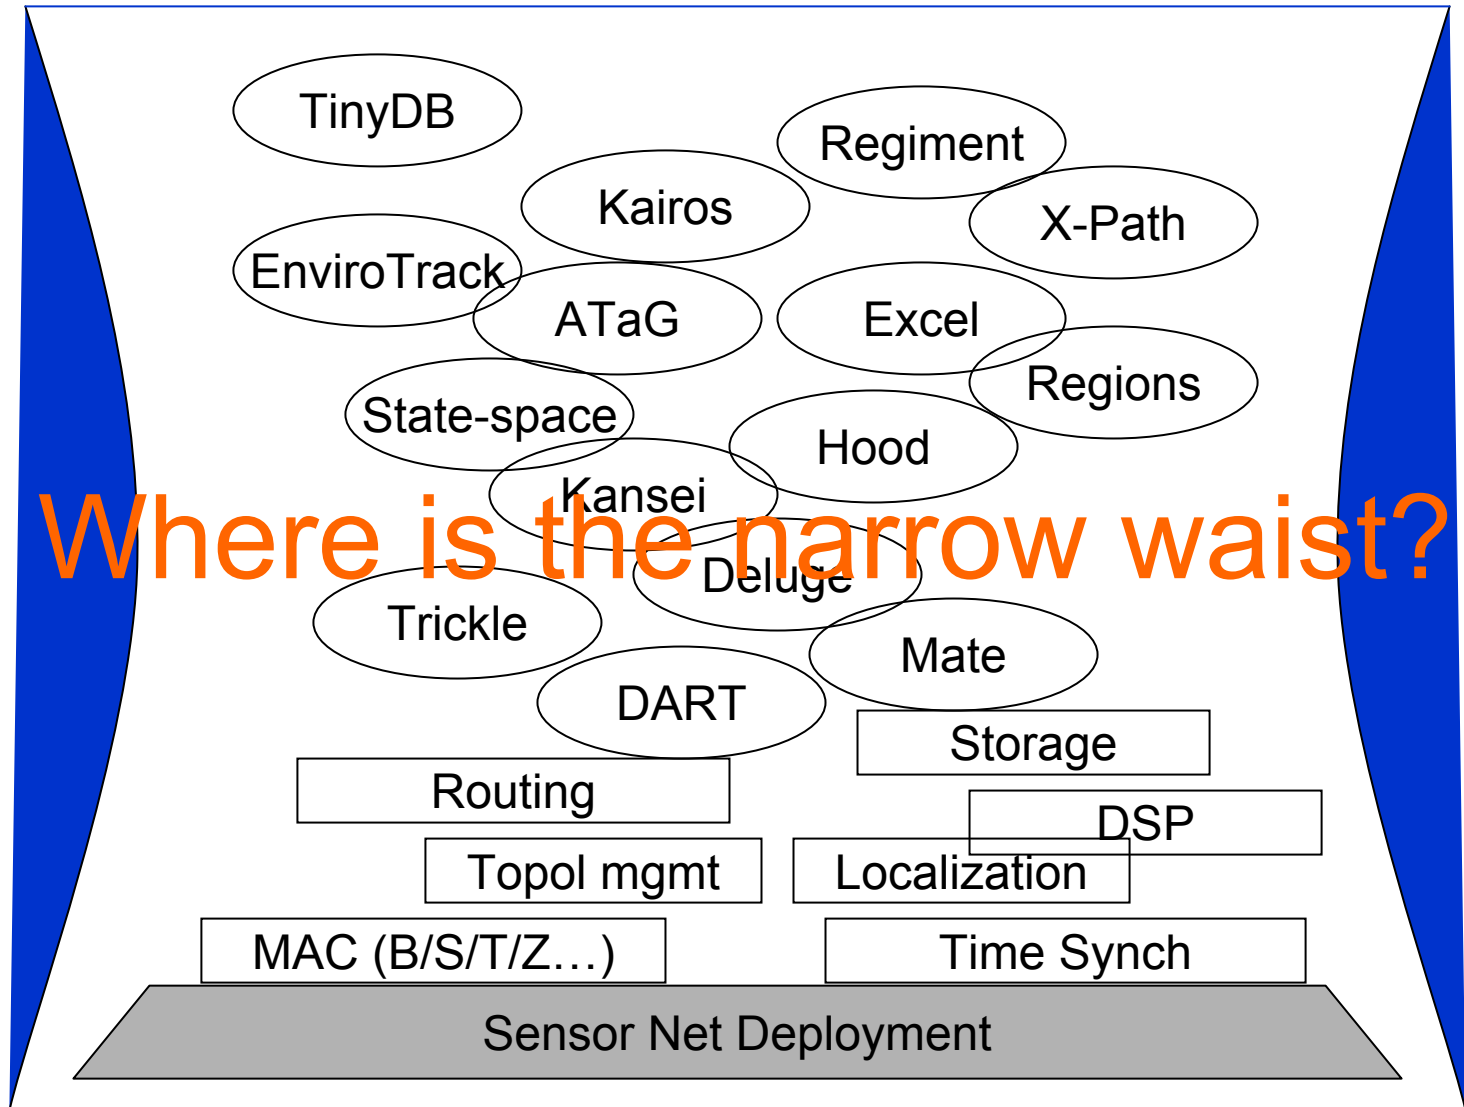
- Need to standardize **interfaces** between apps/db/nw/dsp/hw (e.g., SP link layer abstraction, SensorML)
- Need to treat **uncertainty** (in both data and systems) as first-class entity, for reasoning and system mgmt
- Need simple **tools** for sys config/mgmt, for data collection and vis, for in-situ debugging

# But the problem is made hard by ...

- A great deal of variability
  - Variability in app requirements
    - Use scenarios (data collection, control in the loop, or in-network compression), data rate, hw
  - System dynamism and unreliability
    - Parts of system may be deployed over time by different vendors, using different technologies (e.g., network protocols)
    - Nodes/links come and go
    - Possibly non-replenishable resources
  - Variability in data
    - Uncertain data due to sensor noise, packet loss
    - Incomplete information due to partial observability of the world
- User tasks often specified in high-level semantic queries
  - E.g., “Tell me if you see a red car”, or, “Doing this with that data”.
- As a result, systems are often built vertically
  - With own system abstractions and data models
  - Make sense from efficiency POV, but with relatively small code reuse
  - Phil Buonadonna: “Every sensor net experiment needs 2 PhDs and 5 grad students”, or something like that



Many systems got built over the past few years ...



# From printer land, Mojave Desert, to parking garage ... our own experiences


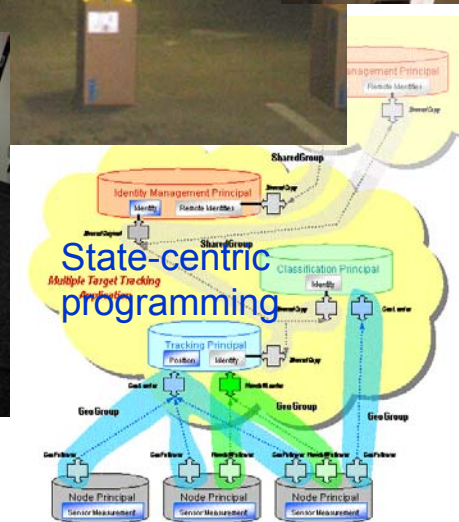
MSR parking garage testbed, June 2004




MSR in-building testbed, 2005



PARC camera testbed, 2003



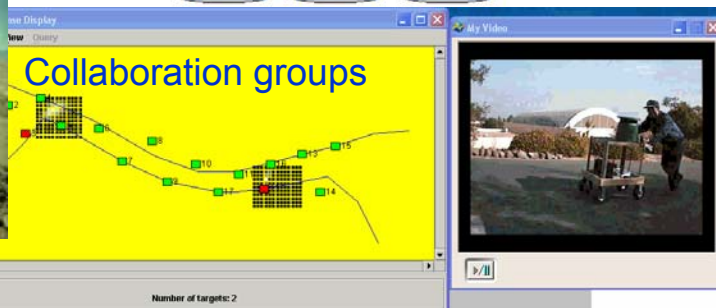
Mote testbed



Machine diagnostics, 1997



DARPA SensIT experiments, 29 Palms, 2000-2001



Outdoor experiment, 2002

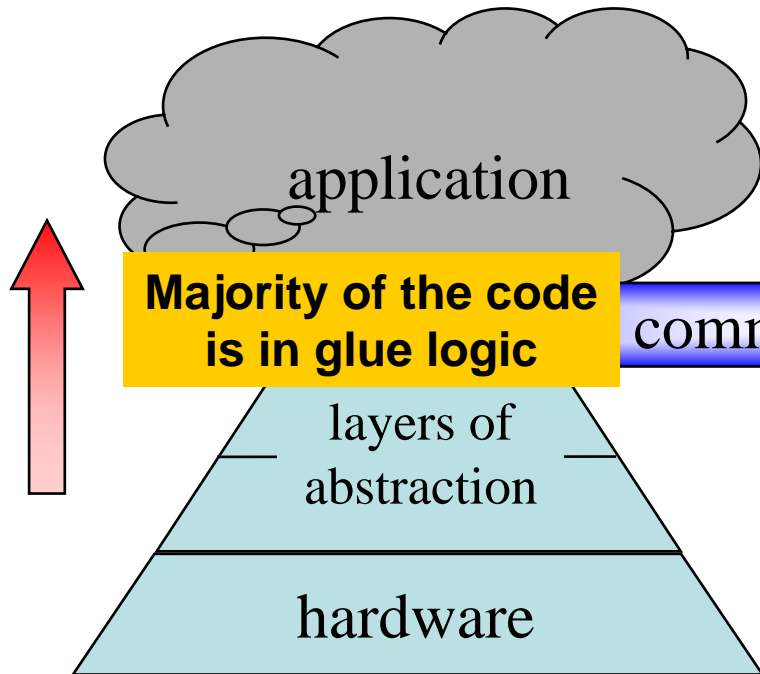
# What we've learned ...

