Challenges in designing information processing sensor networks

Feng Zhao

Palo Alto Research Center
www.parc.com/zhao

Starting March 1, 2004:
Microsoft Research
fz@alum.mit.edu

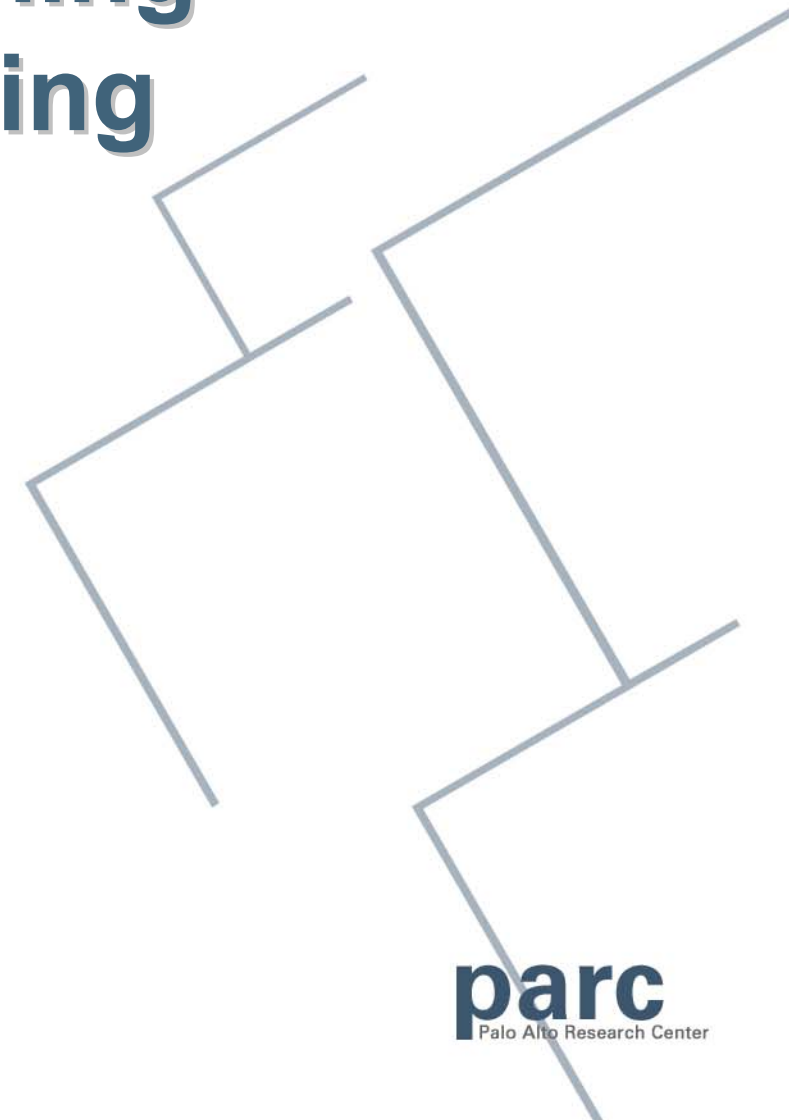

parc
Palo Alto Research Center

Collaborative processing in sensor networks

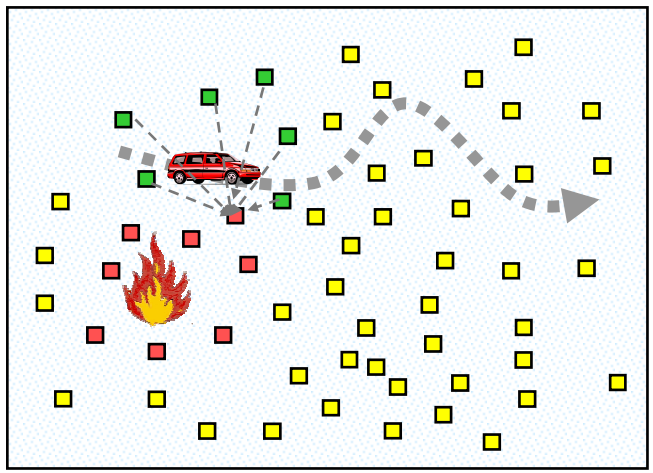

- What information to gather and communicate? And how often?
- Which nodes should participate in sensing, processing, or communication?
- How should the information be migrated?
- What is routing or querying in this context?

Moving from complexity-centric view of algorithm design to utility-based design:

- What information is critical for the high-level tasks?
- What is the cost of accessing the information?

Collaborative processing

– Group formation in sensor networks

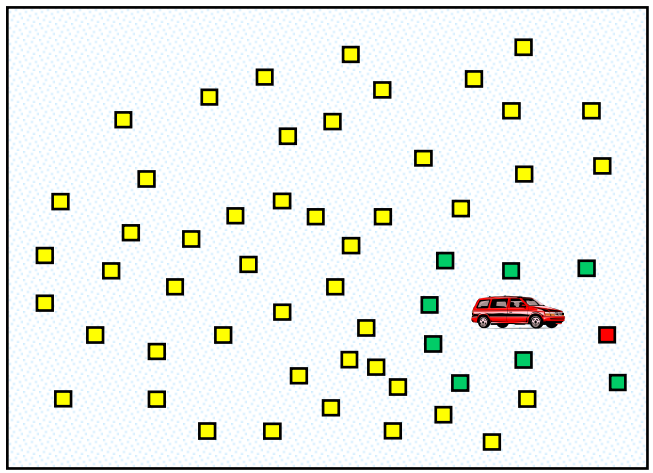

- Information needs and resource constraints define who should participate in the processing groups
- Group membership (e.g. location) defines the behavior of a node

parc
Palo Alto Research Center

Information-based group formation

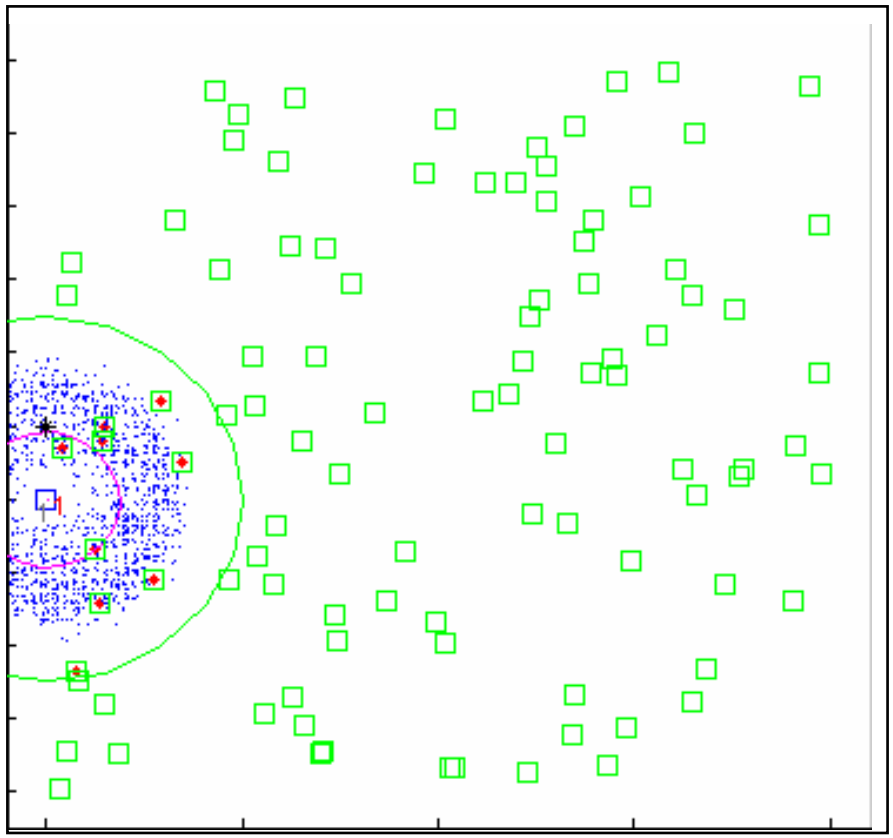

A leader node (blue square) carries belief state

- Choose sensor in the neighborhood with good information
- Hand off current belief to chosen sensor (new leader) and update

