Controlling sensors with physics

Simon Dobson

UCD Systems Research Group
School of Computer Science
and Informatics
UCD Dublin IE

simon.dobson@ucd.ie
http://www.simondobson.org

Overview

Wireless environmental sensing is one of today's most exciting challenges

How do we perform unintrusive sensing while still trusting the results we get?

My goal today

- Motivate entwining physical and control models
- Overview of our *very* early-stage ideas
- Suggest some approaches we might use in the specific case of marine and river sensing

Background: autonomic systems

Self-management of behaviour and changes in that behaviour

- Sensorised response to environment
- Feedback loop
- Stability

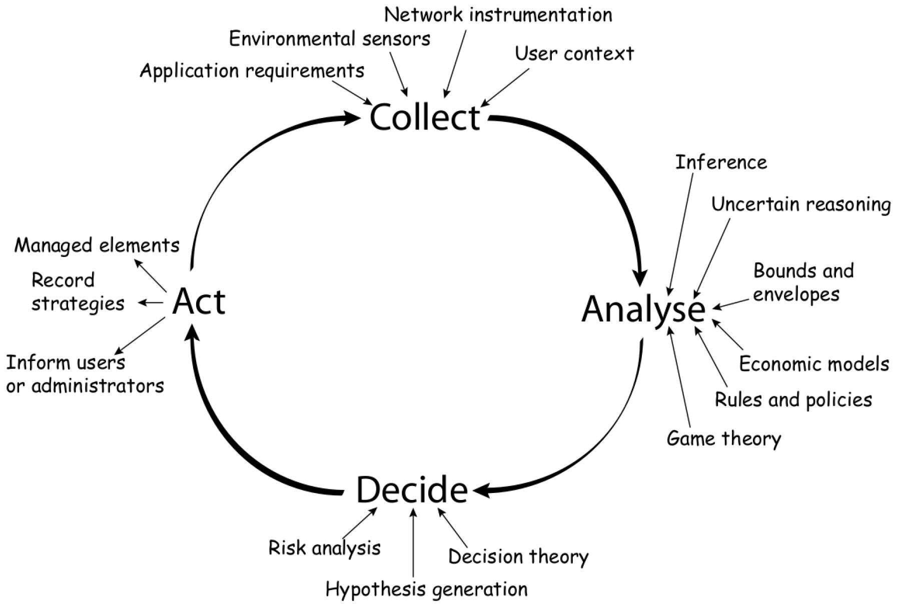

From Dobson *et alia*. A survey of autonomic communications. ACM Trans. Auto. Adapt. Sys **1**(2). 2006

Need to be able to capture the system's desired responses to stimuli and map them to changes in its behaviour in a clean way

Background: environmental sensing

One of the new frontiers of distributed systems

- Lightweight, low-power nodes
- Computation, communication, sensing and (possibly) actuation
- Networks built and maintained *ad hoc*

Challenges

- Deal with the limitations
- Maintain integrity of network in the face of failure
- Adapt behaviour of sensing to what's being sensed

Missions and mission goals – 1

Mission goals are almost always a trade-off

- Provide high-resolution sensing of the area of interest
- ...but also have a long life to get good value
- ...and deal with partial failures in routing, sensing
- ...and don't interfere with the environment being sensed
- ...and did we mention the long life?

Clearly conflicts we have to resolve

Missions and mission goals – 2

In a lot of missions we can't make these trade-offs *a priori*

- Fixed sensing and communication periods (*duty cycle*) makes for predictable battery usage
- Too long a sensing period risks missing phenomena
- ...too short burns power sensing the uninteresting
- Too long a communications period risks losing data through failures, either local or remote
- ...too short runs down everyone's batteries

Adapting seems to make sense

Adaptive sensing – 1

We therefore want to *entangle* the management of a node with its sensing functions

- Make duty cycle etc a function of what;s being sensed
- Increase frequencies when there's “something interesting going on”; reduce them otherwise

Makes things *much* more interesting

- Hard to model power lifetime
- An additional factor to consider in terms of system correctness

