University of St Andrews

Programming for adaptive sensor networks

Back to the future

Simon Dobson
sd@cs.st-andrews.ac.uk
http://www.simondobson.org

IFIP WG2.11 Generative Programming workshop. St Andrews, UK. March 2010.

Overview

- Sensor networks are the new frontier for distributed systems
  - Enormous potential for fascinating research whilst also supporting real scientific experimentation
- Currently weak language support
  - Need to express adaptive sensing and autonomic control, network re-purposing and evolution
- My goal here
  - Explore the issues, and suggest some opportunities

The personal context

- I moved from UCD Dublin to St Andrews in October 2010
- Seems like a good time for a research semi-reset
  - Middleware, programming
  - Pervasive systems, uncertain reasoning, sensor fusion, situation recognition
  - Apply to environmental sensor networks
  - Novel languages (again)
  - Theory backed by experimentation

Context: environmental sensing

- New frontier of distributed systems
  - Small "motes" with limited processing, sensing and comms capabilities
  - Get power from *ad hoc* composition
- Challenges
  - Lots of partial failure
  - Don't get a Moore's Law effect
  - Adapt to what's being sensed
  - ...whilst maintaining scientific validity

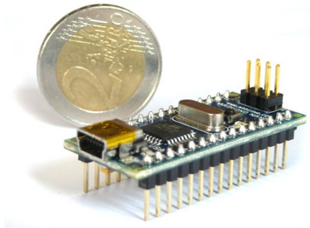

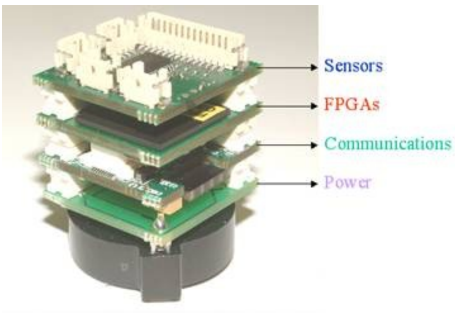

Scientific validity *vs* adaptation

- Environmental sensing has a mission
  - Measure pH/turbidity/elephants/whatever
  - Results must be *valid* in the sense of being a true reflection of the phenomena being observed
  - Must be maintained in the face of any adaptations we make to configuration or behaviour

