GEOMETRICALLY correct the digital **IMAGERY,**

then create a seamless **MOSAIC,**

and semi-automatically **EXTRACT FEATURES,**

completing the perfect **WORKFLOW.**

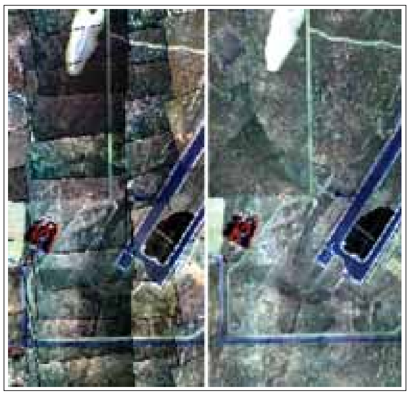

Successfully retrieve valuable information contained in your imagery. Leica Geosystems Geospatial Imaging provides proven technology solutions to maximize your workflow.

Scan your film to produce digital imagery with our high performance photogrammetric film scanner, the DSW700. The resulting output is ready to be orthorectified with Leica Photogrammetry Suite, our workflow-driven, production-oriented photogrammetric software tool. Or, use IMAGINE AutoSync™ and skip the photogrammetric processing step, employing your existing image library to semi-automatically georeference the new images. Easily create your seamless mosaic in Leica MosaicPro, with advanced seam-editing capabilities and increased overall efficiency. And finally, extract features effortlessly using semi-automatic feature extraction in IMAGINE Easytrace™.

From preparation of source material to final product generation, maximize your workflow with solutions from Leica Geosystems Geospatial Imaging. Call us at +1 877 463 7327 or +1 770 776 3400 or send an e-mail to info@gi.leica-geosystems.com

ERDAS IMAGINE®

DSW700

IMAGINE Enterprise Loader™

IMAGINE Enterprise Editor™

IMAGINE AutoSync™

Leica Virtual Explorer

Leica Photogrammetry Suite

Leica MosaicPro

Stereo Analyst® for ArcGIS

Image Analysis™ for ArcGIS

IMAGINE Easytrace™

capture · reference · measure · analyze · present

POWERING GEOSPATIAL IMAGING™

- when it has to be **right**

Leica Geosystems

Great Barrier Reef, Australia
(Courtesy: NASA MODIS)

SENSOR NETWORKS

PAUL HAVINGA EXPLAINS HOW WIRELESS SENSOR NETWORKS ARE BRINGING THE GREAT BARRIER REEF ONLINE, PROVIDING MARINE ENVIRONMENTAL DATA.

The Great Barrier Reef (GBR), located along the north-east coast of Australia is made up of over 3,200 reefs and extends over an area of 280,000 square kilometers. The scale of the fluctuation of environmental parameters in the GBR, ranges from kilometre-wide oceanic mixing to millimetre-scale inter-skeletal currents. Being able to monitor such parameters (e.g. temperature, light, etc.) at real-time and at a high spatial and temporal resolution would enable scientists to better understand the underlying complex environmental processes that help shape the behaviours of the biological and physical characteristics of the GBR.

The Australian Institute of Marine Science (AIMS), which is one of the world's leading research centres focusing on marine environments is at the forefront of carrying out research on various aspects of the GBR. Currently AIMS has a test site at Davies Reef which is located around 80km from the shore of the main AIMS research facility at Cape Ferguson. In the present set-up at Davies Reef, temperature data is logged using one data logger that has two sensors attached to it at two different depths. Samples are taken once every 30 minutes and stored within the data-logger. This data is subsequently transmitted to AIMS via a 3.3MHz HF radio.

The main drawback of the current setup is that it only allows single point measurements. This makes it impossible to get a true representation of the temperature gradients spanning the entire reef