Adaptive sensing – 2

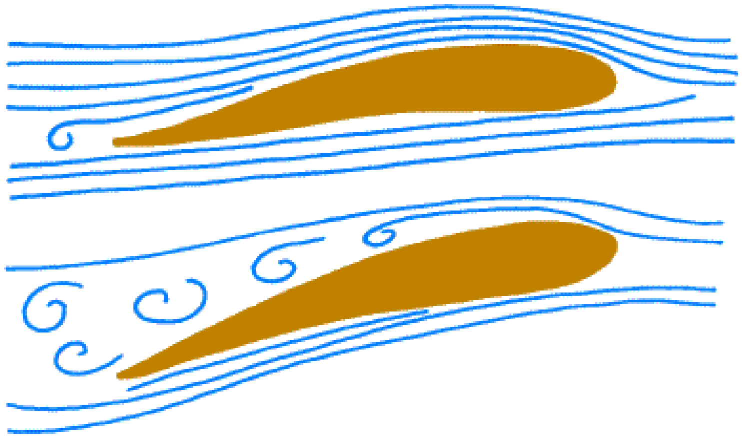

But there are advantages too

- Better sensing of events of interest
- Potentially use *less* power than the simpler system

Linking the sensing to the system

- Potentially increases confidence in results
- Not just random sampling

The uncertainty principle – 1

We don't want sensing to alter what we're sensing

- The Heisenberg uncertainty principle applied to sensing, perhaps?

This places limits on many things

- The size and intrusiveness of sensors – must be small enough not to interfere
- Their number – can't flood an area to the detriment of other uses

The uncertainty principle – 2

The uncertainty principle – 3

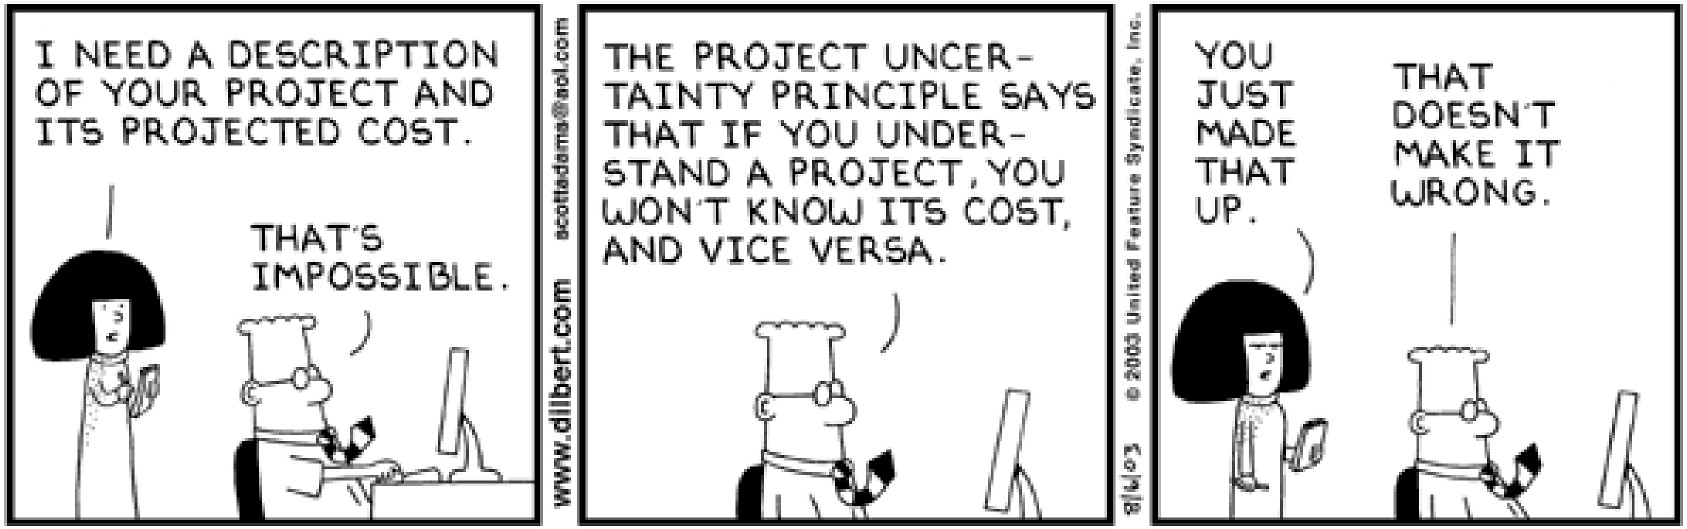

The Irish aspect – 1

Much of Ireland's income comes from tourism and fishing, so we have a major interest in water quality