- Proposed tracking as a canonical problem for sensorset problem (circa 2000)
  - Dynamic data, close the loop between sensing & decision making (not just data collection)
  - Created SensIT data sets (with ground truth!), app scenarios, experimental testbed (used by others)
  - Many projects since then: SensIT 29P demo, NEST demo, EnviroTrack, ...
  - But the utility of tracking app is often a function of good signal processing
- Developed the IDSQ (Information Driven Sensor Querying) framework/Sensor tasking (circa 2001)
  - Pay more attention to data: value of information
  - Interface btw app ↔ routing: value of info and routing decision at each node
  - Move beyond individual nodes: group/neighborhood management
    - Many other work in this space: Hood, Regions, ...
  - But ours (and others) with it own interface assumptions and data representation

# More on Lessons Learned

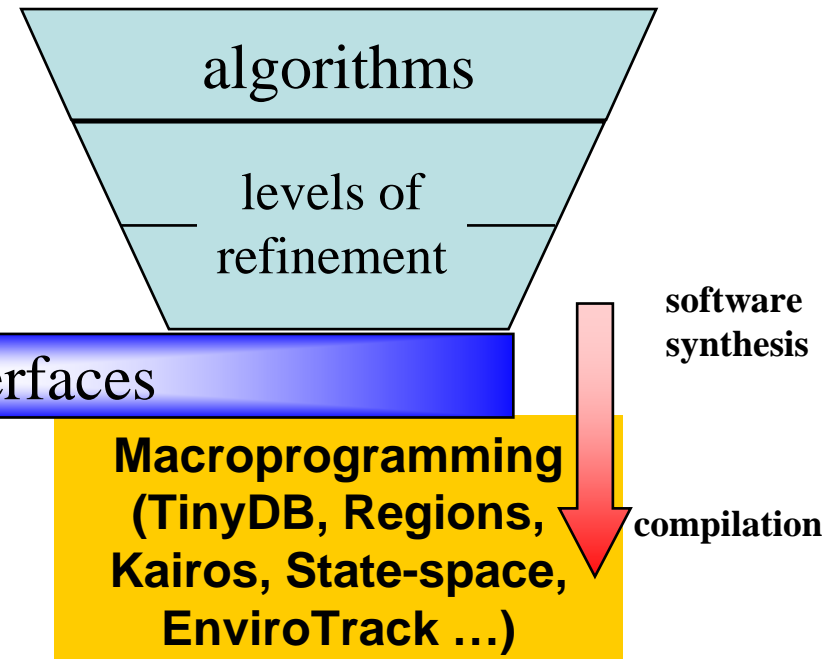
OS/Network-centric view

Designing component-level  
abstractions



Information-centric view

Designing application-level  
abstractions



# A Parking Garage Example




Break beam sensors

## Sense

- Speed, length, and direction of vehicles
- Magnetic signature of vehicles
- Image of vehicles