An example of collaborative groups

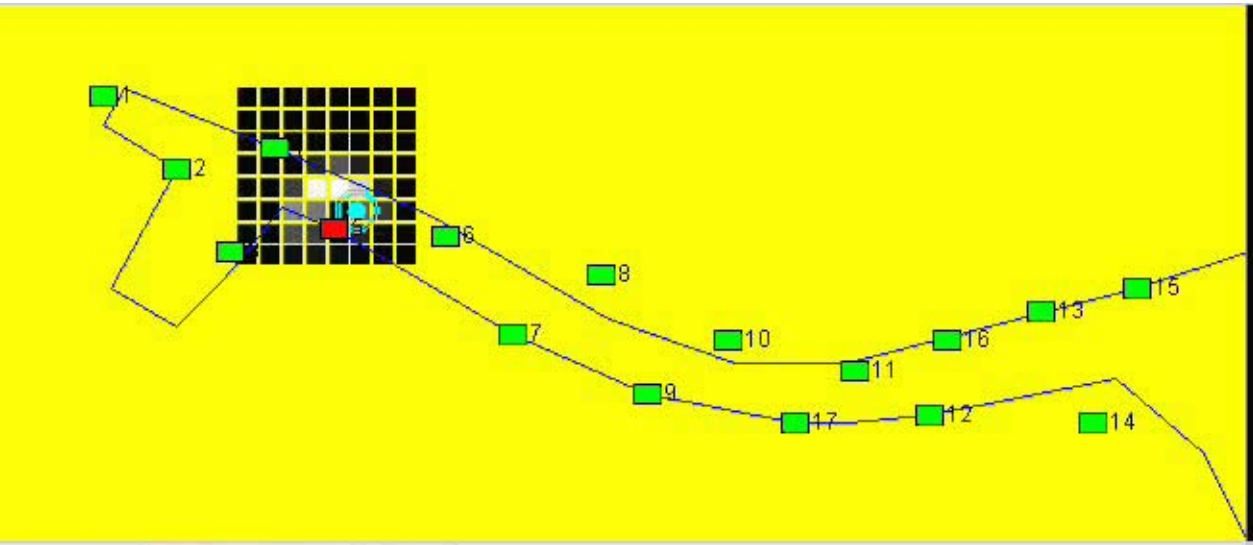

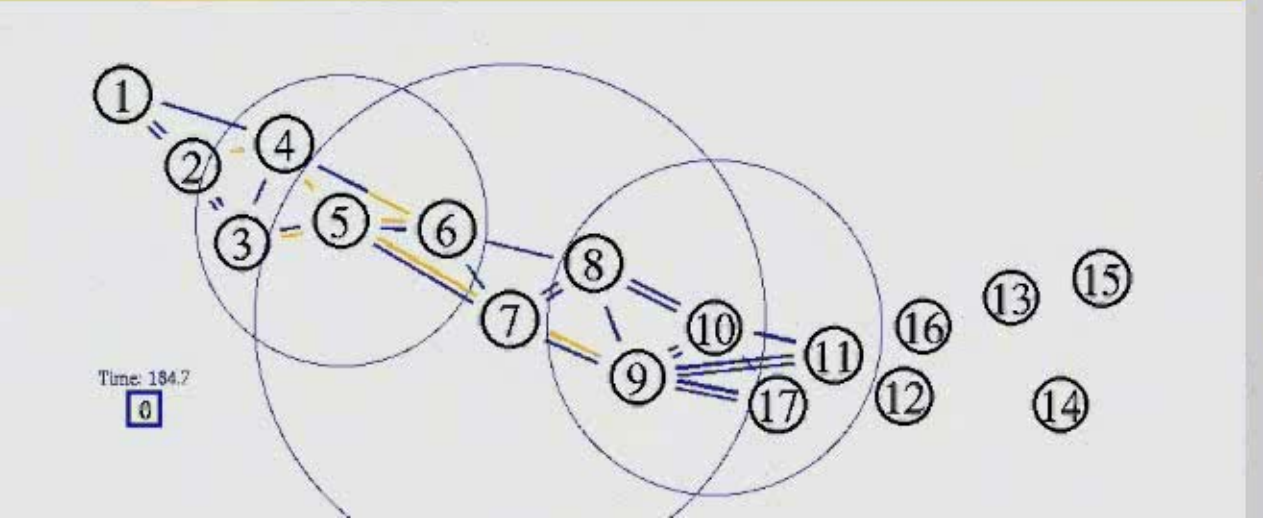

Leader Node

Non-leader Node

GPS ground truth

parc
Palo Alto Research Center

ISI

Data Packet

Control Packet

Broadcast control

Cross-layer interactions

In sensor networks, networking is intimately coupled with sensing, interaction, and control needs and hence application semantics

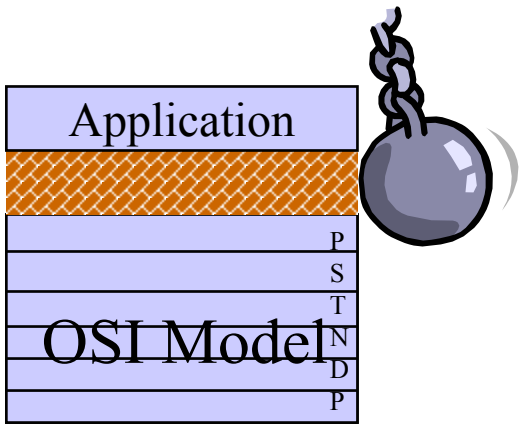

- Break down traditional barriers of network stack
  - Consider both communication cost *and* application requirements to plan routes and task sensors
- Data-centric and ad-hoc
  - Address nodes based on geography and capability, not by name
- Group management vital to scalability
  - Limit data propagation to sensors relevant to measurement at hand