Much of the pollution comes from farming run-off (nitrates) from inland

How does the pollutant reach the sea? How does it disperse once it's there? What effects does it have on the sea and the coastline

We – like every other country – *need* to know

The Irish aspect – 2

Focusing research on areas off the west coast of Ireland

- Galway bay
- Shannon estuary

How can we mount an effective sensing mission in these busy areas?

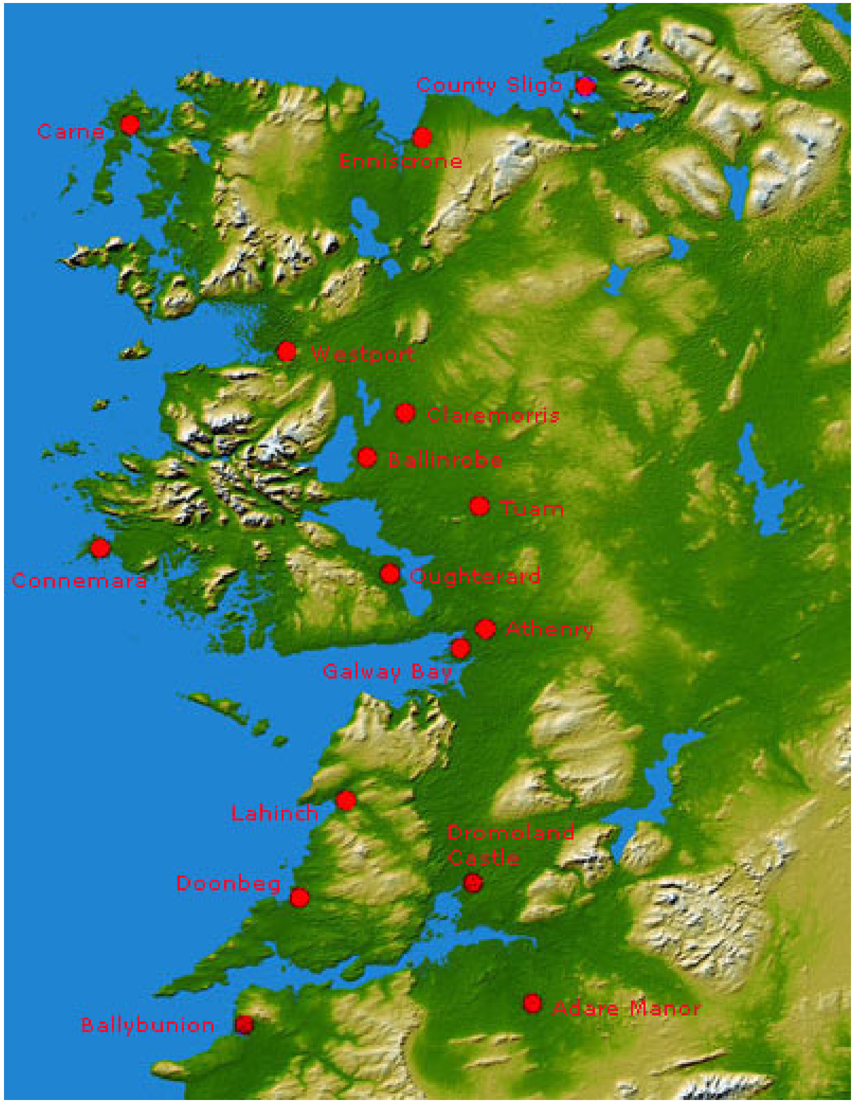

Options

Network of static sensors

- Position in “interesting” places (or at random)
- In reality, constrained to stay out of the shipping lanes, scenic areas, fisheries, ...