Micro-server

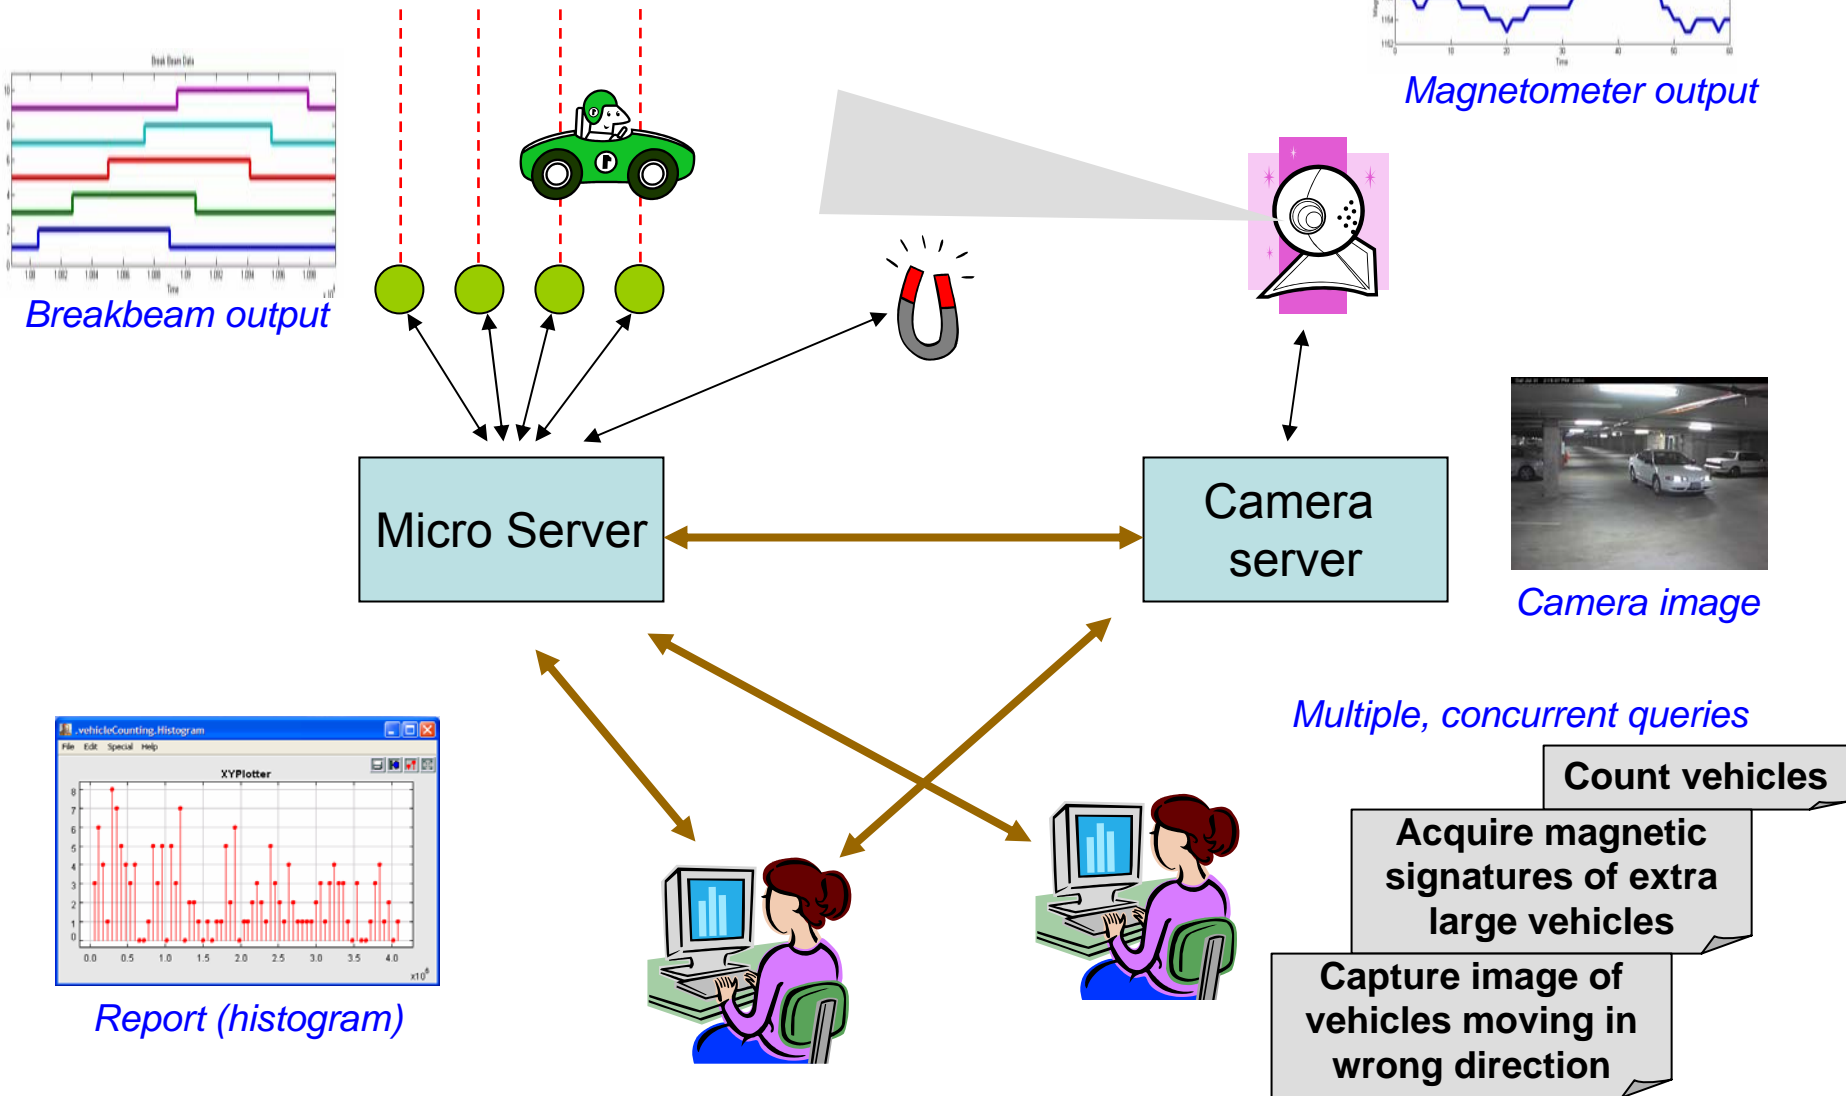


Magnetometer

Camera

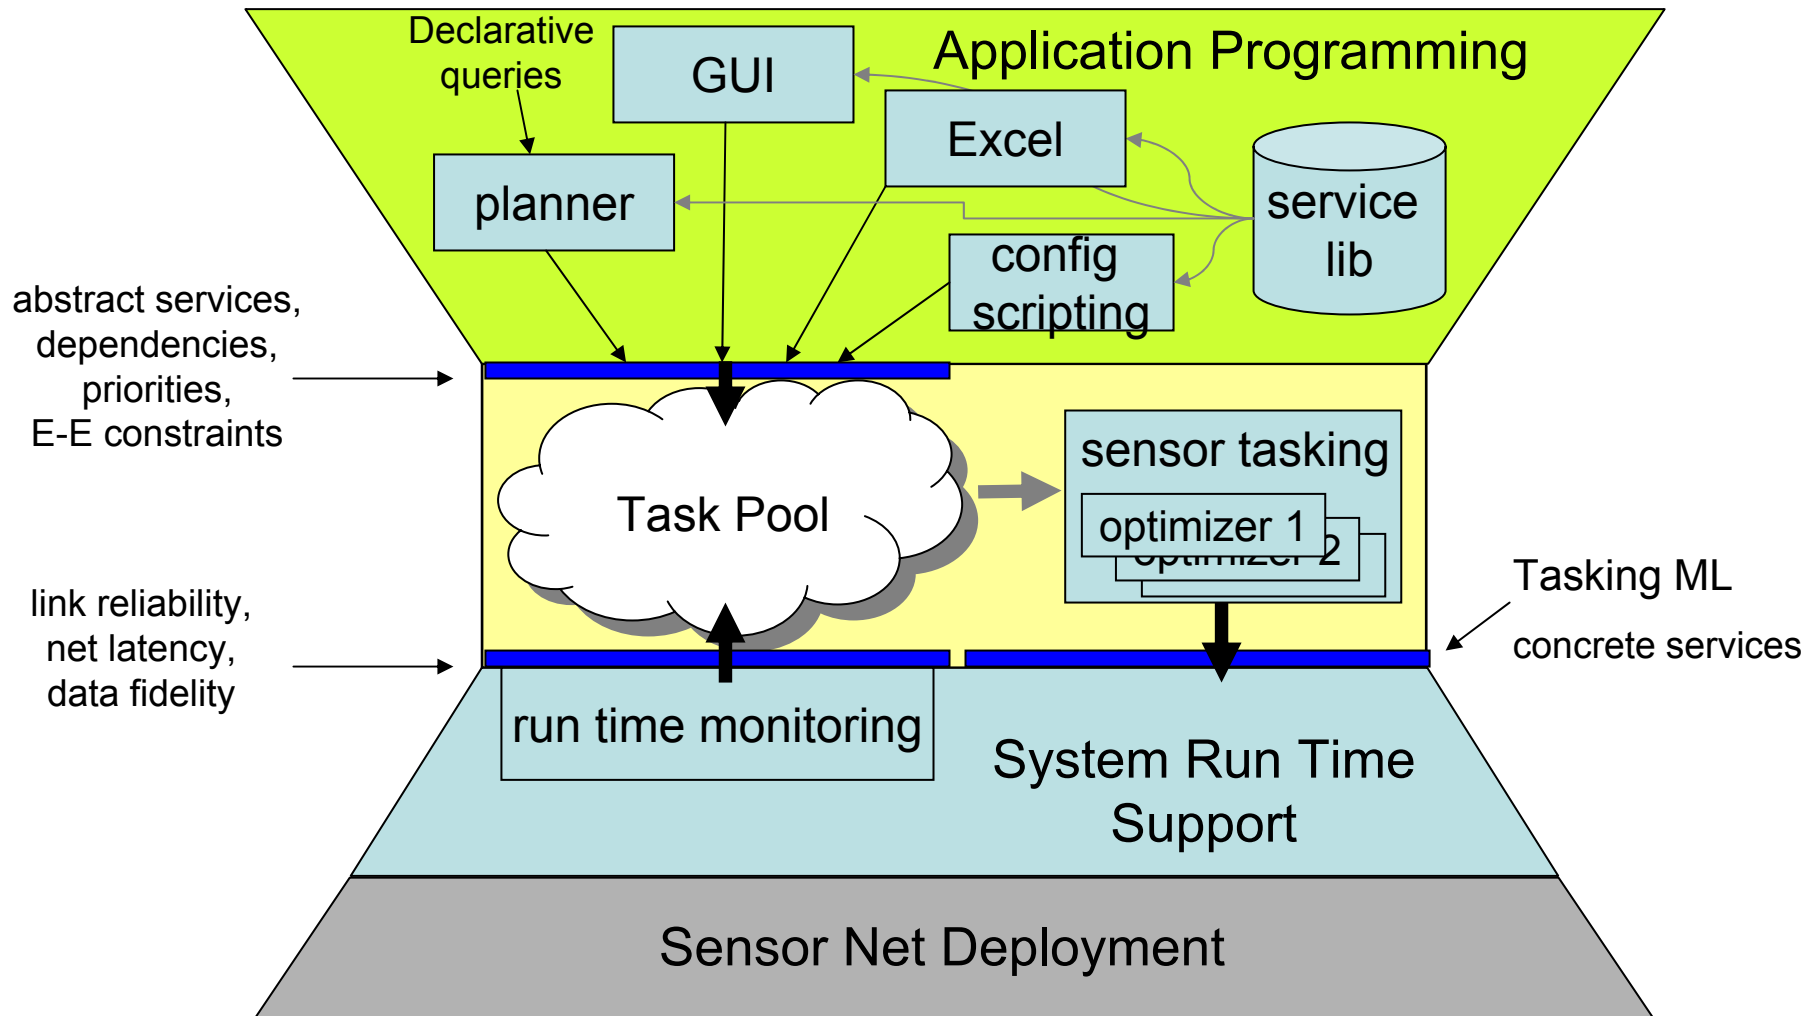
User interface

# Concurrent Uncoordinated Tasks




Tasks are sent to microservers at uncoordinated times, running for unpredictable duration. Tasks may partially overlap.

# Interfaces Between Apps and Run-Time



# Task Pool Abstraction




- Application programs inject abstract services and their dependencies into task pool, specifying what/how
  - Services have priorities
  - Applications have end-to-end constraints such as latency, data quality, energy usage
- Sensor Tasking embed services onto nodes, instantiate where/when
  - Need information about physical topology, link quality, latency, data fidelity from run time
  - Search for optimal assignment satisfying EE constraints
    - Using a variety of tools such as CLP(R), LP, or Monte Carlo
  - Tasking ML describes concrete services instances
- Task pool is agnostic to application programming environments and run-time system support, provided that
  - Services are described in a common intermediate language
  - Run-time provides system and data reliability info