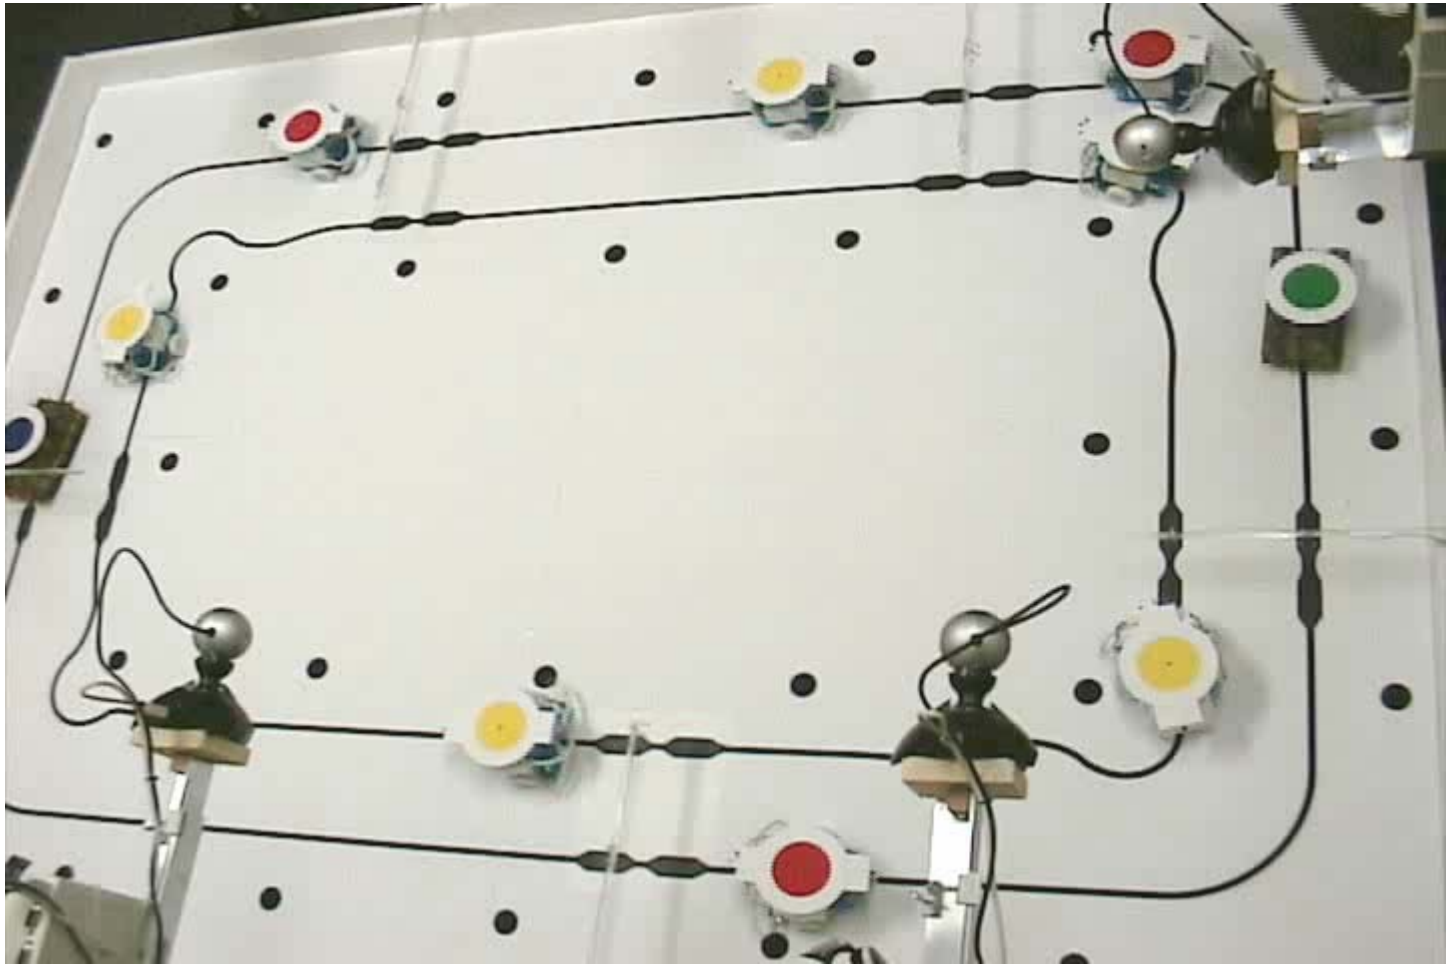

What objects to pay attention to? What to ignore? Which aspects of the objects (e.g., location, identity) are of interest to the high-level tasks?

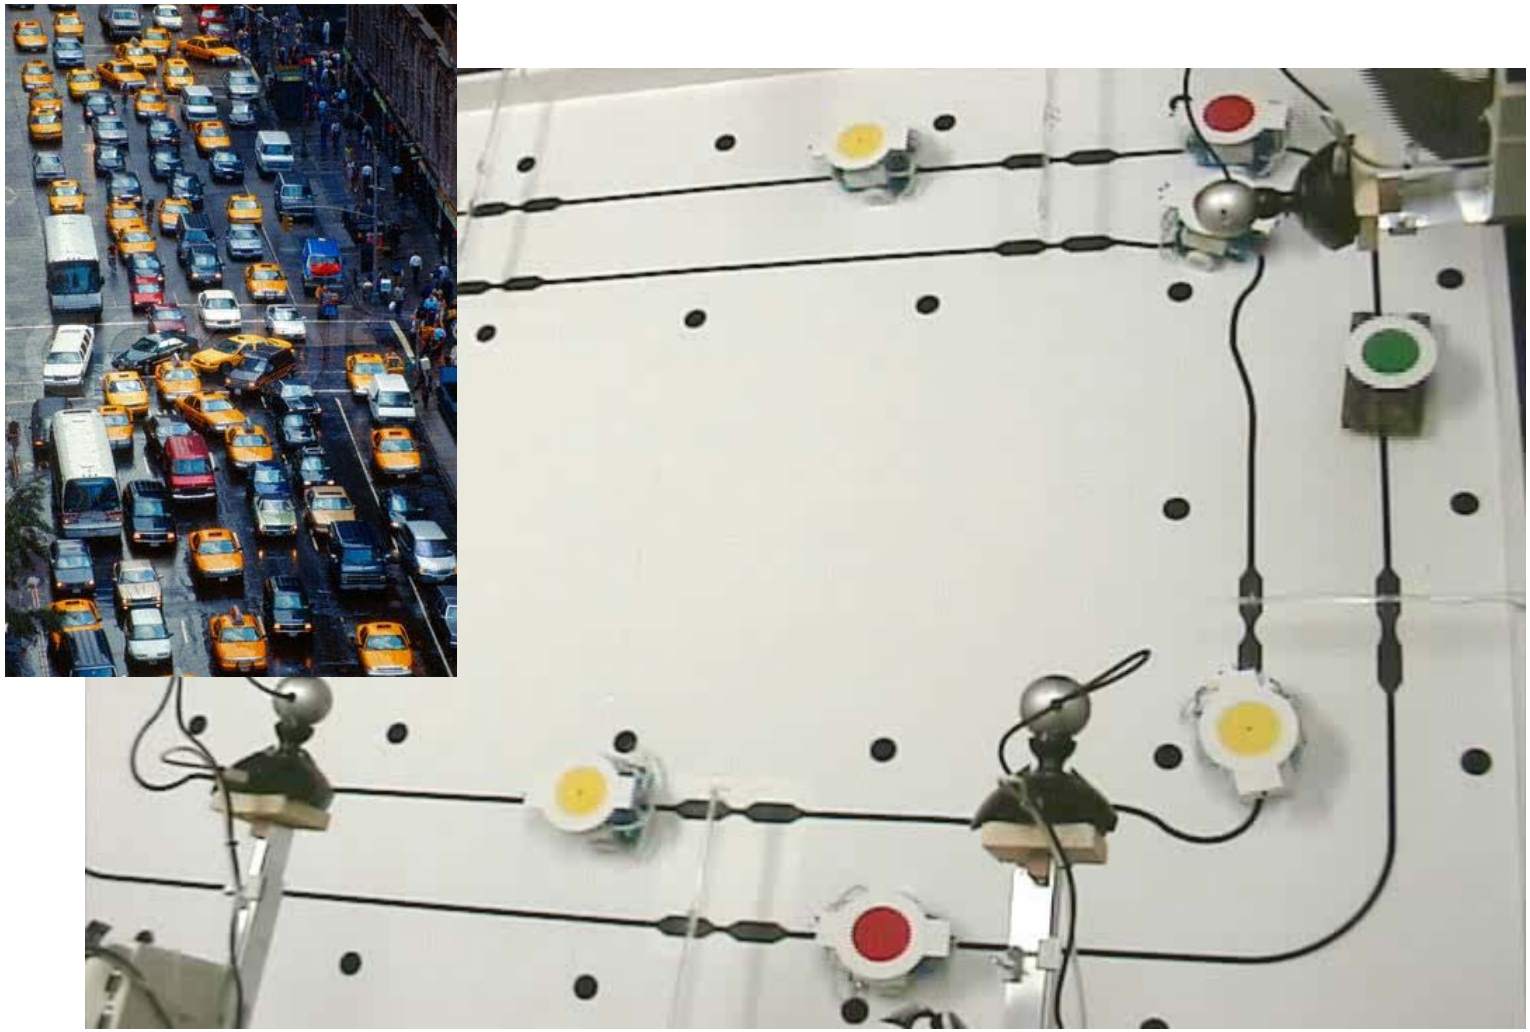

What objects to pay attention to? What to ignore? Which aspects of the objects (e.g., location, identity) are of interest to the high-level tasks?