Mobile sensors

- Move around, purposefully or at random
- *Try* to stay out of everyone's way, or be small enough to be run down without a problem
- Much harder control problem

Dealing with power

A typical mobile sensor requires power, both for its sensing/computing/communication and for its motion

Remove the latter by using yachts

- Wind power to move and recharge
- Indefinite lifetime
- Major planning problem in terms of how to move from *a* to *b* in given wind conditions
- Big enough for “real” sensors

Mission architecture

We envision a network of 1m (or larger) model yachts with sensor packages

- Why models? They're small, cheap, an already rigged for computer control by remote control
- Maintain communications either as a mesh or through some or all having longer-range radio

Through the Clarity centre we have access to sensor experts as well as to hardware and users

CLARITY
clarity-centre.org

Marine Institute
Foras na Mara

sfi
science foundation ireland
fondúireacht eolaíochta éireann

UCD DUBLIN

Challenges

Too many to mention...

1. How can we sail a yacht under computer control so it goes where we want it to go?
2. How to we decide where we want to go?
3. How do we express this goal in a way we can analyse?
4. What is the best programming approach and/or language for highly sensorised adaptive systems?

For this talk we'll focus on the second and third

How to sail?

The “how” is horrificially complicated, but can be simplified

Anderson. The physics of sailing. Physics Today. Feb 2008.

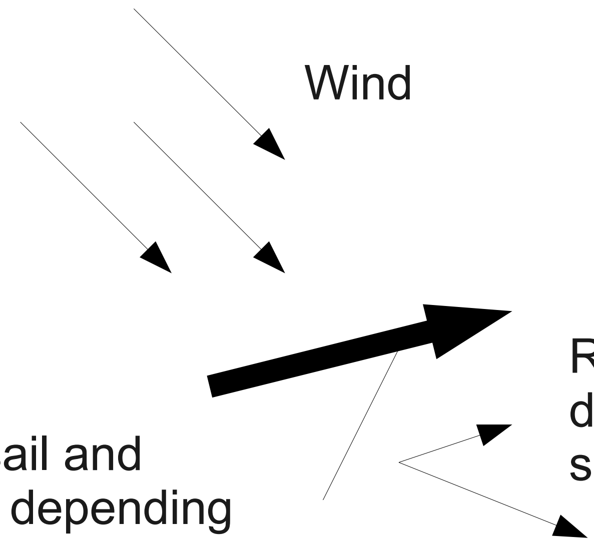

A small number of sailing manoeuvres depending on direction of wind relative to desired direction

- A fairly classical AI planning problem

Where to sail?

Where would we want to sail to?

- Random direction – might find something interesting
- Static search pattern – can be tailored
- Dynamic pattern – need to know how to plan the pattern

Analogy: if you randomly sample an airflow over a wing, you'll get mostly laminar flow

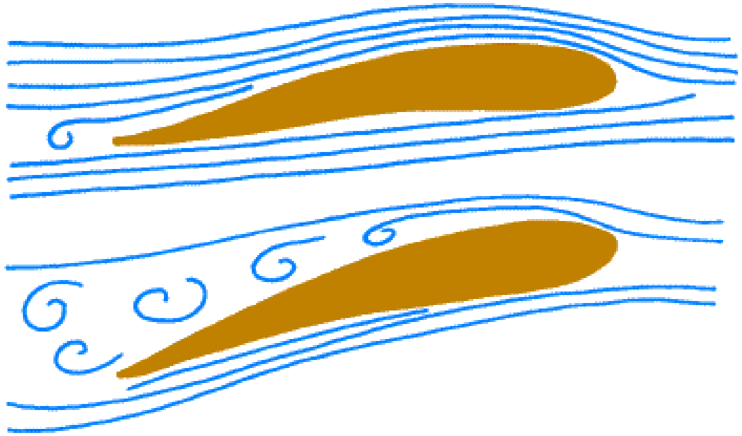

Knowing the physics

In order properly to plan a search pattern, we need to understand the physics of what we're searching for

- What constitutes an “interesting” place?
- How do these places evolve?