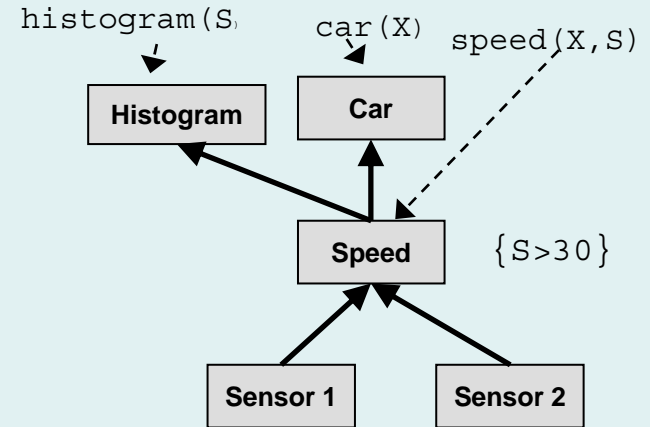


# SONGS: Service Oriented Networked ProGramming of Sensors

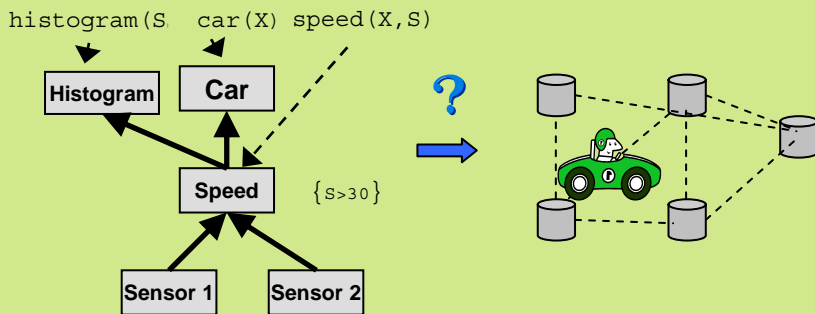
## Service Abstraction and Interface




## Service Planning



## Service Embedding



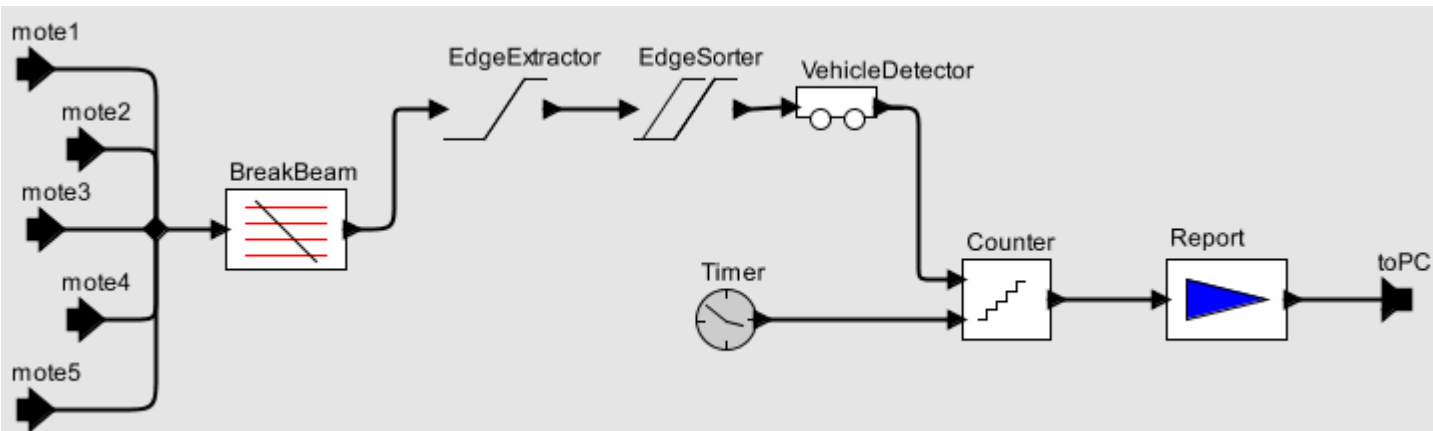

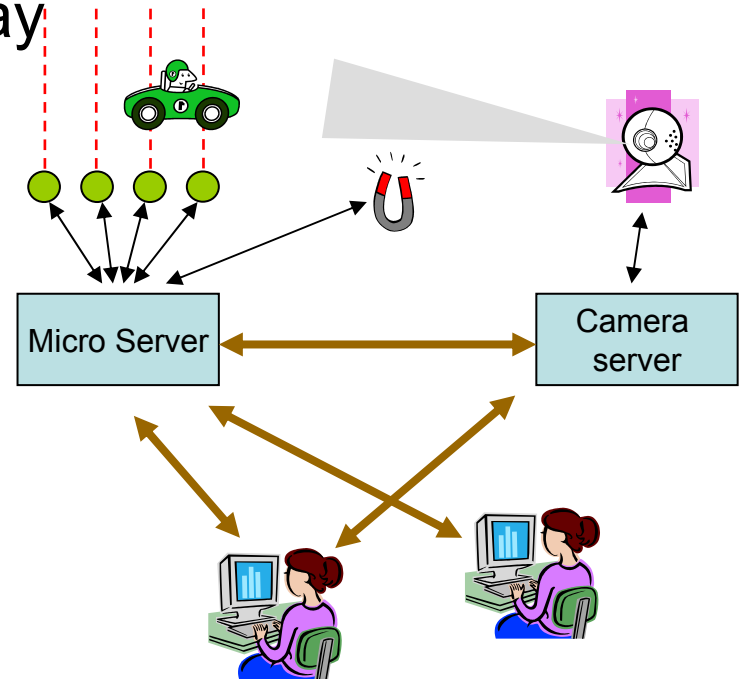
## Service Scheduling and Execution



# Example of services and their composition

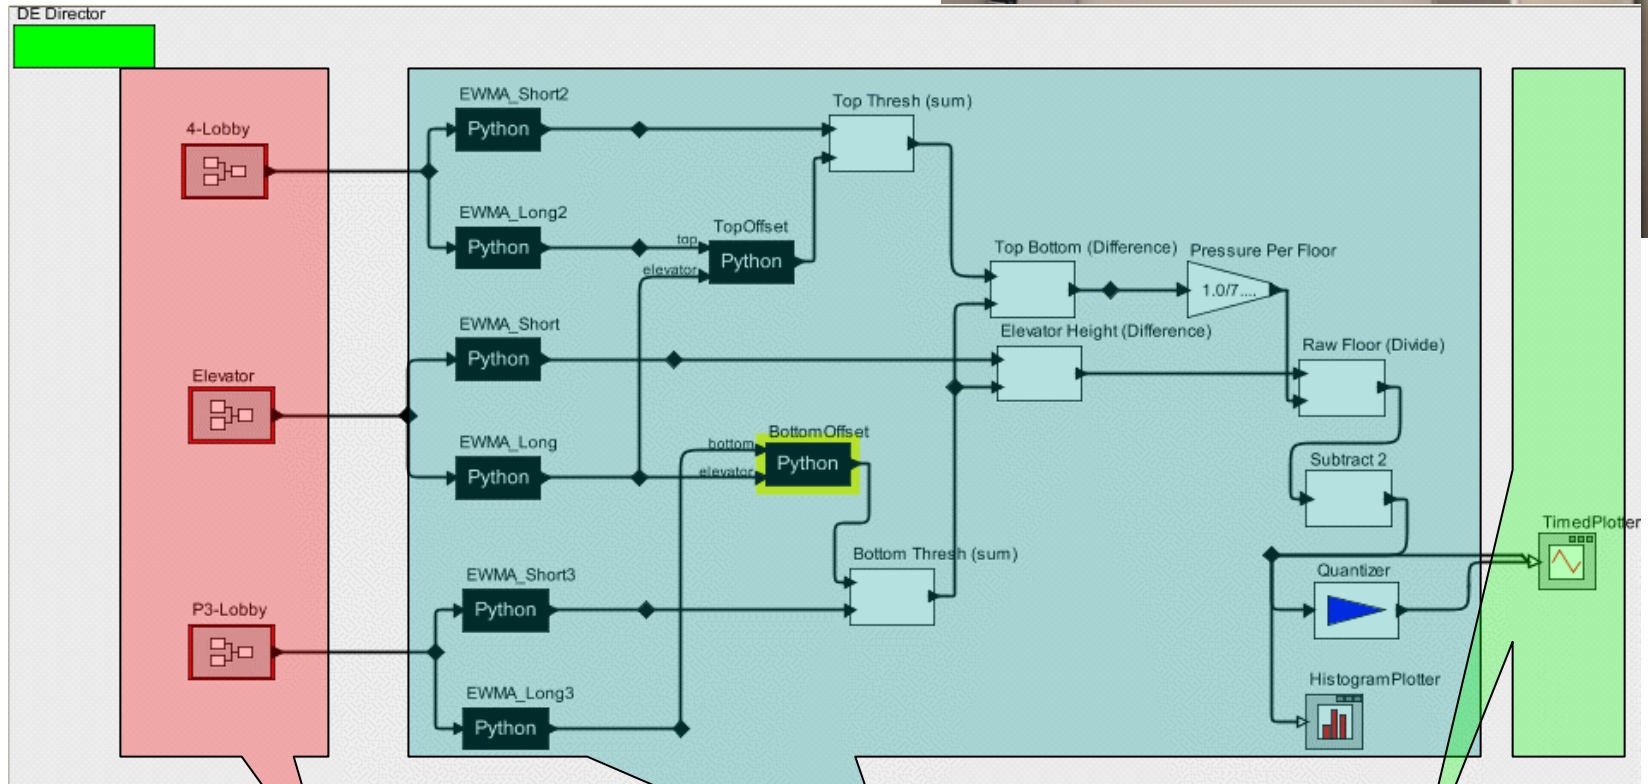
## Counting vehicles with a sensor array

- Extract edges from break beam detections
- Sort edges into consecutive detections
- Detect vehicles based on timing relations among detections
- Count vehicles
- Generate an arrival histogram report



# Elevator app

Jeremy Elson and Andrew Parker



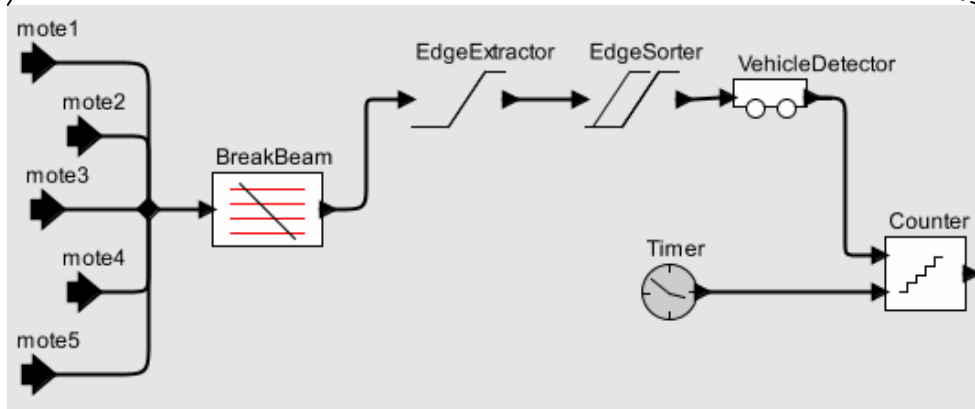

Data Sources

Intermediate Processing

Output

# An Example of TML

- Task graph for counting vehicle query
  - Ports
  - Services
  - Wiring services together
- Description in Tasking ML (micro-server tasking markup language)



```

<entity name="WebServiceComposite" class="microsoft.necomp.emws.EMWebServiceComposite">
  <property name="hostName" value="t-elaine" />
  <property name="portNumber" value="6000" />
  <port name="motel" type="AMHandler">
    <property name="input" />
    <property address="101:11" />
  </port>
  <port name="port2" type="Socket">
    <property name="output" />
    <property address="172.28.1.1" />
  </port>
  <port name="mote2" type="AMHandler">
    <property name="input" />
    <property address="102:11" />
  </port>
  <port name="mote3" type="AMHandler">
    <property name="input" />
    <property address="103:11" />
  </port>
  <port name="mote4" type="AMHandler">
    <property name="input" />
    <property address="104:11" />
  </port>
  <port name="mote5" type="AMHandler">
    <property name="input" />
    <property address="105:11" />
  </port>
  <entity name="SlowTimer" type="MicroserverTimer">
    <property name="rate" value="1000" />
  </entity>
  <entity name="BreakBeam" type="BreakBeamService">
    </entity>
  <entity name="EdgeExtractor" type="EdgeExtractorService">
    </entity>
  <entity name="EdgeSorter" type="EdgeSorterService">
    </entity>
  <entity name="VehicleDetector" type="VehicleDetectorService">
    </entity>
  <entity name="BreakBeam" type="BreakBeamService">
    </entity>
  <relation name="relation8" />
  <relation name="relation" />
  <relation name="relation2" />
  <relation name="relation5" />
  <relation name="relation3" />
  <relation name="relation6" />
  <relation name="relation7" />
  <relation name="relation4" />
  <link port="motel" relation="relation4" />
  <link port="mote2" relation="relation4" />
  <link port="mote3" relation="relation4" />
  <link port="mote4" relation="relation4" />
  <link port="mote5" relation="relation4" />
  <link port="BreakBeam.Input" relation="relation4" />
  <link port="Count.output" relation="relation" />
  <link port="Report.input" relation="relation" />
  <link port="SlowTimer.output" relation="relation2" />
  <link port="EdgeExtractor.Output" relation="relation3" />
  <link port="EdgeSorter.Input" relation="relation3" />
  <link port="EdgeSorter.Output" relation="relation6" />
  <link port="VehicleDetector.Input" relation="relation6" />
  <link port="VehicleDetector.Output" relation="relation7" />
  <link port="EdgeExtractor.Input" relation="relation7" />
  <link port="BreakBeam.Output" relation="relation8" />
</entity>
    
```

**<port name="mote4" type="AMHandler">  
 <property name="input" />  
 <property address="104:11" />  
 </port>**

**<entity name="BreakBeam" type="BreakBeamService">  
 </entity>**

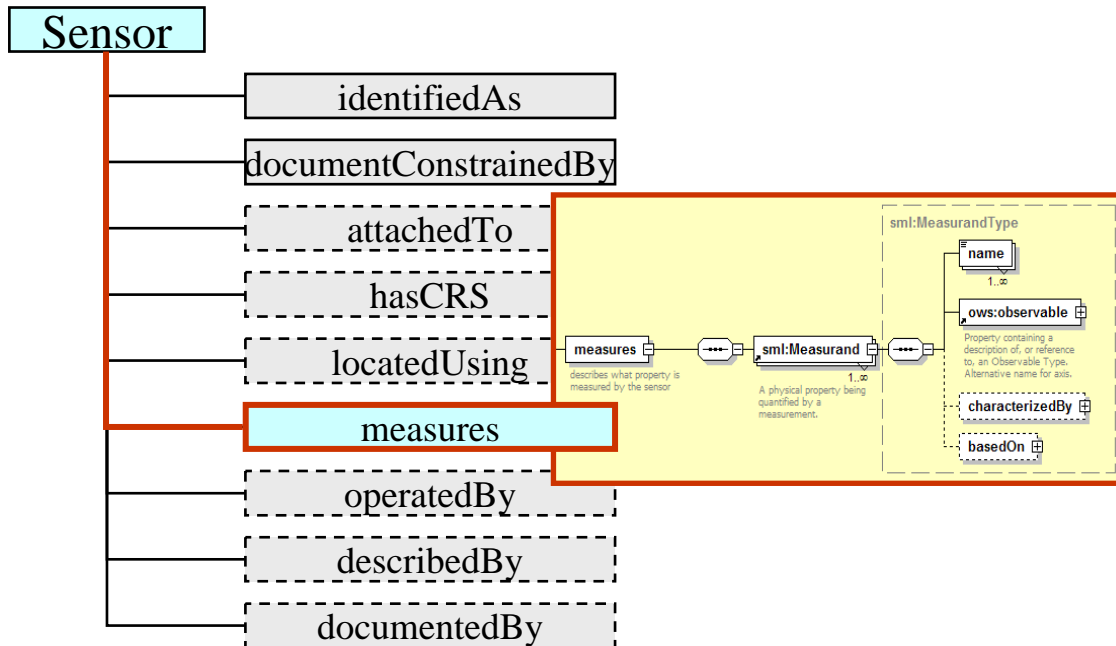
**<relation name="relation4" />**


**<link port="BreakBeam.Input" relation="relation4" />**

**Connecting services together**


# Standards for data/services/interoperation

- SensorML from Open Geospatial Consortium (<http://vast.nsstc.uah.edu/SensorML/>) ?
- Other interface standards?
  - Service description, data publishing, tasking






CLICK HERE



April 2003

Intelligent Systems



Connectivity

## A Sensor Model Language:

Moving Sensor Data onto the Internet

A new XML encoding scheme may make it possible for you to remotely discover, access, and use real-time data obtained directly from Web-resident sensors, instruments, and imaging devices.