parc
Palo Alto Research Center

Mixing of events

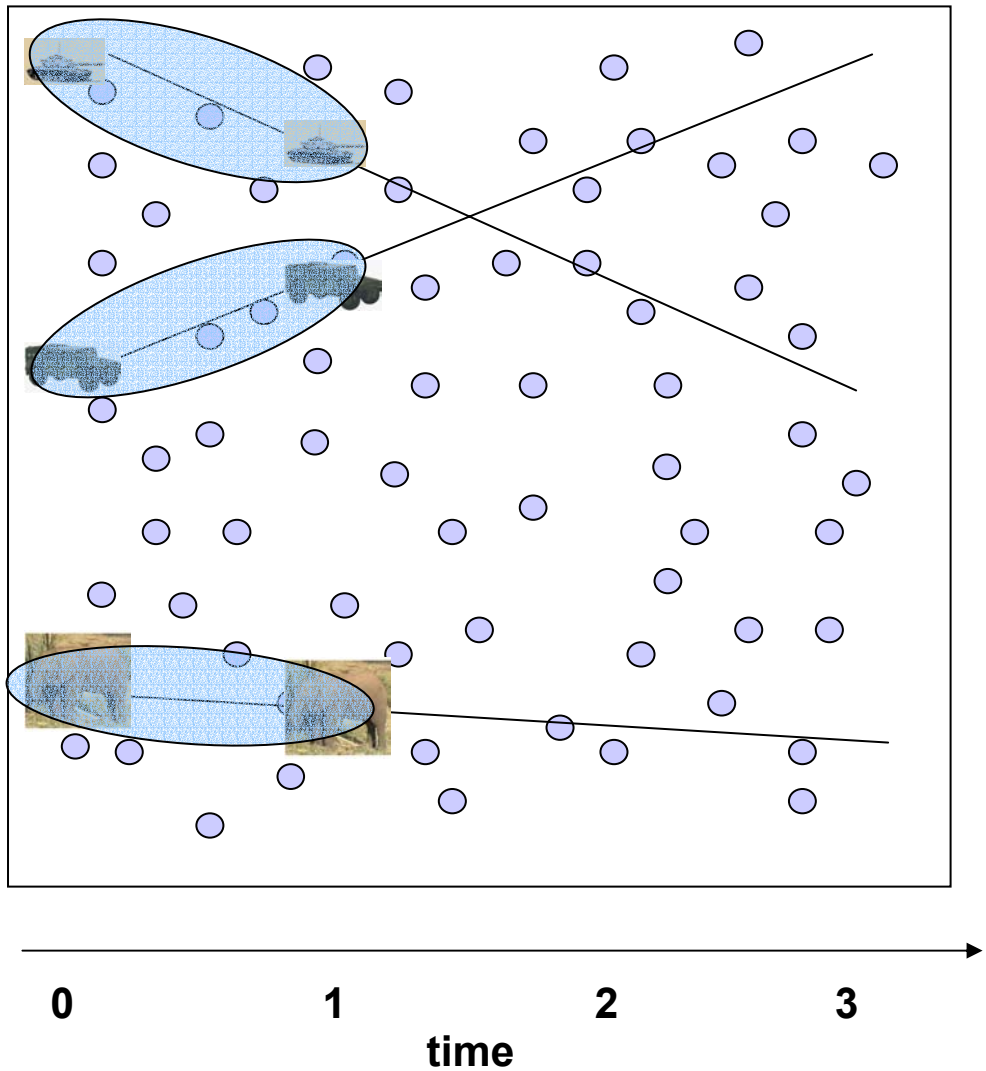

parc
Palo Alto Research Center

Mixing of events

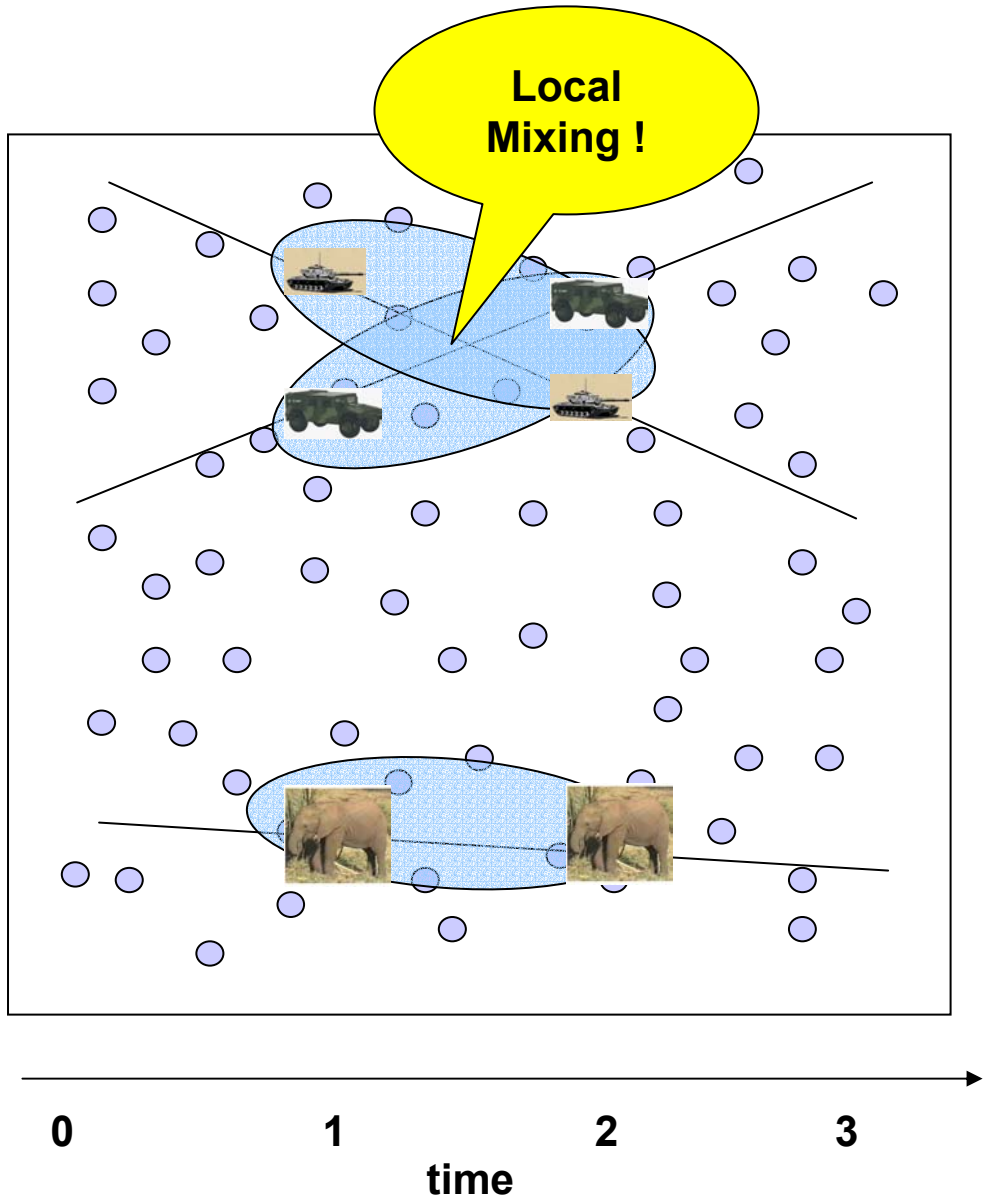

parc
Palo Alto Research Center

Mixing of events

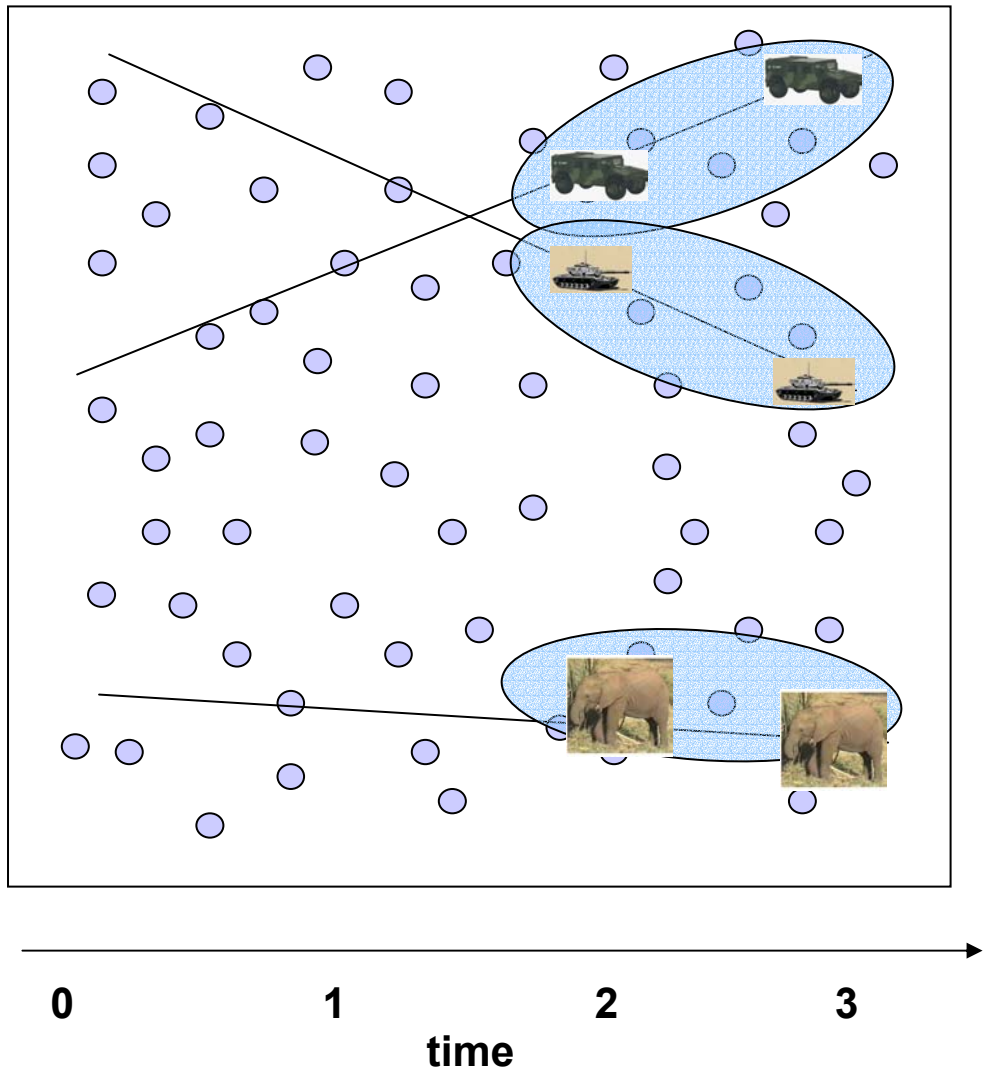

parc
