



IMAGE

EVERYTHING, EVERYWHERE

Tiny computers that constantly monitor ecosystems, buildings and even human bodies could turn science on its head. **Declan Butler** investigates.



“What are you doing in the lab? Why aren't you out working in the field?” These are not the sorts of question you usually put to your computer. But they should be, according to the proponents

of a new type of information technology sometimes known as ‘smart dust’.

In their current, mostly desktop, incarnation, computers used for science usually come into their own quite late in the process of inquiry. Questions are asked, the data that might answer them identified, that data gathered — and only then does the computer start to play a role. In the future, this set up could be reversed. Computers could go from being back-office number-crunchers to field operatives. Twenty-four hours a day, year-in, year-out, they could measure every conceivable variable of an ecosystem or a human body, at

whatever scale might be appropriate, from the nanometric to the continental.

These new computers would take the form of networks of sensors with data-processing and transmission facilities built in. Millions or billions of tiny computers — called ‘nodes’, ‘nodes’ or ‘pods’ — would be embedded into the fabric of the real world. They would act in concert, sharing the data that each of them gathers so as to process them into meaningful digital representations of the world. Researchers could tap into these ‘sensor webs’ to ask new questions or test hypotheses. Even when the scientists were busy elsewhere, the webs would go on analysing events autonomously, modifying

their behaviour to suit their changing experience of the world.

If this scenario sounds far fetched, imagine the owner of a mainframe in the 1970s asking why it wasn't sitting on millions of desks and laps worldwide. An absurd question — to which the answer was “it's just a matter of time”. The world's stock of computing power, and the number of devices over which it is distributed, has increased exponentially since then, as has the capacity of networking technology. These trends show no sign of slowing down, and that makes pervasive sensor nets not so much possible as inevitable. One does not need to be a visionary to see that soon, tiny devices with the power of today's desktops will be cheap enough to put everywhere.

Gaetano Borriello, a computer scientist at the University of Washington in Seattle, argues that such widely distributed computing power will trigger a paradigm shift as great as that

“We will be getting real-time data from the physical world for the first time on a large scale.”

— Gaetano Borriello

brought about by the development of experimental science itself. "We will be getting real-time data from the physical world for the first time on a large scale."

Instead of painstakingly collecting their own data, researchers will be able to mine up-to-the-minute databases on every aspect of the environment — the understanding of diseases, and the efficacy of treatments will be dissected by ceaselessly monitoring huge clinical populations. "It will be a very different way of thinking, sifting through the data to find patterns," says Borriello, who works on integrating medical sensors — such as continuous monitors of heart rate and blood oxygen — with their surroundings. "There will be a much more rapid cycle of hypothesis generation and testing than we have now."

Mallikarjun Shankar, who works on sensor webs for military and homeland security at Oak Ridge National Laboratory in Tennessee, agrees. "If one looks at the landscape of computing, this is where it will link with the physical world — where computing science hits the tangible, the palpable world. It is the next frontier."

From virtual to actual

Things don't get much more palpable than the experience of having an offshoot of Europe's biggest ice sheet grinding you into the granite below. That is the lot in life of the sensor web that Kirk Martinez, of the University of Southampton, UK, has been running for the past few years. He is helping glaciologists to study the dynamics of the Briksdalsbreen glacier in northwest Norway, part of the Jostedalbreen ice field, in the hope of better understanding the impact of climate change and weather patterns on the ice field¹.

Martinez's team uses a hot-water 'drill' to place pods — a dozen at any one time — at different locations deep under the ice. Each pod is equipped with sensors that measure variables such as temperature, pressure, and movement; the data collected are used to work out the rate of flow of the glacier and to model subglacial dynamics. The sensor web can be programmed in such a way that the individual pods cooperate. "You can get the pods talking to each other, and deciding that nothing much has happened recently as most of our readings have been the same, so lets the rest of us go to sleep and save batteries, with one waking us up if something starts happening," says Martinez.



Watchful eyes: tiny computers called motes may one day be in everything.

"This is where where computing science hits the tangible, the palpable world. It is the next frontier."

— Mallikarjun Shankar

Martinez himself is a computer scientist, not a glaciologist. He was drawn to the task of remotely monitoring a hostile environment around the clock because he wanted the challenge of trying to bring together the various different technologies such sensor webs need. "This is very, very technologically tricky stuff," he says.

For non-computer scientists, it is even trickier. Researchers can already buy pods the size of cigarette packets or credit cards off the shelf from a slew of new companies such as CrossBow and Dust Networks (both based in the Bay Area of California). But that's only the start. At present, creating a sensor web for a specific scientific application requires extensive customization, particularly in the programming.