Mike Botts, University of Alabama in Huntsville  
Lance McKee, Open GIS Consortium, Inc.

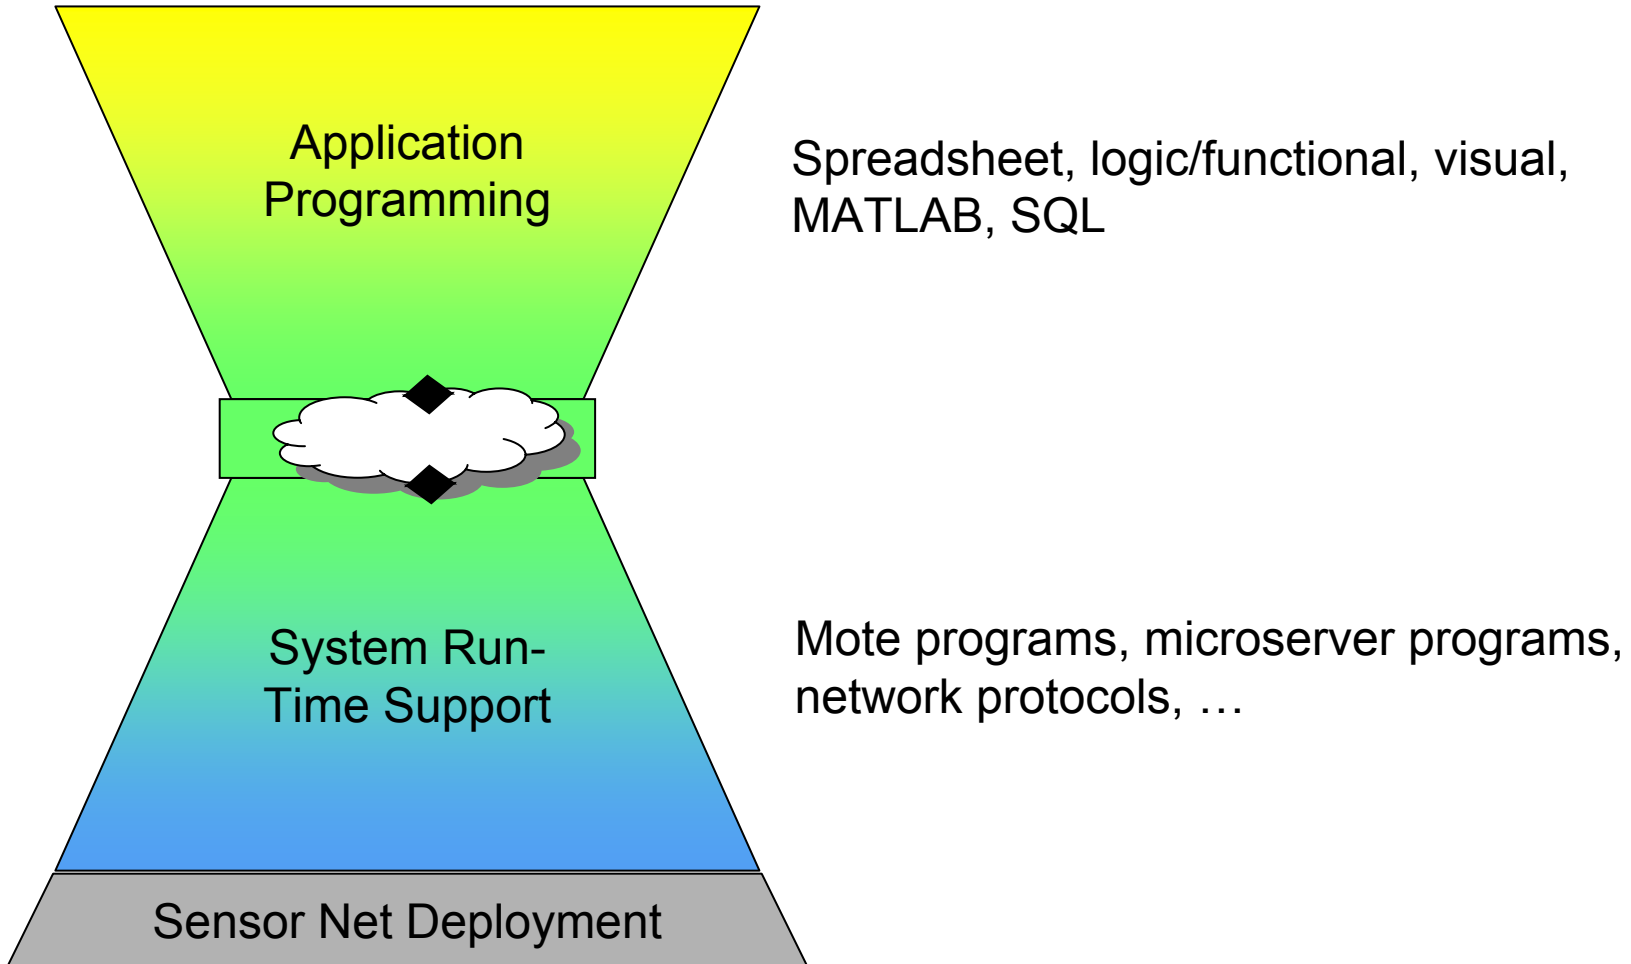
**M**embers of the Open GIS Consortium, Inc. (OGC), including NASA, the National Imaging and Mapping Agency, and EPA, are developing a standard XML encoding scheme for metadata describing sensors, sensor platforms, sensor tasking interfaces, and sensor-derived data (see [Editor's Note](#)). The goal is to make all types of Web-resident devices (e.g., flood gauges, stress gauges on bridges, mobile heart monitors, Web cams, and satellite-borne earth imaging devices) discoverable and accessible using standard services and schemas. The Sensor Model Language (SensorML) is a vital component that provides sensor information necessary for discovery, processing, and georegistration of sensor observations.

```

<MAP NAME="TOPNAV">
  <AREA SHAPE="Rect" COORDINATE="10 10 100 100">
  <AREA SHAPE="Rect" COORDINATE="100 10 100 100">
  <AREA SHAPE="Rect" COORDINATE="10 100 100 100">
  <AREA SHAPE="Rect" COORDINATE="100 100 100 100">
</MAP>
<MAP NAME="CONTACT">
  <AREA SHAPE="Rect" COORDINATE="10 10 100 100">
  <AREA SHAPE="Rect" COORDINATE="100 10 100 100">
  <AREA SHAPE="Rect" COORDINATE="10 100 100 100">
</MAP>
<MAP NAME="COPYRIGHT">

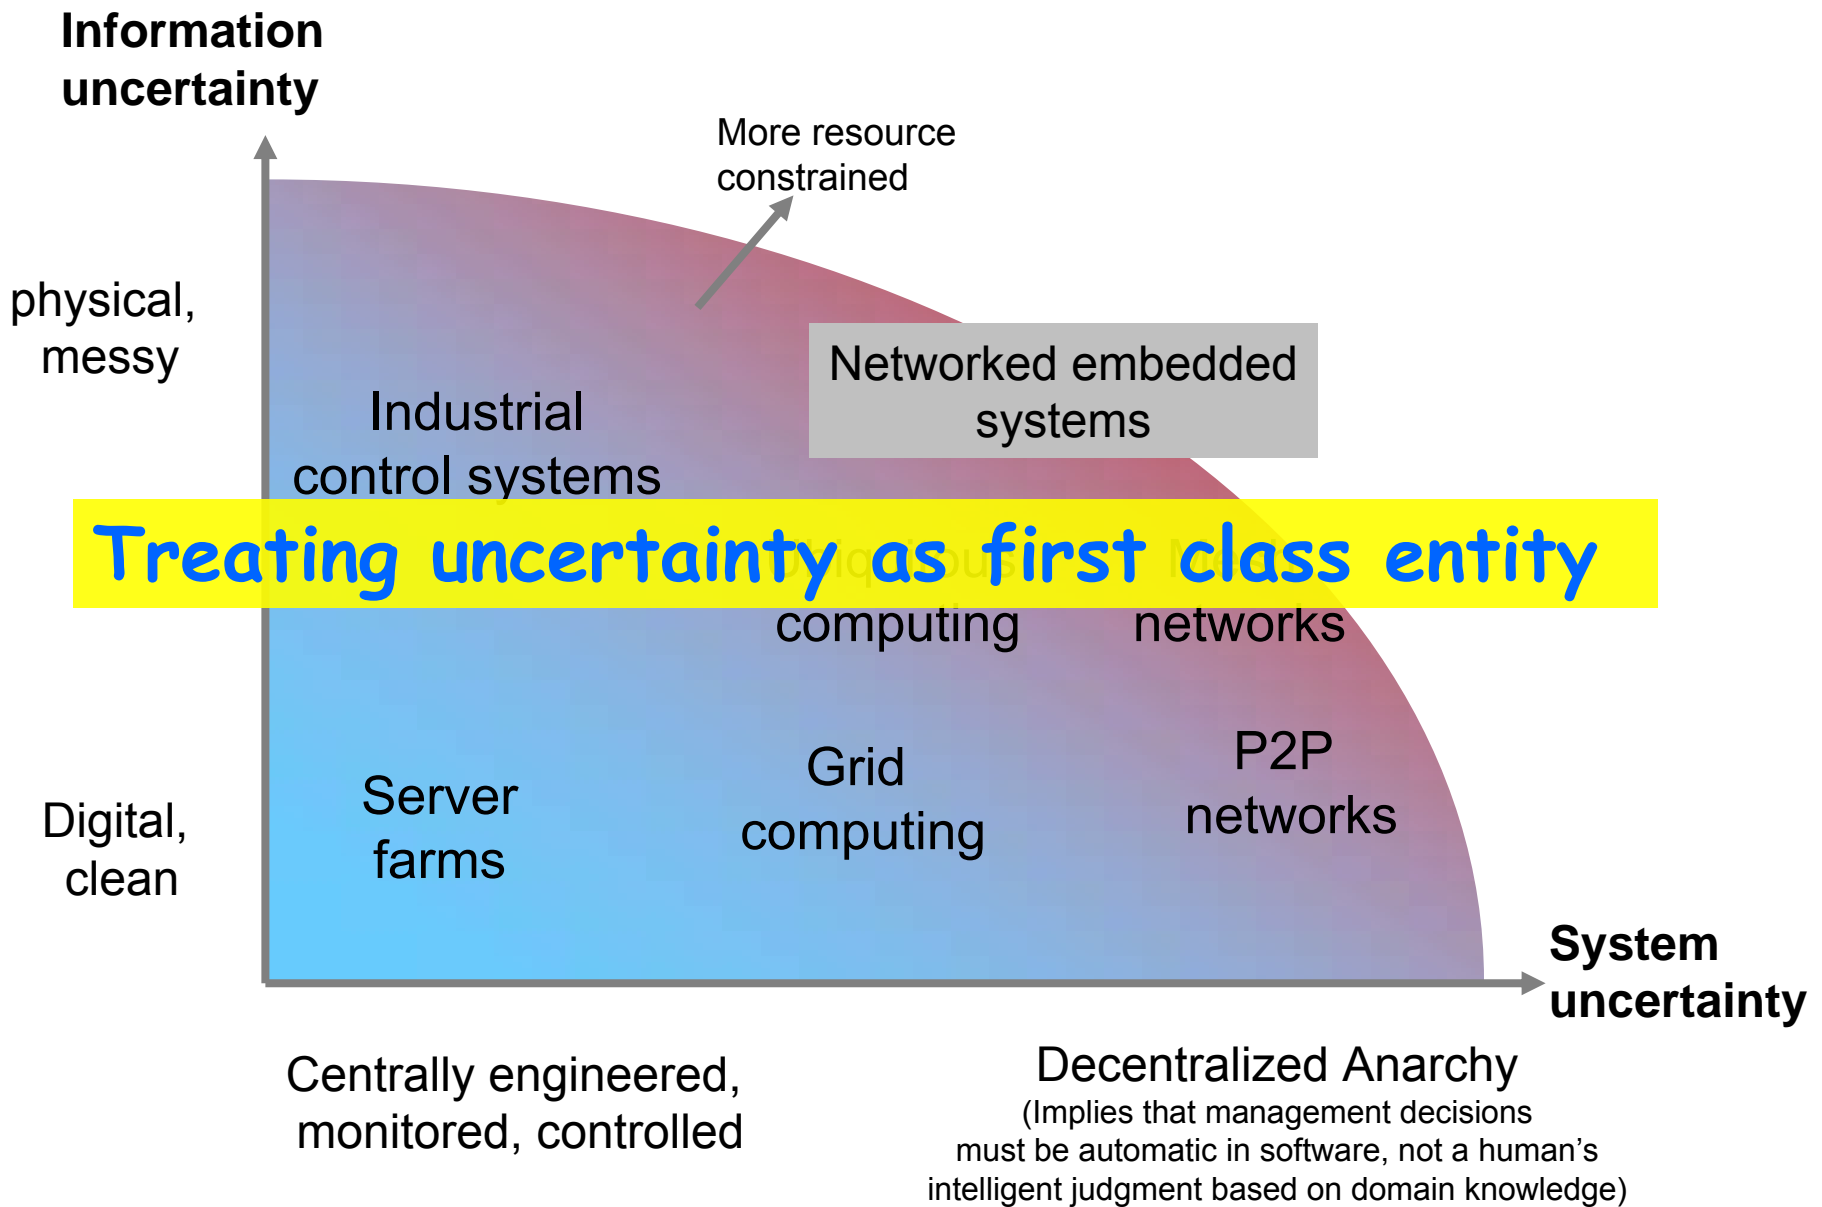
```