Palo Alto Research Center

Mixing of events

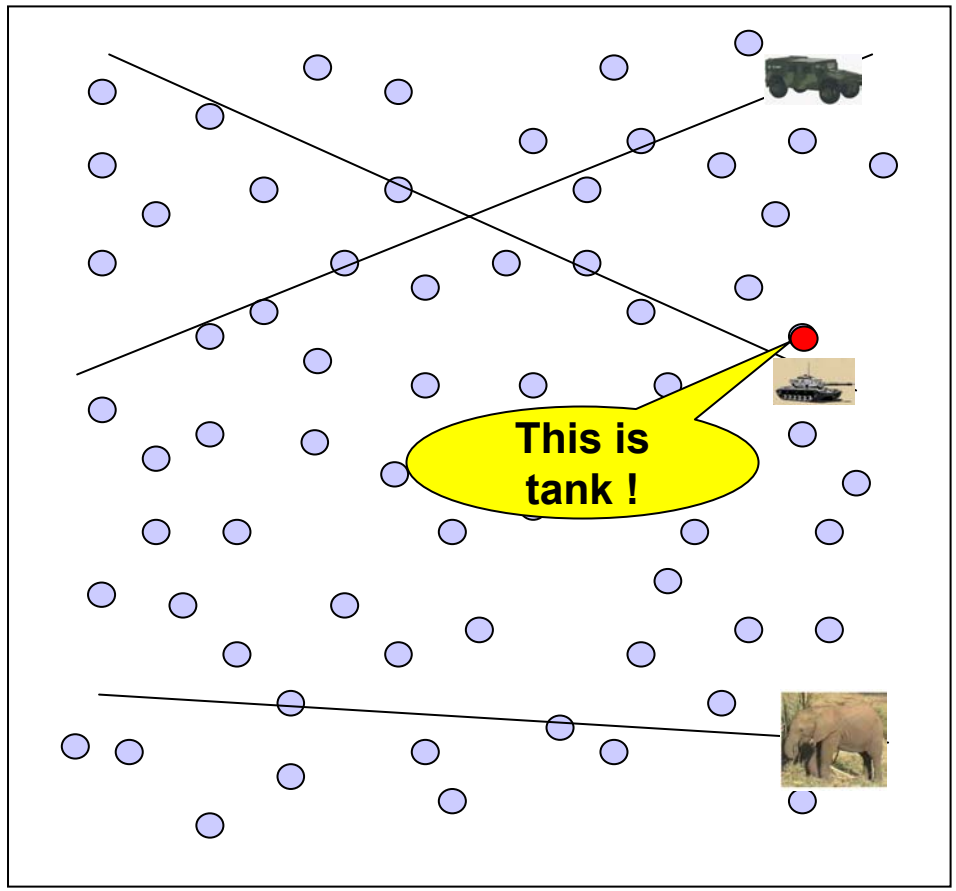

parc
Palo Alto Research Center

De-mixing of events

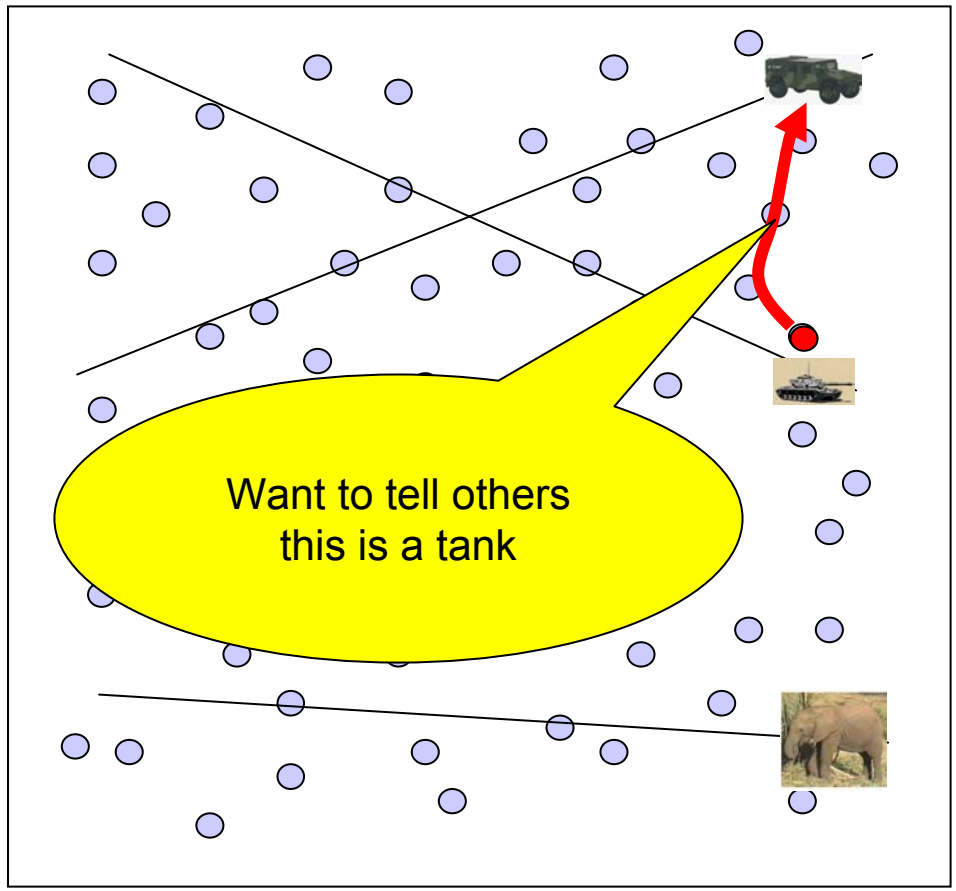

parc
Palo Alto Research Center

Another example of tracking multiple, interacting events