Fig 1

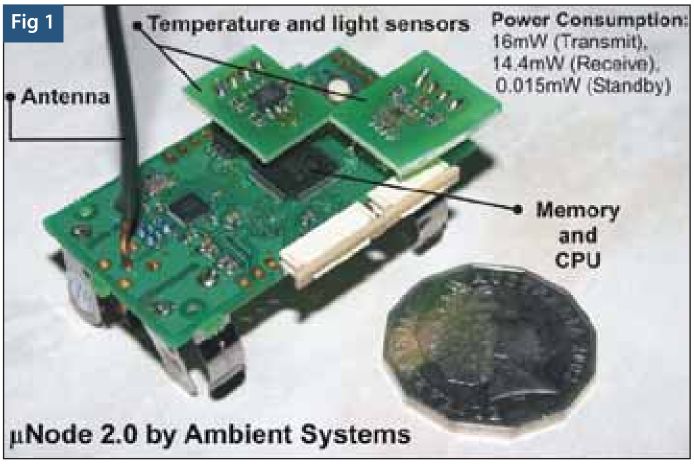

A typical example of a wireless sensor node

Fig 2

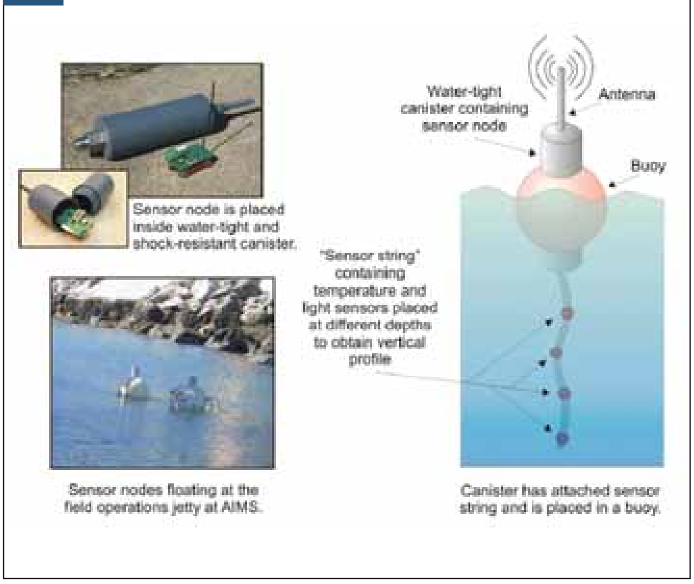

Buoys fitted with "sensor strings"

Fig 3

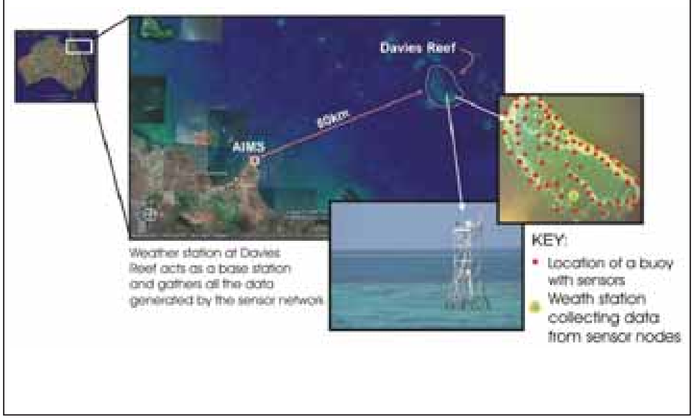

Data collected by the sensor network is transmitted to AIMS via a microwave link.

Fig 4

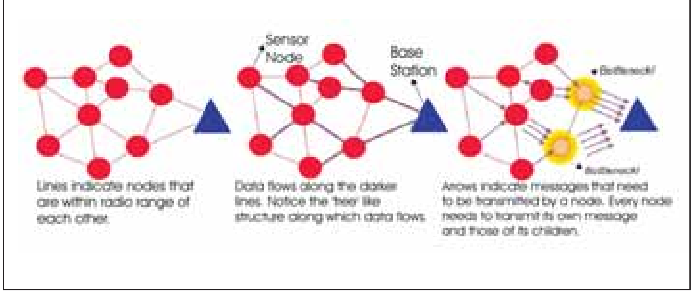

The disadvantage of collecting raw data.

which is approximately 7 by 3 kilometres. The bandwidth limitation of the HF radio also makes it impossible to study the fluctuations of various environmental parameters in real-time.

Having high-resolution data streaming in from the reef would not only improve understanding of various environmental processes (e.g. coral bleaching) but would also have immediate benefits to society at large. For example, as the data will be made available to the public, the tourism industry will be able to tap into the information to better understand diving conditions. The fishing industry will also benefit from knowing how temperature changes and gradients affect fish behaviour. Data collected by the sensors could also be used to study the effects of excess fertiliser washed off from the agricultural lands on the coast of Queensland.