Narrow waist should allow applications to be specified independent of system configurations



# How to engineer reliable behaviors out of unreliable components?

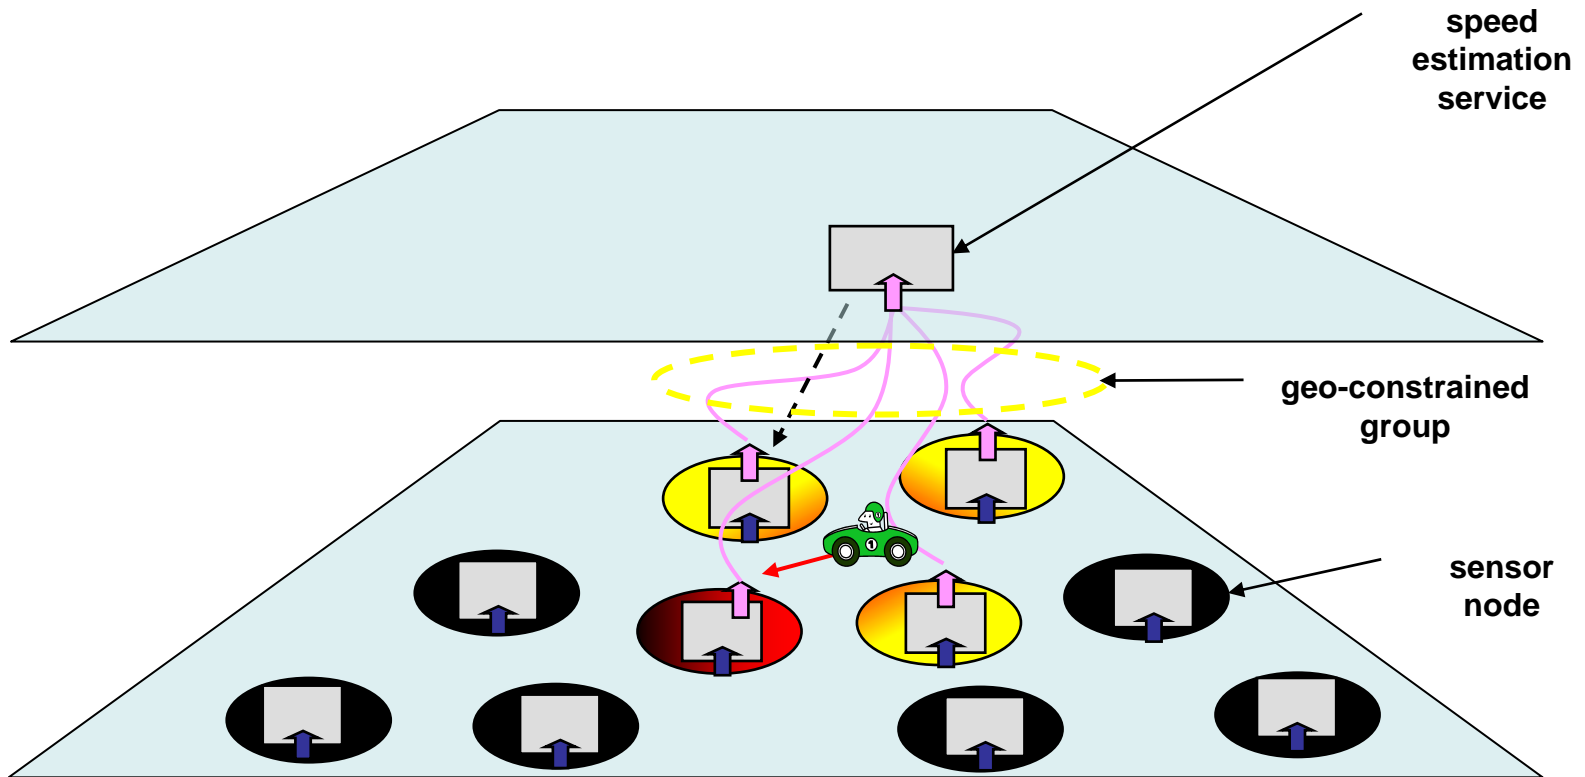
- Sensor net today: reliable components, unreliable networks/systems, and unpredictable apps
- Biology: unreliable components, reliable systems, predictable behaviors





# Manage both system and event uncertainty


## A tracking example



# Migration ...

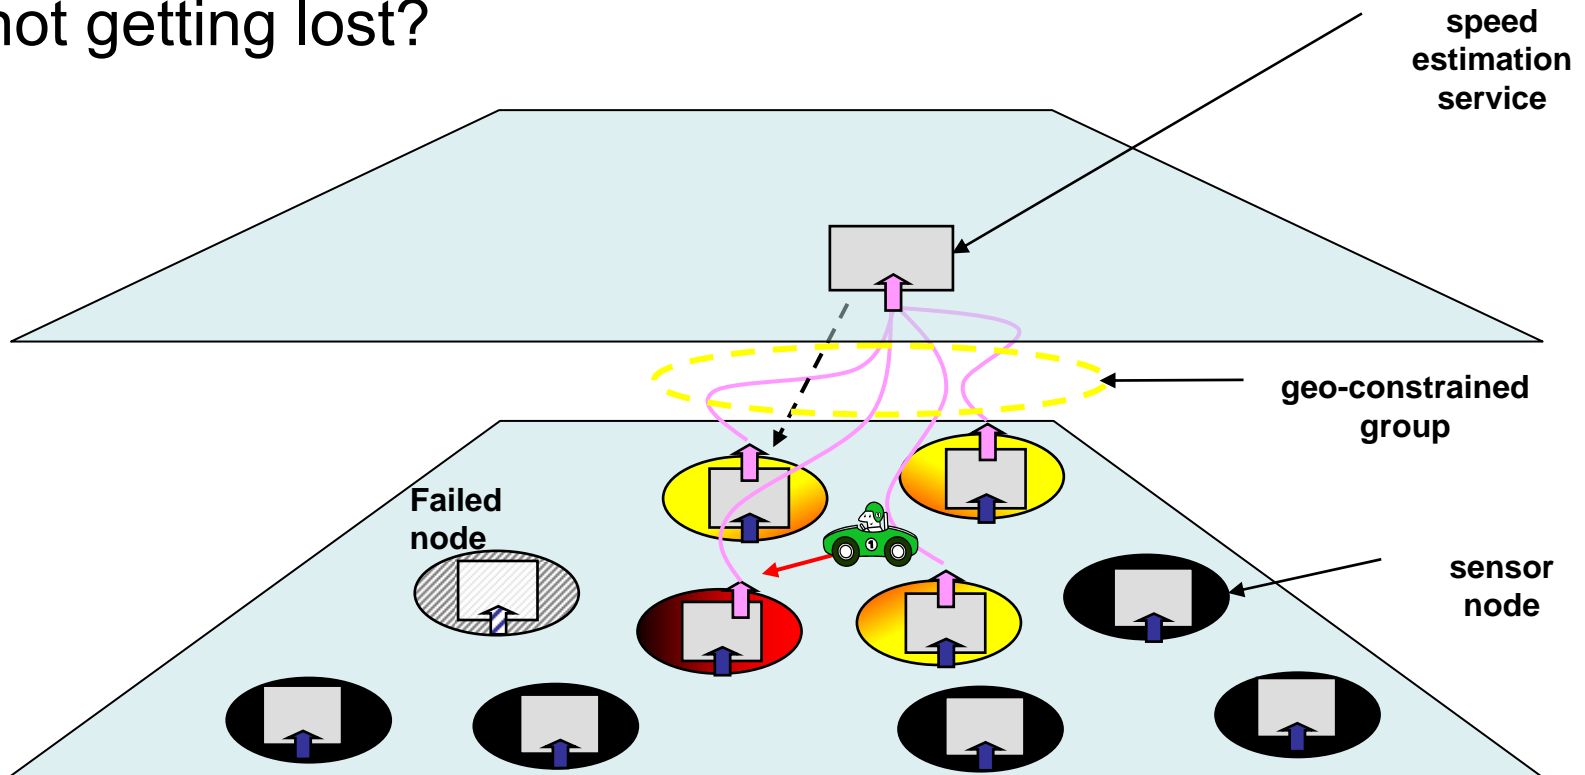
As the event moves, one needs to move the code and/or state