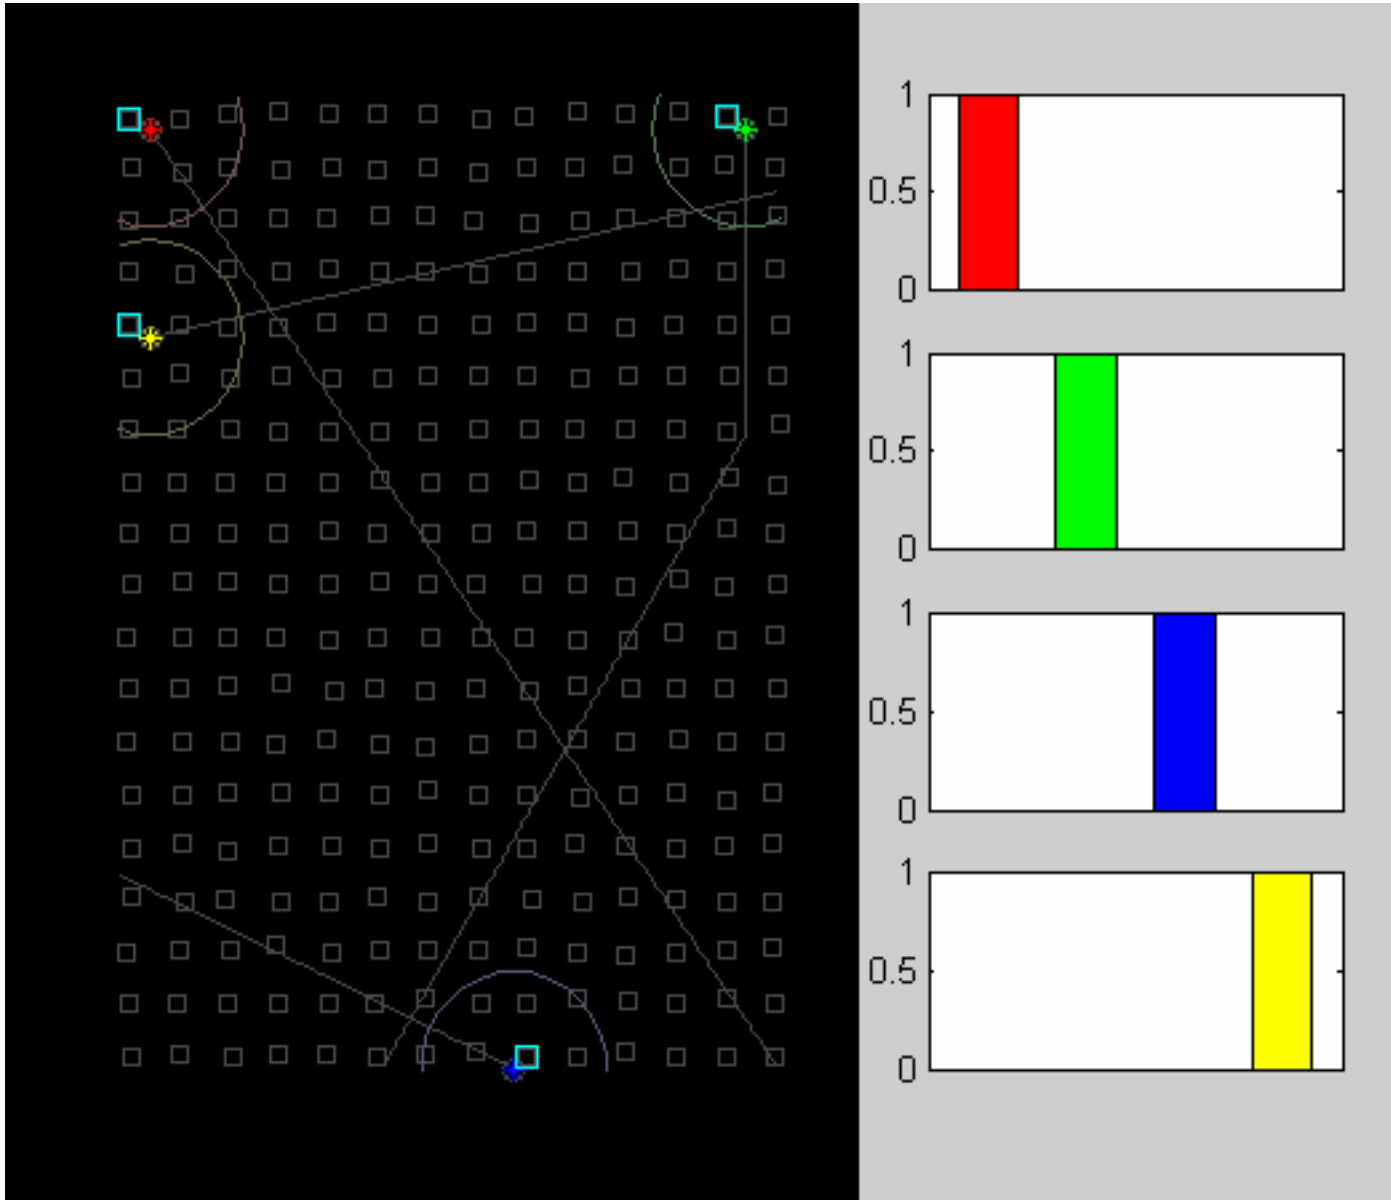

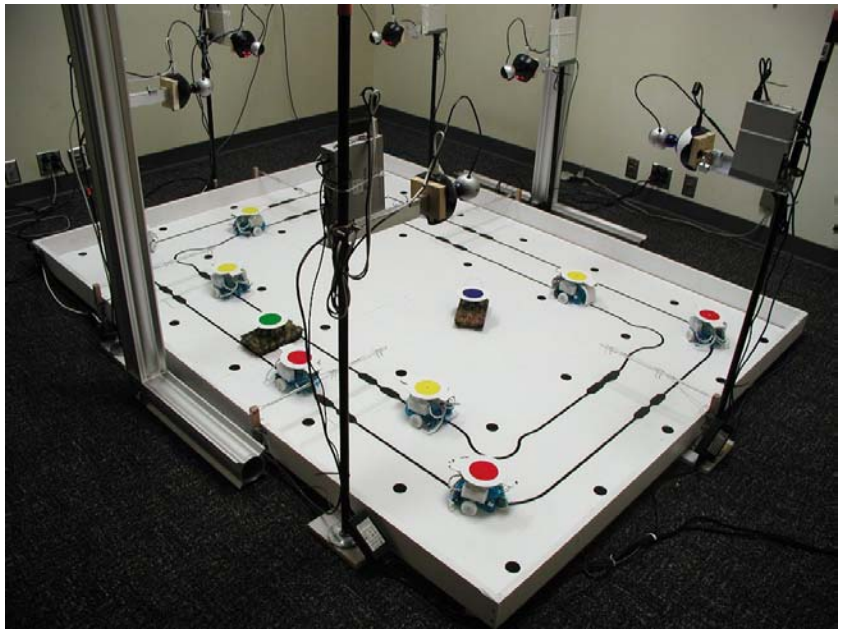

Video sensor network testbed

- 6 networked pan-and-tilt cameras
- FOV of cameras; each camera can see two FOV's on either side; overlapping FOVs
- Calibration points (1.5" diameter)
- Robot tracks (thickness 3/8")