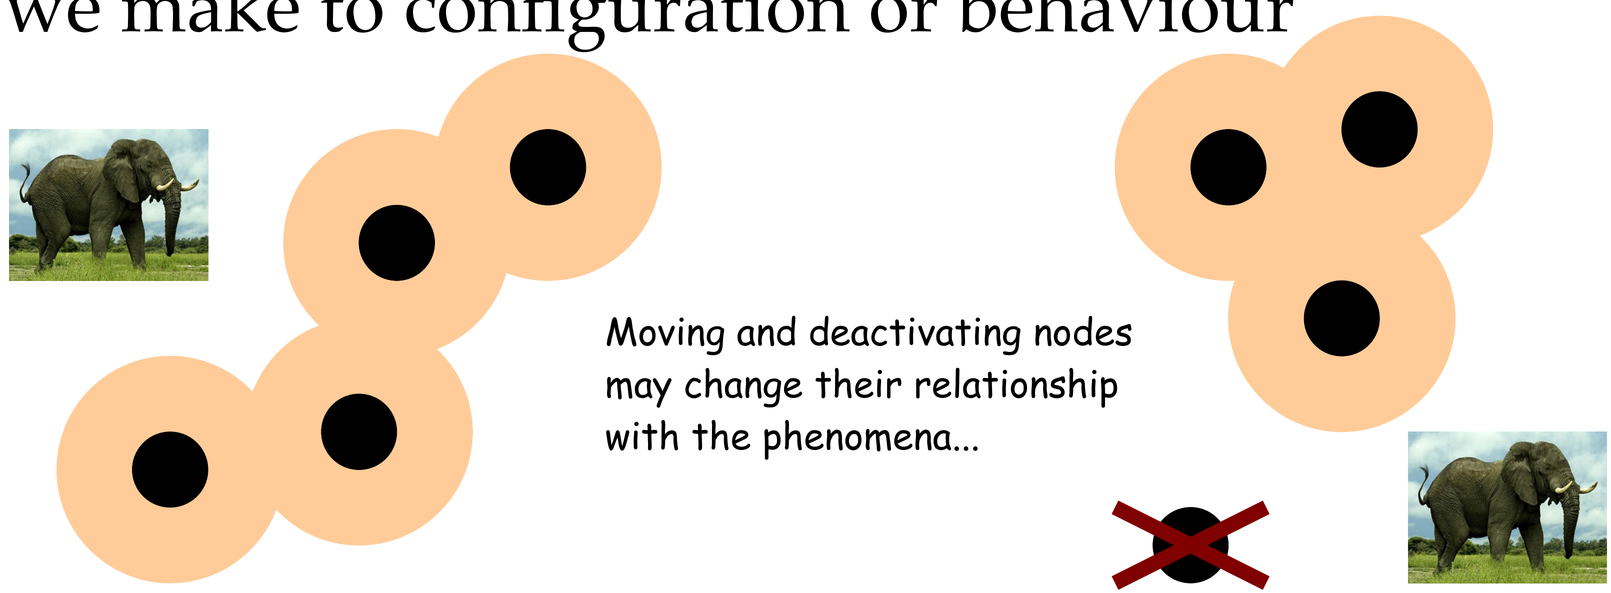

Missions and mission goals

Mission goals are almost always trade-offs

- Provide high-resolution sensing of the area
- ...but also have a long life to get good value
- ...and deal with partial failures in routing, sensing

- Often can't be made *a priori*
  - Frequent observation, mostly see nothing, run everybody's batteries down
  - Infrequent observation, better lifetime, miss the elephant
- *Adaptive* sensing is clearly desirable

Adaptive sensing

- Entangle the *scientific* functions with the *management* functions
  - *How* we sense depends on what we *have* sensed and what we *suppose we will* sense
  - Network becomes an active participant rather than a passive observer
  - Bound large-scale behaviour, allow adaptation within it

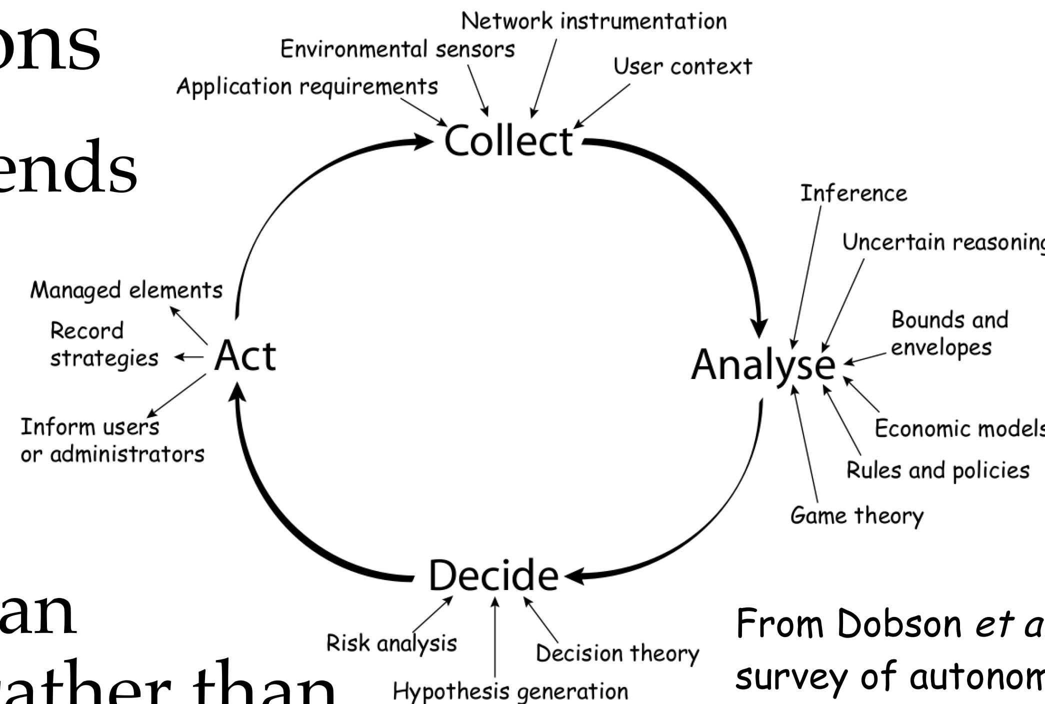

From Dobson *et alia*. A survey of autonomic communications. ACM TAAS 1(2). 2006.