Katalin Szlavetz, an ecologist at Johns Hopkins University, in Baltimore, Maryland, works on soil biodiversity and nutrient cycling. She was driven to sensor webs by frustration caused by trying to solve complex problems with limited data. Soil is the most complex layer in land ecosystems, but it is poorly understood, because sampling is limited by the fact that technicians must collect soil samples by hand, and analyse them back in the laboratory. "Wireless sensor networks could revolutionize soil ecology by providing measurements at temporal and spatial granularities previously impossible," she says.


Data stream

Last September, Szlavetz deployed ten motes along the edge of a little stream just outside the university campus. Each mote measures soil moisture and temperature every minute, and the network transmits its data periodically to a computer in her office.

It sounds simple, but just to get the pods up-and-running she had to create a multidisciplinary team, involving computer scientists, physicists and a small army of student programmers. Szlavetz says that she has "no doubt" that pod networks will take off widely, but that it won't happen until they are easier to deploy. And cheaper: "Each unit costs around US\$300, but if you include all the hours of student time, each



1976 At Los Alamos, Seymour Cray installs the first Cray supercomputer, which can process large amounts of data at fast speeds.



DALE BOYER / SCIENCE PHOTO LIBRARY

1980s >>

1983 Danny Hillis develops the Connection Machine, the first supercomputer to feature parallel processing. It is used for artificial intelligence and fluid-flow simulations.

1985 After receiving reports of a lack of high-end computing resources for academics, the US National Science Foundation establishes five national supercomputing centres.

1989 Tim Berners-Lee of the particle-physics laboratory CERN in Geneva develops the World Wide Web — to help physicists around the globe to collaborate on research.



MICHAEL QUAN/ZUMA PRESS

1990s >>

1990 The widely used bioinformatics program Basic Local Alignment Search Tool (BLAST) is developed, enabling quick database searches for specific sequences of amino acids or base pairs.

1996 George Woltman combines disparate databases and launches the Great Internet Mersenne Prime Search. It has found nine of the largest known Mersenne prime numbers (of the form $2^n - 1$), including one that is 9,152,052 digits long.





KNUDSEN

Science of the future: researchers can keep a constant eye on the flow of a Norwegian glacier by tracking miniature sensors buried beneath the ice.

works out closer to \$1,000." Despite this, Szlavetz says, the Microsoft-funded pilot was a success, revealing previously unobserved variations in soil microclimate, and showing how rain affects wetting and drying cycles². (Szlavetz's colleague Alexander Szalay is an author of the commentary on page 413).

Handful of dust

Difficulties like Szlavetz's are all too common, says Kris Pister, founder and chief technology officer of Dust Networks. It was Pister who, at the University of California, Berkeley, coined the term smart dust to describe his vision of sensors smaller than the eye could see joined into networks larger than the mind could comprehend. Pister built prototype sensor webs with funding from the Defense Advanced Projects Research Agency, which is interested in the technology's military applications. But he says it was the desire to create more usable systems that led him to get the commercial backing to create Dust Networks. "It was very frustrating that while we could do spectacular demos, we couldn't give scientists something off-the-shelf, to put in a tree or a river. What people want is the ability to just put sensors out in the environment and get data back."

He likens the current stage of sensor-web development to the early days of computing. "There are a group of experts at the cutting edge of sensor webs, who have the time and expertise to go in and learn how to use the tools

and do all the neat stuff," he says. "But for people who are not experts it has been difficult to get in and use it." He predicts, however, that just as the first web browser, Mosaic, and its successor, Netscape, sparked mass take-up of the World Wide Web, so future, more user-friendly sensor-web tools will generate interest.

Although Pister is interested in scientific applications, his key target is the lucrative industrial market for control systems. He has been contracted by the US Department of Energy to help build 'intelligent' self-monitoring lighting systems for factories, offices and homes; the 600 billion kilowatt-hours of lighting used for this purpose account for 30% of total building electricity use. "The next step is about really getting some standards and commercial adoption," he says. "That will drive cost down and performance up, and then scientific uses will take off."

Even with today's sensor-web technology, applications for research are proliferating. The Jet Propulsion Laboratory (JPL) in Pasadena, California, which is responsible for most of NASA's planetary science, has been running

nine large sensor webs to study among other things, flooding, pollution and microclimates, in settings ranging from botanical gardens to the Seville National Wildlife Refuge in central New Mexico. This month, Kevin Dellin, the head of the JPL sensor-web programme, spun it off into a company, Sensorware Systems.

They mote be giants

But sensor webs currently have major limitations for people doing science in the field, says Deborah Estrin, director of the Center for Embedded Networked Sensing in Los Angeles, California, which operates a suite of land- and sea-based monitoring projects in collaboration with university groups. Estrin says that sensor webs alone are often not sufficient for all monitoring needs, and that the cost of sensors prohibits researchers from obtaining the pod densities often needed for detailed field experiments.