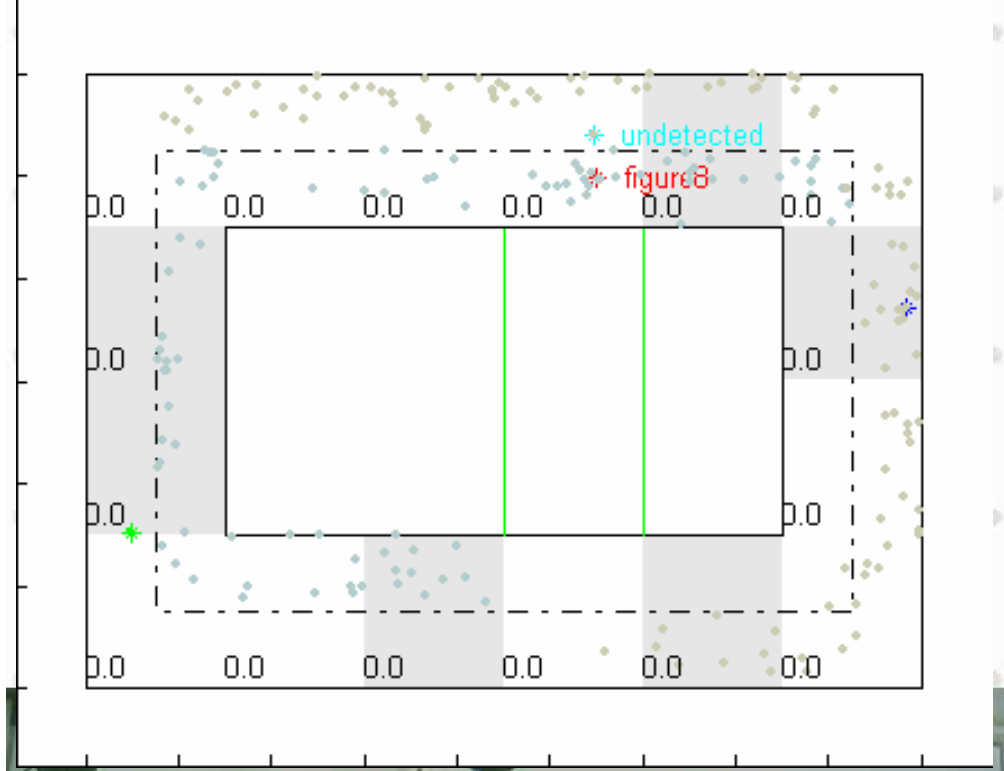

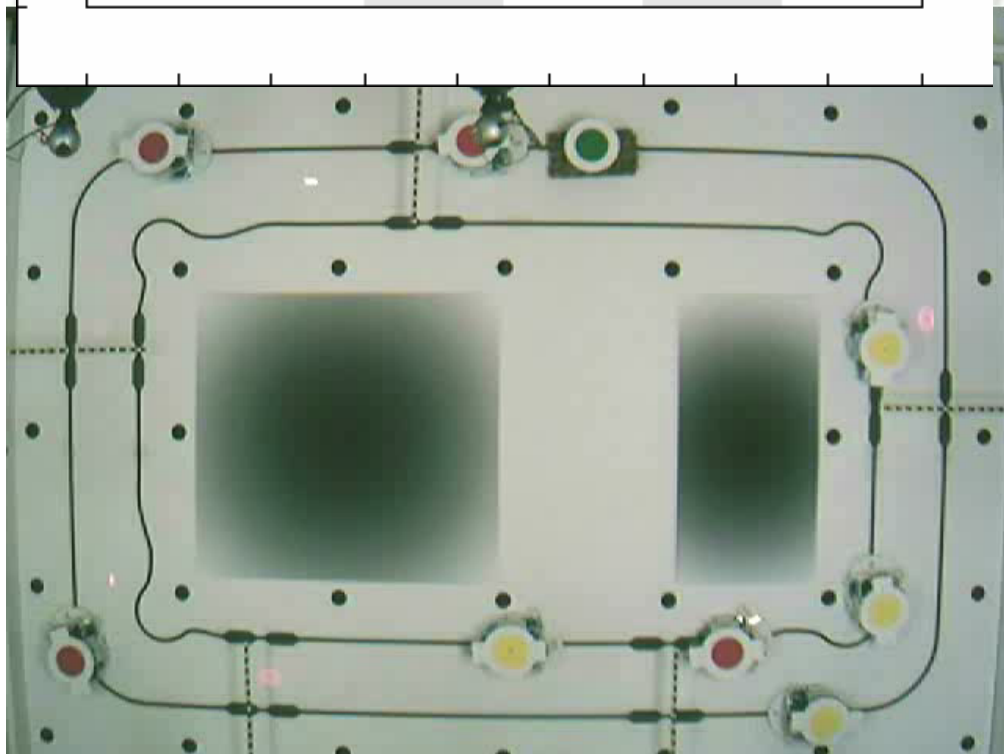

Examples of Group-Based Communication Patterns

- Acquaintance group
  - Roaming members keep persistent connectivity

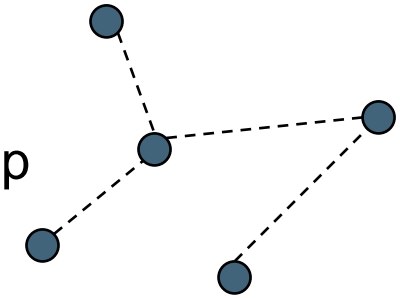

- Geographically constrained group
  - Defined by geographic extent

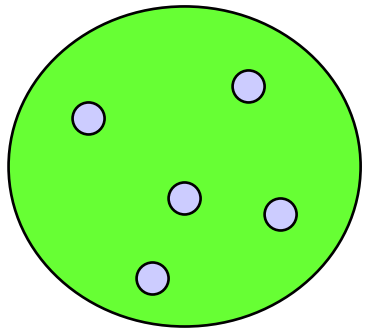

- Publish-subscribe group
  - Defined by producers and consumers of shared interests