But: the state of the art

- Limited languages and OSs
  - Some variant of C
  - Micro-kernel, limited database and comms function

Some variants: see Mainland, Morrisett and Welsh. Flask: Staged Functional Programming for Sensor Networks. Proc. ICFP. 2008.

Most common example is TinyOS and TinyDB for Berkeley/Crossbow motes

- Most innovation has occurred in comms
  - Robust self-routing protocols: AODV *et alia*
- Significantly less advanced in terms of programming and analysis
  - Need to program with large volumes of very uncertain data, in a way that's dependable

Concept mission: marine sensing

- Networks of mobile sensors
  - Move around to look at “interesting” places (or at random)

Title:river.eps
Creator:GIMP PostScript file plugin V 1.
CreationDate:Mon Feb 23 12:39:16 2009
LanguageLevel:2

Dobson, Coyle, O'Hare and Hinchey. From physical models to well-founded control, Proc. IEEE EASe. 2009.

How can we do this?

- Have to control the swarm of sensors as a whole
  - Patterns we're interested in lead to tactics for adaptation
  - Piecewise dynamics
  - ...but analysed at the swarm (network) level
- Has been demonstrated for simple cases, but needs to be generalised

"Move against the gradient of the vector field to find greater pollutant"

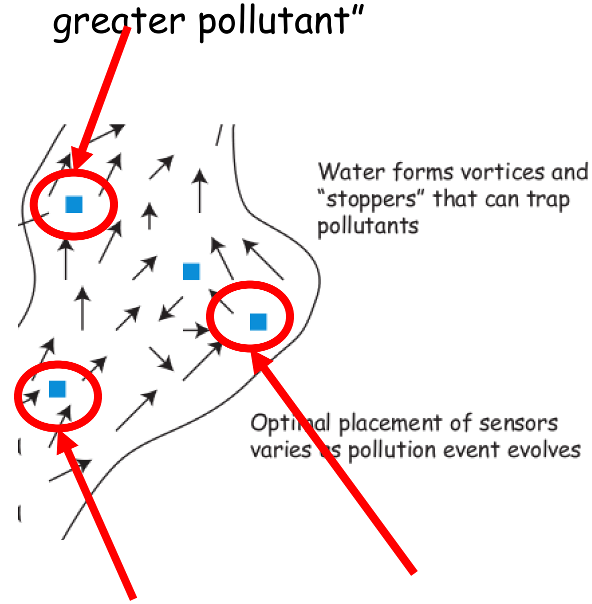

"Go somewhere no other node is"

"Move perpendicular to the curl to find the edges of vortices"

Programming with situations

- Semantically meaningful abstractions of what's being observed
  - Translate raw data using domain knowledge
  - Reasoning and machine learning

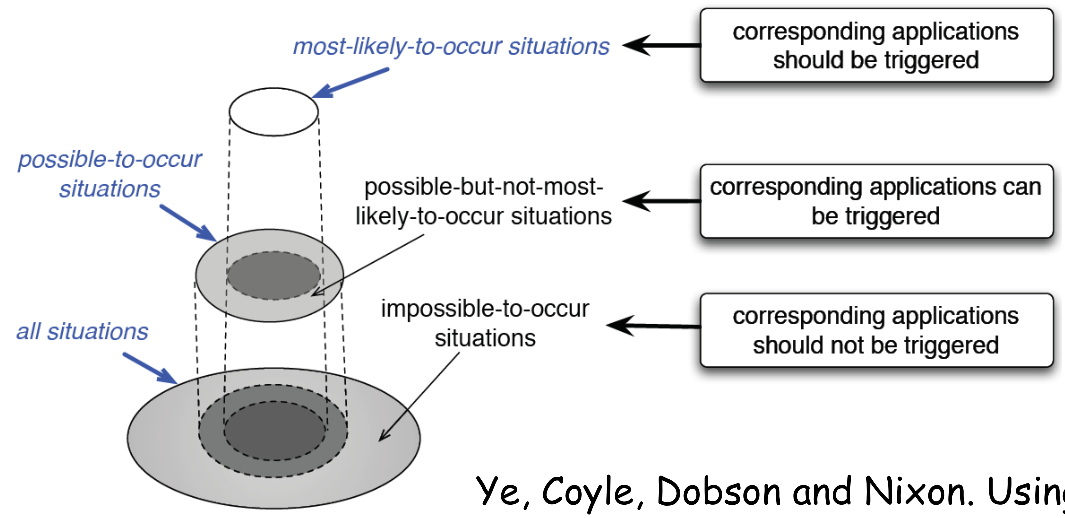

Ye, Coyle, Dobson and Nixon. Using situation lattices in sensor analysis. Proc. Percom'09.