Although the detailed understanding of water flows is extremely complex, a naïve understanding will (perhaps) suffice for our purposes

A naïve understanding

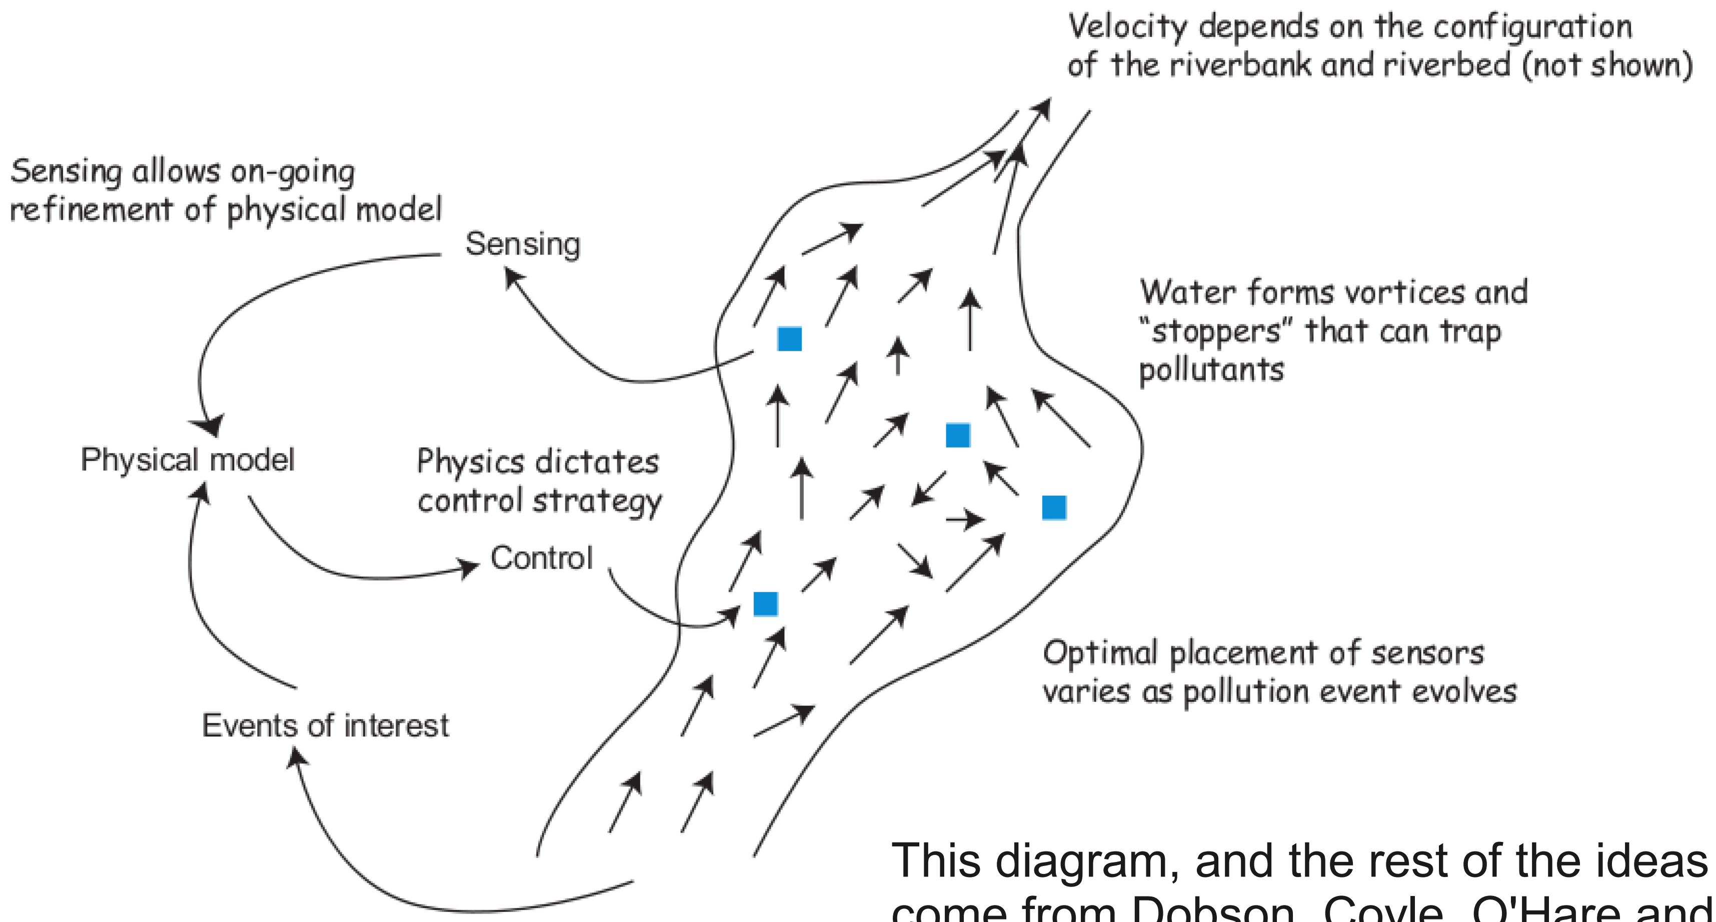

This diagram, and the rest of the ideas in this talk, come from Dobson, Coyle, O'Hare and Hinchey. From physical models to well-founded control, Proceedings of IEEE EASe. 2009. To appear.

Controlling the swarm

We formulate this problem as one of maximising a value function for the swarm of yachts

Parameters

- The wind (vector)
- The flow field of the river (vector)
- The pollutant level at each point (scalar)
- The locations of the yachts (from GPS)

Value function

- Is each yacht sensing something interesting? Is the area being covered?

Defining the value function

We don't let have a really good definition

- Areas of high pollution
- Areas downstream of areas of high pollution
- Areas not being observed, to avoid missing other events

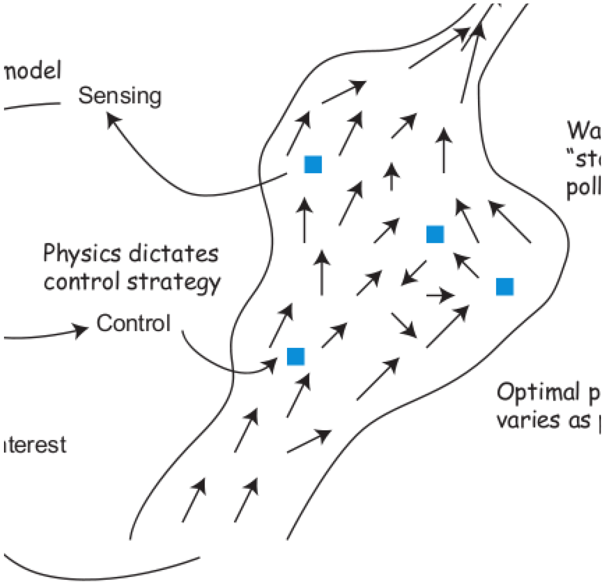

All seem to need a combination of local and global information, and a considerable amount of data exchange

Approach, once we have one

For any given scenario we can assign a value to any particular sensor constellation

- Non-unique values

For any particular sensor constellation we can define a movement that moves towards a better (or no worse) constellation

- Mostly a local operation, but requiring a global view of the scenario
- Must be balanced against what's possible, in terms of sailing against the wind etc

Software engineering properties

It's important to realise that, although the model is defined globally, its implementation is neutral in terms of local and global decision-making

- May get better or worse results, and better or worse consumption of resources
- Provides a *semantics* against which to judge any solution, and against which to prove properties
- Decouple *specification* from *solution* to get better analysis – not common in the autonomics literature

Methodology

1. Obtain a physical model
2. Define capabilities of sensor nodes
3. For each model configuration define a “good” or “best” positioning constellation
4. Apply tactics to move current constellation towards a better (or no worse) one
5. Evaluate tactics by (for example) time to converge to best constellation – even though this will change in reality, and never be reached

Current state

- We can sail, in a straight line, downwind-ish, on our lake
- We can define simple models of fluid flow
- We have tried defining value functions, none of which is really great
- We are starting to evaluate the combination of model and control integration
- We have a basic set of tactics for sailing

Three things to take away

Wireless sensing using mobile nodes can address issues in the depth of sensor coverage it's possible to achieve

Using physical models may lead to data in which we can have better confidence

Tying physics to control can give an analytic framework within which to explore the solution space