speed  
estimation  
service



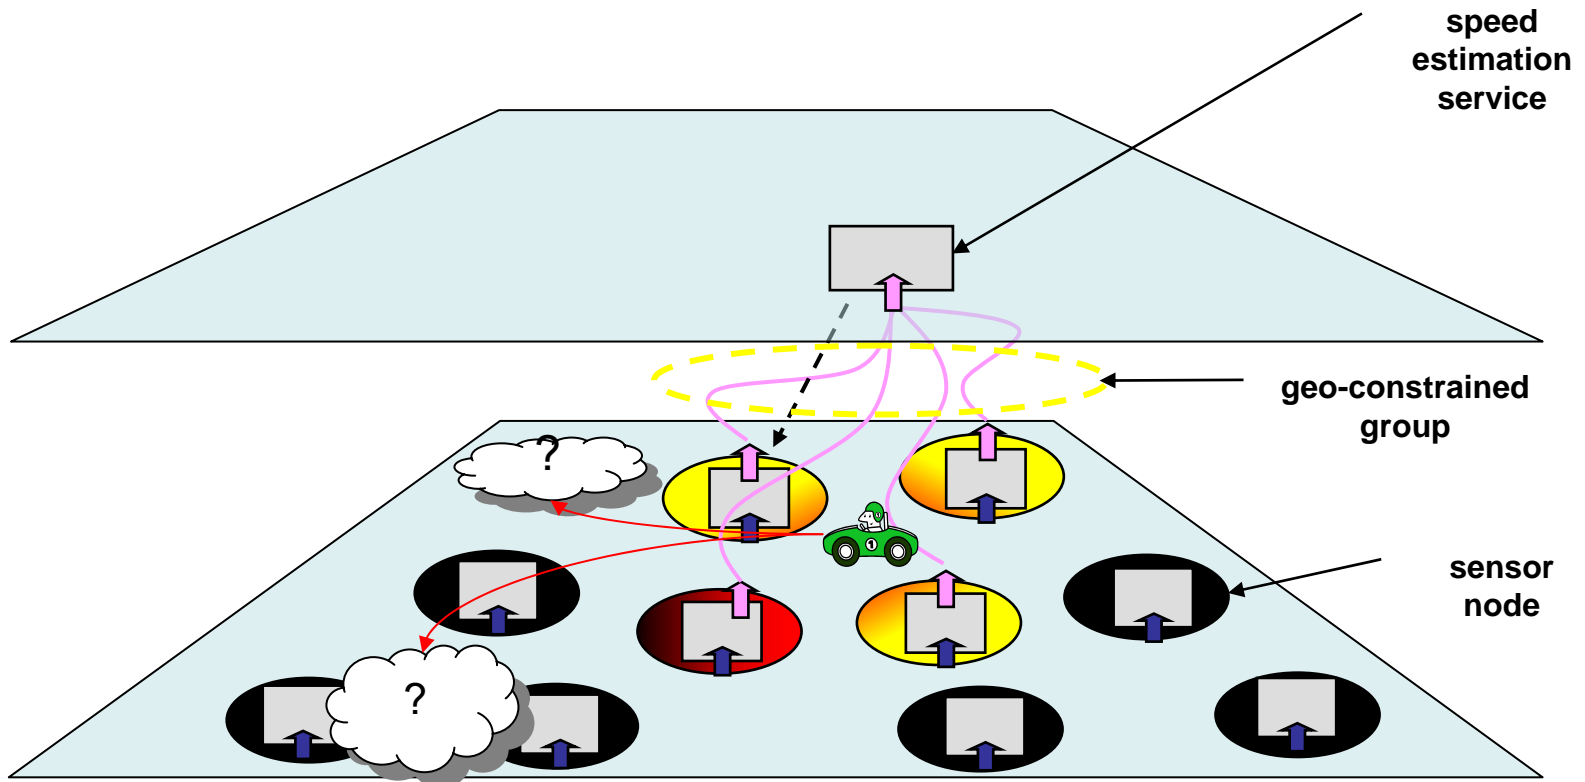
# What if nodes fail?

- Contingency plan?
- How many copies do I send around, to maximize odds of not getting lost?



# ... and world knowledge is incomplete?

- Follow the clouds
- Probabilistic tasking?
- Proactive or reactive?

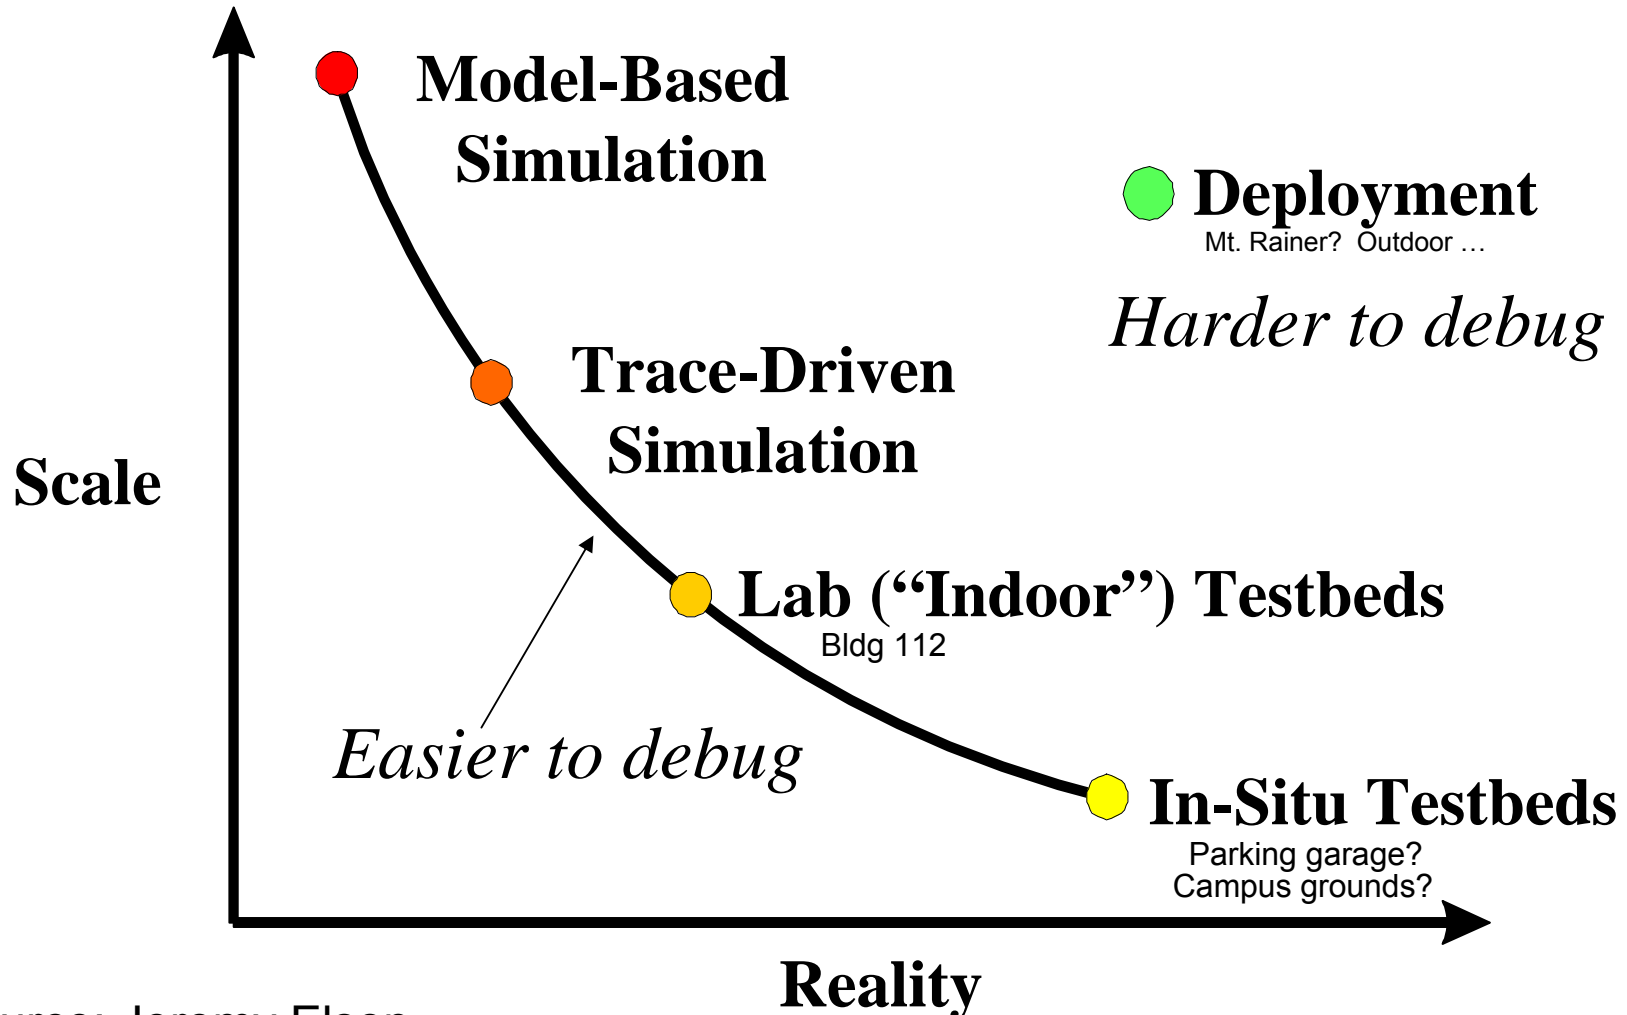


# Manage system and event uncertainty

- Event uncertainty
  - Build predictive models of events (e.g., probabilistic)
- System uncertainty
  - Build models of system behaviors (e.g., component failure, processing latency, link quality)


# Visibility to systems

*A spectrum allows high-visibility debugging before jumping into low-visibility deployment*



# Data Collection

Archiving Events



**DataBase**  
(MS Access / SQL Server 2005)

SQL Query / Report

Raw Data + Processed Data

XML packets

Microserver Tasking

Raw Data Streaming

Status / Sensor Readings  
(TinyOS Packets)




**Gateway** (MicroServer)

Visualize Events/ Process Data




**User Interface / Data Processing**  
(MS Excel)



**Sensor Net** (Tmote Sky)




Mote (Id = 30)




Cold air from AC



Hot air from Laptop heat sink



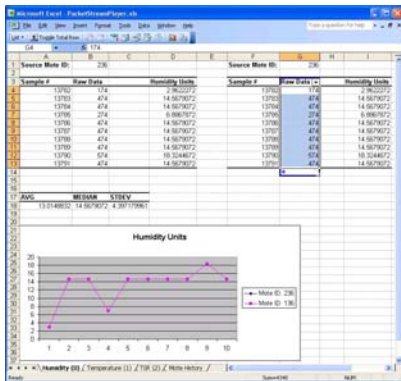
Mote (Id = 40)