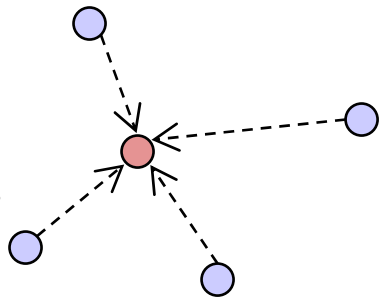

- N-hop neighbor group

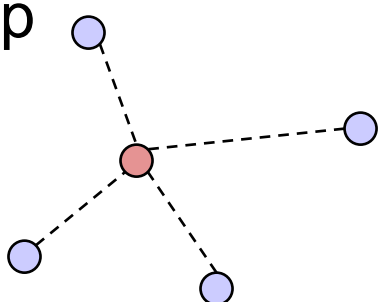

parc
Palo Alto Research Center

RoamHBA: maintaining connectivity among roaming agents

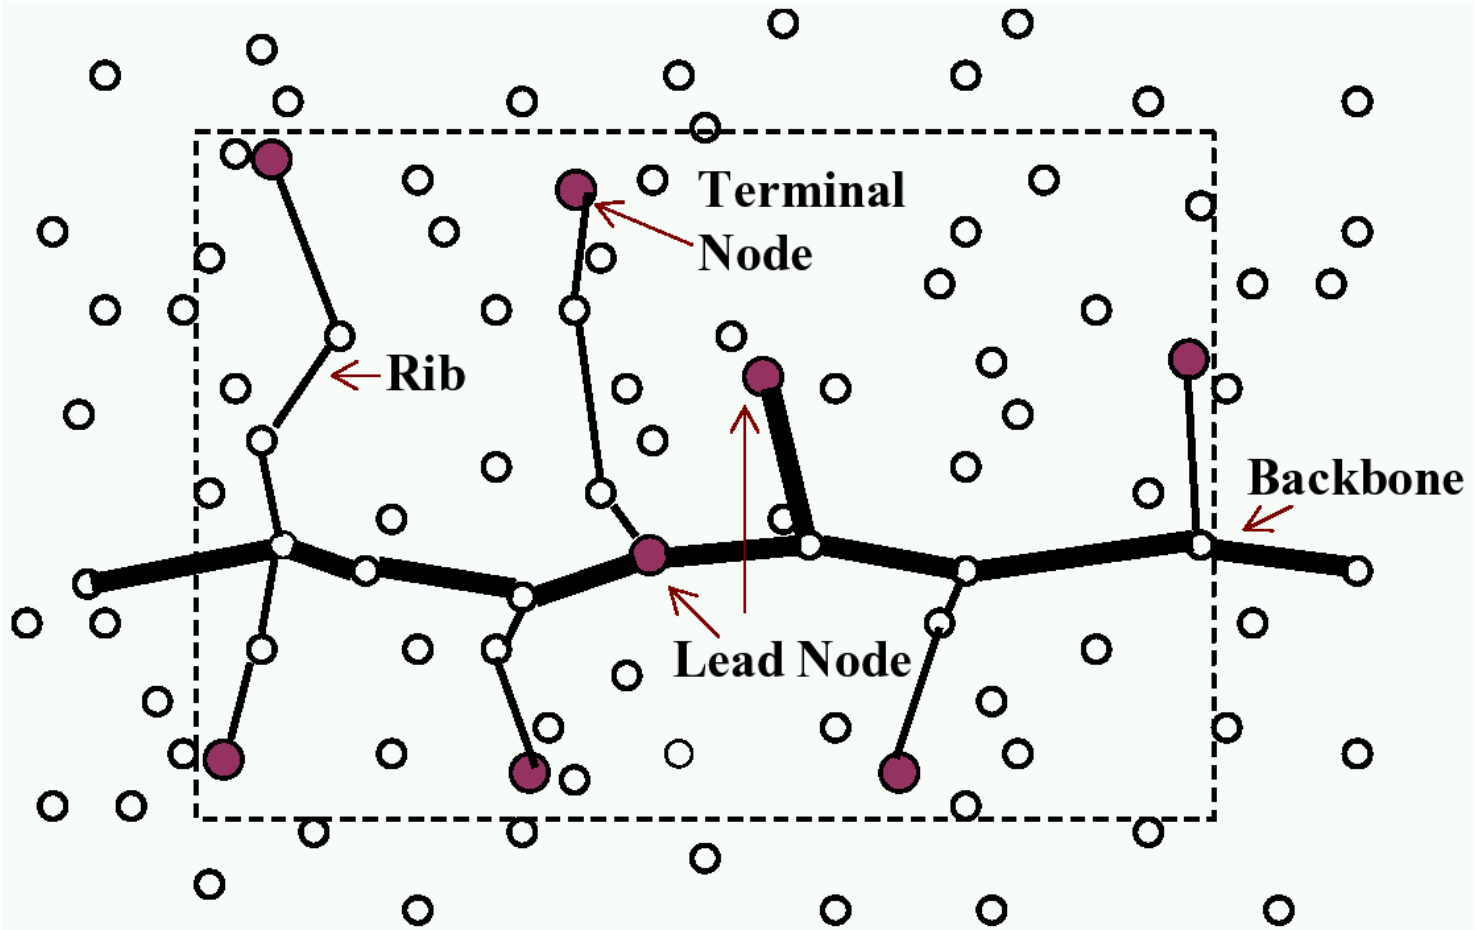

Any-time algorithms for sensor networks

- Incremental behaviors are important
  - Nodes may die, links may drop
  - Resources may be depleted
  - Stopping criteria may not be known a priori
- Quality of information should monotonically improve as more data become available
  - Without the need to re-compute from scratch every time
  - More graceful degradation
- Examples: localization, tracking

An Example: Incremental localization

- Least square formulation of localization

$$z_i = f(|x_t - x_i|) + \Delta z_i \quad \Rightarrow \quad a_i^T x_t = b_i$$

- Can be solved incrementally.
  - Each node only needs to pass a 2x2 matrix and a 2-vector, regardless of mote count
  - Incremental, anytime algorithm

- iLS node localization:
  - Propagate location information from small number of known anchors to non-anchor nodes.
  - Error control avoids overweighting of single measurement by propagating error estimates
- iLS Tracking
  - Identical algorithm

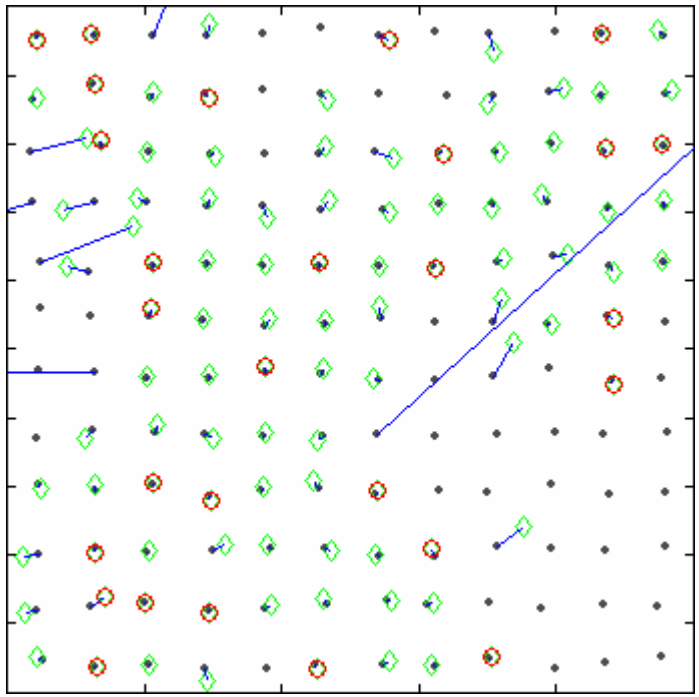

parc
Palo Alto Research Center

Collaborative processing in sensor networks

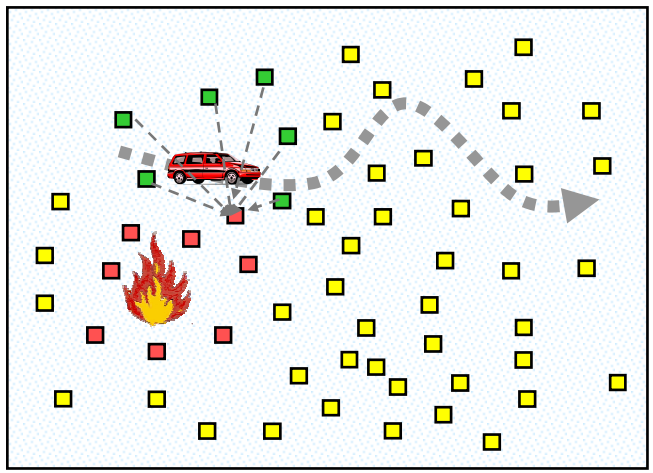

Information-centric design:

- What information is critical for the high-level tasks
- What is the cost of accessing the information
- Which nodes should participate in sensing, processing, or communication?
- How should the information be migrated?

The information processing needs largely determine the roles of and are supported by other layers of a sensor network

parc
Palo Alto Research Center