Challenges in Distributed Management

Rolf Stadler

Laboratory for Communication Networks
KTH Royal Institute of Technology

Stockholm, Sweden

20[th] NMRG Meeting, Utrecht, Oct 18-19, 2006

Objective and Implications

Goal

- scalable, robust management system with low complexity
  for large-scale, dynamic network environments

Approach

- decentralization of functionality, self-organization
- split the management plane into two
  - a thin, pervasive layer with end-to-end functionality
    part of the managed system
  - traditional FCAPS functions

Implications

- *research methodology*
  more formal methods, more probabilistic functions
- *overlap with other communities*:
  sensor networks/embedded systems, streaming databases, data mining
- later: equipment manufacturers, standardization

What and How

Functionality

– *efficient network-wide monitoring*
  polling aggregates, top k queries; continuous monitoring,
  trigger exceptions when system state leaves operational envelope
  self-tuning monitoring functions

– *independent local control actions* following global policies
  auto-configuration

Design space

– engineering a light *self-organizing layer*

– create and *maintain management* overlay with suitable topology

– *protocols* for processing management information:
  tree-based, flooding, gossip, random walks
  deterministic vs. probabilistic

The Network as a Management Computer
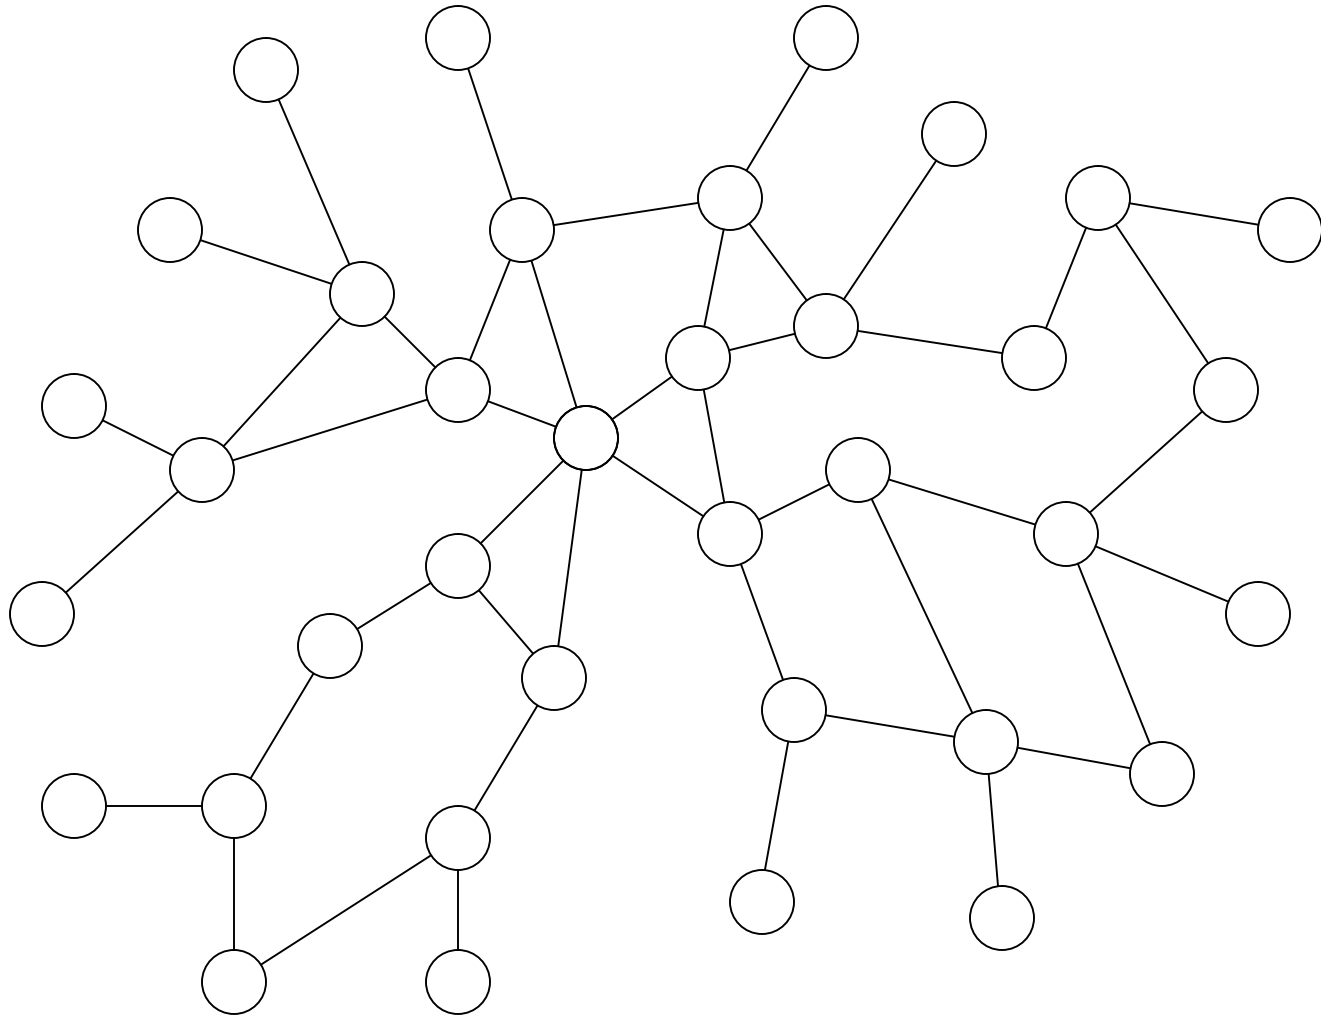
The hard Problems

Challenges
- understand and control *trade-offs* between
  completion time, accuracy, overhead, robustness.
- *interdomain management*
- *safety and security*

Other communities with similar research focus
- *sensor networks* "in network processing"
  our domain has different constraints
  - minimize energy needed for an operation
    vs. having a bound on load on nodes
  - building link-level overlays
    vs. mapping general overlays on physical infrastructure
  - aggregating a single variables
    vs. aggregating hundreds or thousands
- *streaming databases*
  - aggregate time series of a single variable
    vs. aggregating of 100s of local variables, aggregating over space