# Features

- Excel 2003
  - Worksheets
  - Xml Maps
  - Cell Functions
- Packet Stream Player
  - Familiar, simple interface for streaming data
  - Similar to other media-centric players, i.e., Connect, Play, Record, Next, Previous, etc.
- Packet Database
  - Session data
  - Packet data
- Microserver
  - Data provider



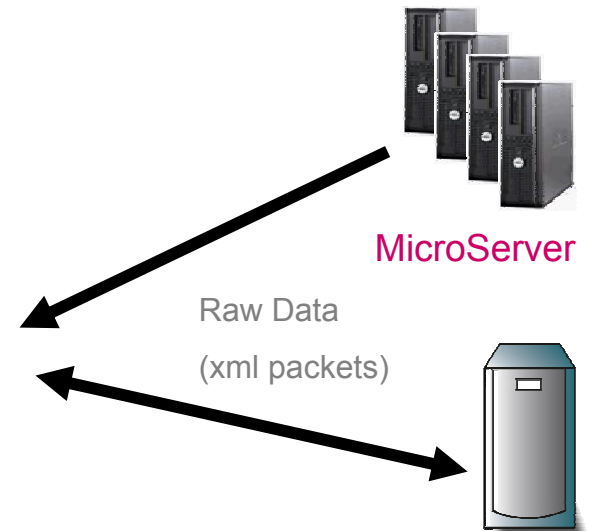
Excel



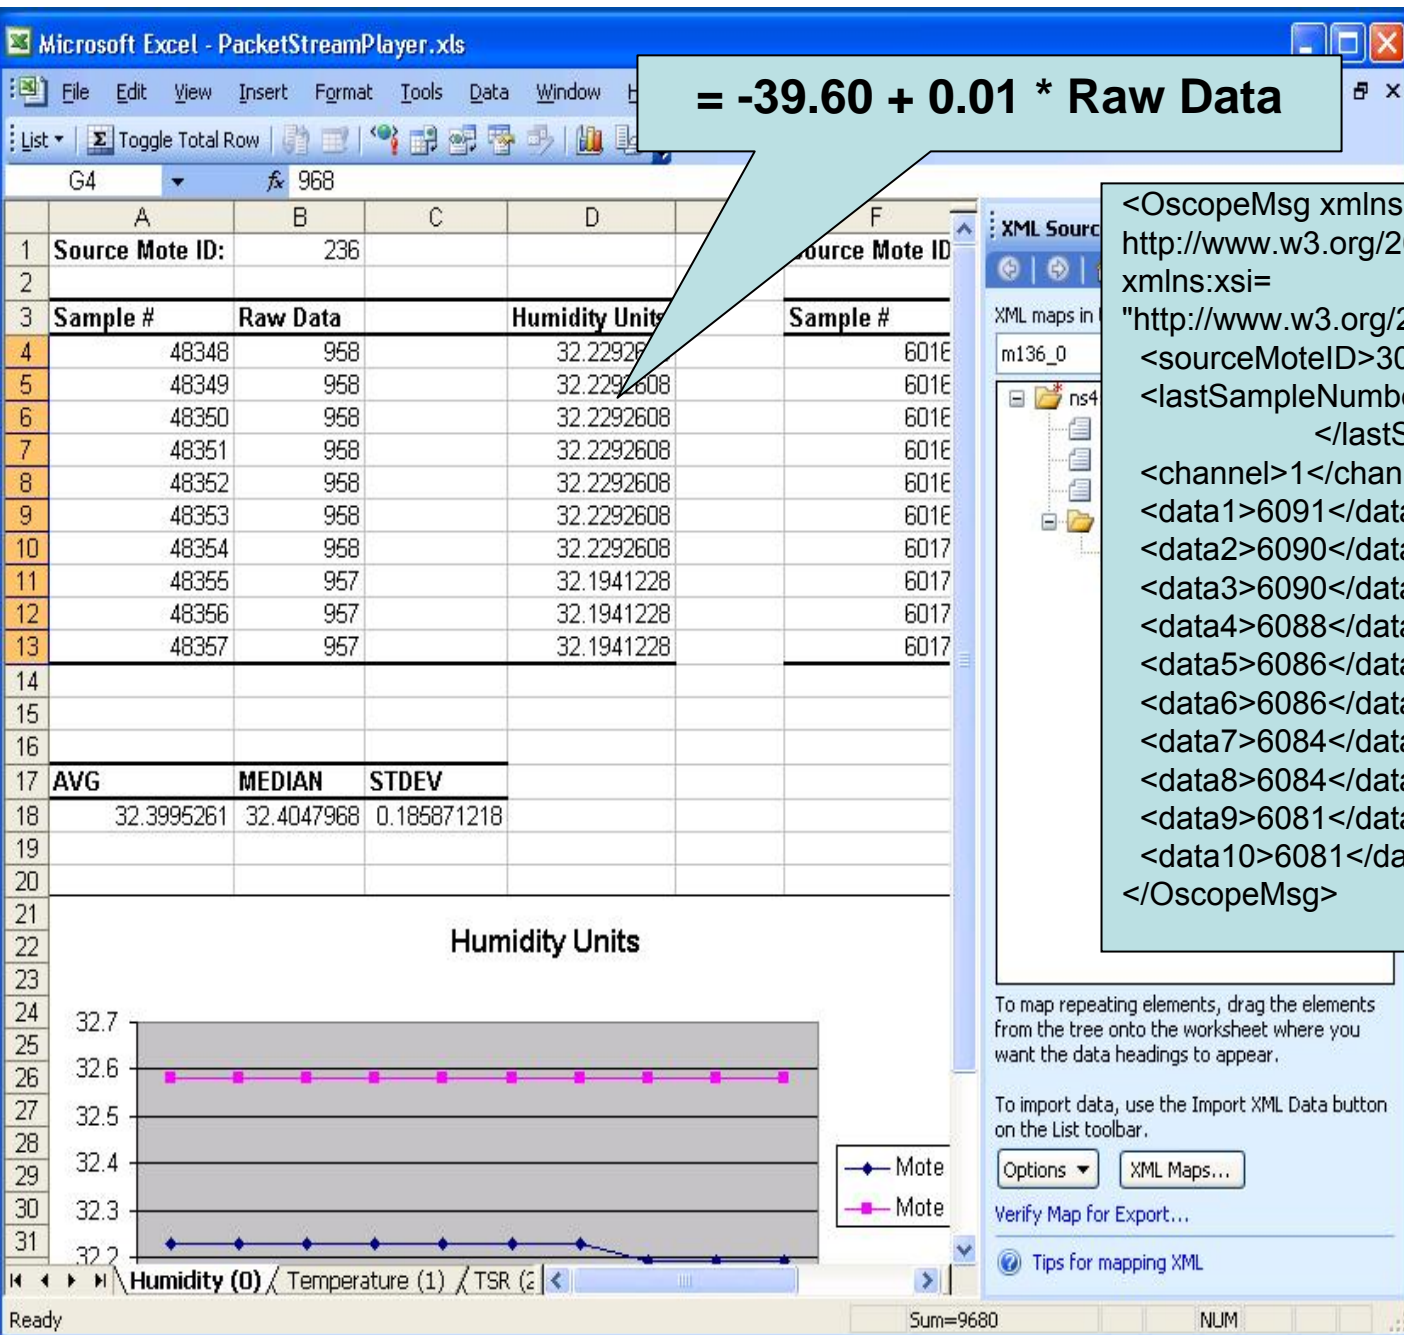
Transformed xml



Packet Stream Player



SQL 2005



$$= -39.60 + 0.01 * \text{Raw Data}$$

```
<OscopeMsg xmlns:xsd=
http://www.w3.org/2001/XMLSchema
xmlns:xsi=
"http://www.w3.org/2001/XMLSchema-instance">
  <sourceMoteID>30</sourceMoteID>
  <lastSampleNumber>40550
    </lastSampleNumber>
  <channel>1</channel>
  <data1>6091</data1>
  <data2>6090</data2>
  <data3>6090</data3>
  <data4>6088</data4>
  <data5>6086</data5>
  <data6>6086</data6>
  <data7>6084</data7>
  <data8>6084</data8>
  <data9>6081</data9>
  <data10>6081</data10>
</OscopeMsg>
```

To map repeating elements, drag the elements from the tree onto the worksheet where you want the data headings to appear.

To import data, use the Import XML Data button on the List toolbar.

Options XML Maps...

Verify Map for Export...

Tips for mapping XML

# Our tools will be available

- In source code from <http://research.microsoft.com/>
- Oct. 2005 (alpha)
  - **Microserver execution environment ( $\mu$ SEE v0.1)**
    - On Windows XP or XP Embedded, .NET Framework 1.1
    - MSTML 1.1 based on MoML
    - Interfacing with MicaZ or Telos (802.15.4 radio packets)
    - Per-node tasking
    - No runtime service sharing
  - **Microserver Interaction Console ( $\mu$ SIC v0.1)**
    - Generic microserver configuration
    - MSTML editing and data collection
    - Runs locally on the microserver machine
  - **Service composition GUI based on Ptolemy II**
    - Released as a patch to Ptolemy II 0.5
  - **Excel interface for data collection and visualization**
    - Require Visual Studio 2005 BETA and .NET Framework 2.0
    - Require SQL Server 2005 BETA
- Dec. 2005 (beta)
  - Move to .NET 2.0, VS 2005, SQL2005 when they are final.

# Wrap up: Challenge Problems

- As a community, we need grand challenge problems
- Three focus areas:
  - A community effort to standardize **interfaces** between various sensornet components/layers
  - Build models for **uncertainty** that can be utilized by data and system management
  - Create more **tools**, for sys config/mgmt, data collection and vis, debugging

# To probe further

- Some of the relevant papers on the SONGS architecture
  - "[Semantic Streams: a for Declarative Queries and Framework Automatic Data Interpretation.](#)" Tech Report MSR-TR-2005-45, Microsoft Research, April 2005 (Also Sensys'05 Poster)
  - "[Semantics-Based Optimization Across Uncoordinated Tasks in Networked Embedded Systems,](#)" EMSOFT 2005, Sept 2005
  - "[Towards Semantic Services for Sensor-Rich Information Systems,](#)" Basenets 2005, October 2005
- Conferences/Journals
  - ACM Sensys05, San Diego, Nov 2005
  - ACM/IEEE IPSN06, Nashville, April 2006
  - IEEE DCOSS, SECON, EWSN, ...
  - New ACM Trans. Sensor Networks,  
<http://www.acm.org/tosn>
- MSR Project: <http://research.microsoft.com/nec>