Estrin sees the sensor-web revolution as an important thread in a grander tapestry of global monitoring, which involves billions of dollars being poured into projects to monitor the continents and oceans. The US National Science Foundation's Ocean Observatories Initiative (OOI), for example, plans to spend \$300 million over the next six years on gigabit 'backbones' — fibre-optic carriers of data — across the floor of the Pacific Ocean. On land, the planned National Ecological Observatory Network would enable research on terrestrial

"We could do spectacular demonstrations but we couldn't give scientists something off-the-shelf, to put in a tree or a river."

— Kris Pister

ecosystems at regional to continental scales in real time. And underneath the surface, the EarthScope project would explore the four-dimensional structure of the North American continent from crust to core. Integrating local sensor webs and all these other networks is one of the biggest challenges facing the development of observational science, Estrin says.

Such mega-observatories may seem very different from Szlavetz's handful of little sensors alongside a stream in Baltimore. But the principles behind them are strikingly similar: to suck the best real-time data out of the environment as possible. "Instead of handling individual data files you will be handling continual streams of data," says Robert Detrick, chair of the OOI's executive steering committee. "You will be combining inputs from different sensors interactively to construct virtual observatories: sensors will be in everything."

The OOI's aim is to get around the fact that oceanographers tend to see only what their research vessels happen to be traversing at any given time. Checking in on the ocean fibre-optic backbone will be swarms of tiny autonomous submarines. These will carry sensors, go off on sampling missions and return to the backbone to recharge their batteries and upload data. "They can all communicate with one another acoustically," Detrick enthuses. "One can say, 'Hey, I've found an anomaly over here, so come on over'". Static sensors will monitor tectonic plates continuously along their entire length. Episodic events such as earthquakes, volcanic eruptions, and instabilities in ocean currents will



Ecologist Katalin Szlavetz's mini computers tell her about minute-by-minute changes in soil moisture.

be captured in real time, something that is impossible to do using ships.

The existence of such large networks points to some major challenges down the line, says Estrin. Sensor webs will frequently be just single layers in a stack of data-collecting systems. These will extract information at different temporal and spatial scales, from satellite remote-sensing data down to *in situ* measurements.

Managing these stacks will require massive amounts of machine-to-machine communication, so a major challenge is to develop new standards and operating systems that will allow the various networks to understand each other. Sensors and networks of sensors will need to be able to communicate what their data are about, how they captured and calibrated them, who is allowed to see them, and how they should be presented differently to users with different needs. The lack of standards is not an insoluble problem for sensor webs, says Shankar "but it is slowing the field down by several years".

Catching the moment

Despite the difficulties, the use of sensor webs continues to grow. For all the trouble her first ten pods caused her, Szlavetz is upgrading to a network of 200. And experts see no fundamental obstacles to their eventual ubiquity. "We are well on the road to getting there, and I would argue that on many points we are already there," says Borriello. By 2020, says Estrin, researchers using visualization tools like Google Earth will be able to zoom in, not just on an old satellite image, but on "multiple *in situ* observations of the Earth in real time".

Data networks will have gone from being the repositories of science to its starting point. When researchers look back on the days when computers were thought of only as desktops and portables, our world may look as strange to them as their envisaged one does to us. Although we might imagine a science based so much on computing as being distanced from life's nitty gritty, future researchers may look back on today's world as the one that is more abstracted. To them the science practised now may, ironically, look like a sort of virtual reality, constrained by the artificialities of data selection and lab analysis: a science not yet ready to capture the essence of the real world. ■

Declan Butler is a senior reporter for Nature based in Paris.

- Martinez, K., Riddoch, A., Hart, J. & Ong, R. in *Intelligent Spaces* (eds Stevenson, A. & Wright, S.) Ch. 9, 125-138 (Springer, 2006).
- Musualoiu-E., R. et al. *Life Under your Feet: A Wireless Soil Ecology Sensor Network* (submitted) The Third IEEE Workshop on Embedded Networked Sensors (2006).



1996 Craig Venter develops the shotgun technique, which uses computers to piece together large fragments of DNA code and hastens the sequencing of the entire human genome.

1998 The first working quantum computers based on nuclear magnetic resonance are developed.


21st CENTURY >>

2001 The National Virtual Observatory project gets under way in the United States, developing methods for mining huge astronomical data sets.



2001 The US National Institutes of Health launches the Biomedical Informatics Research Network (BIRN), a grid of supercomputers designed to let multiple institutions share data.

2002 The Earth Simulator supercomputer comes online in Japan, performing more than 35 trillion calculations each second in its quest to model planetary processes.



2005 The IBM Blue Gene family of computers is expanded to include Blue Brain, an effort to model neural behaviour in the neocortex — the most complex part of the brain.

2007 CERN's Large Hadron Collider in Switzerland, the world's largest particle accelerator, is slated to come online. The flood of data it delivers will demand more processing power than ever before.

Jacqueline Ruttimann