- Identify situations where adaptations are needed, ensure they occur only at "safe" and "meaningful" points

A programming approach

- A programming approach with appropriate properties inherently
  - Structured according to mission and environment
    - Physically-inspired language constructs and patterns
  - Scalable in terms of nodes and data volumes
    - Generate the node code from reasoning
    - Move the reasoners into the network?
  - Deal with intrinsically uncertain/contradictory data
    - No `if` statement
    - Gradual, reversible decisions where possible

Semantics

- View the system globally as an *adaptive space*

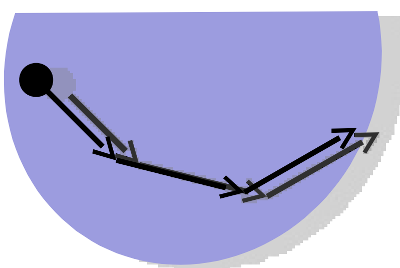

We can plot the ball's x-y position in the bowl and describe how it'll move, eventually coming to rest at the origin

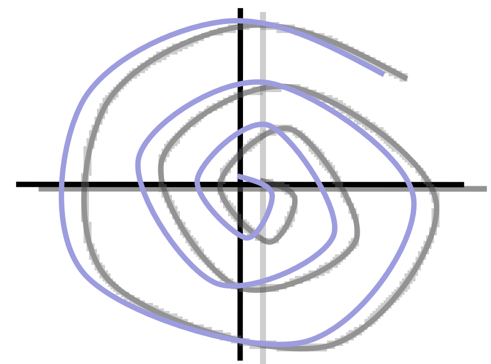

- Changing the environment changes the dynamics we see for the same actions we take
  - Still determined
  - Robust to small changes
  - Regions become situations

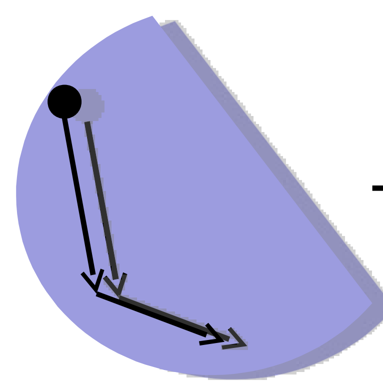

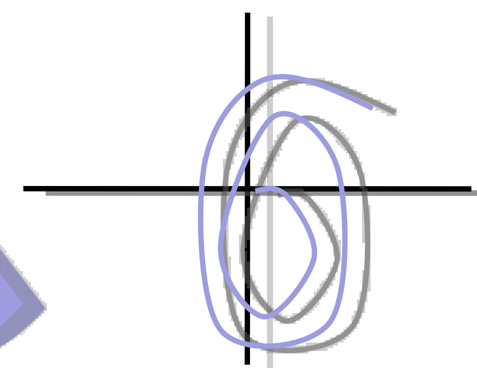

Generation

- Need to map this semantic model across the collection of nodes
  - Reasoning at the node and region level
  - Use the topology of the adaptive and real spaces
- Pluses and minuses
  - New programming model
  - Hard to co-ordinate in the face of limited comms
  - Robust and reflecting reality
  - Well-founded view of adaptation

Zhang, Nixon and Dobson. Multi-criteria adaptation mechanisms in homological sensor networks. Proc. IEEE ICCS. 2008.

In some ways the dual of classical dynamical systems: engineer a system with the given dynamic properties

Three things to take away

- Sensor networks need global analysis and behaviour generation
- Base behaviour on reasoning, and on a strong model of adaptation that's robust to noise
- A systems theory for adaptive computing

Dobson, Sterritt, Nixon and Hinchey. Fulfilling the vision of autonomic computing. IEEE Computer **43**(1). January 2010.