Researchers at the University of Twente and Ambient Systems in the Netherlands have been working closely with scientists at the Australian Institute of Marine Science (AIMS) to develop a large-scale wireless sensor network (WSN) that will enable the GBR to be monitored at extremely high spatial and temporal resolutions. Monitoring the physical parameters of an environment at such high resolutions is something that is currently not possible using existing data logging hardware.

What are Wireless Sensor Networks?

WSNs are typically made up of hundreds or even thousands of tiny energy-efficient, battery operated sensor nodes with built-in wireless transceivers (Fig. 1). Every sensor node also has a CPU, a small amount of RAM and a number of general purpose I/O connections to connect various types of sensors. It is important to realise that sensors nodes are different from sensors which simply have radio transmitters attached to them. Unlike sensors with wireless transmitters, wireless sensor nodes are able to *process* the data within them before transmitting the sampled data due to their built-in computational capabilities. So for example, if a sensor node acquires a reading that is faulty, it could decide to drop the message instead of wasting energy transmitting it.

Similarly, a sensor node may even drop duplicate messages. As individual sensor nodes may have limited communication range (e.g. around 50-100 meters) communication in a WSN is typically done through a *multi-hop* network. So instead of transmitting data directly from a sensor node to the base station, data from a sensor node can be relayed through a number of intermediate sensor nodes, before it finally reaches its destination.

This enables the network to cover a much larger geographical area. Figure 2 illustrates how the sensor node (from Ambient Systems) will be placed in a weather proof canister and then into a buoy. Every canister will have a *sensor* string attached to the bottom which will allow a vertical temperature profile to be created.

These buoys, which will then be spread around the entire area of Davies Reef, will transmit the data over to base station located at a tower that has been erected on the sea floor. The microwave communication system on the tower subsequently transmits the collected data to the AIMS research facility 80 km away as shown in (Fig. 3). The fact that sensor nodes are usually battery powered means that the amount of power available for a node to operate properly is highly limited.

Such tight energy usage restrictions mean that energy efficiency has to be of paramount importance when designing any protocols for WSNs. There are generally *two* significant sources of energy consumption in sensor nodes: the operation of the wireless transceiver and the operation of the sensors attached to the sensor nodes. The sensor network platform we are deploying on the GBR takes both of these sources into consideration to ensure that the

Fig 5

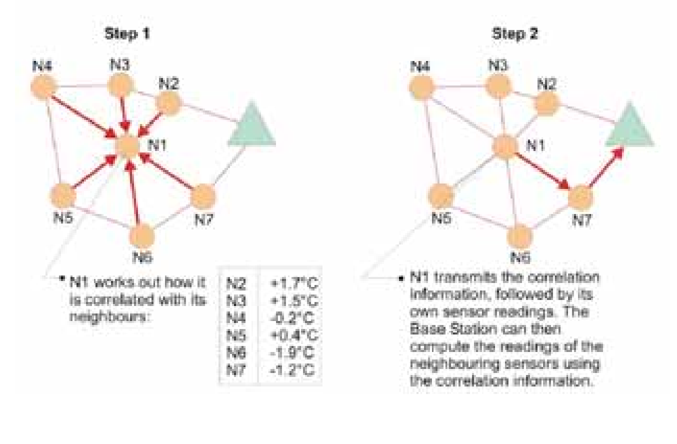

Reducing the amount of data that needs to be transmitted by taking advantage of spatial correlations of sensor readings.

Fig 6

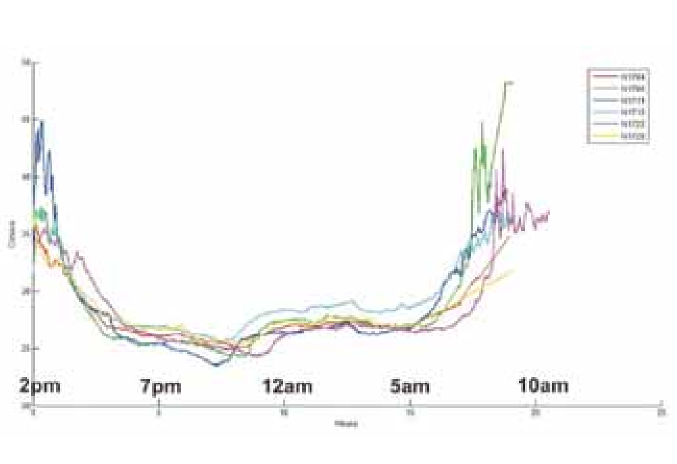

Temperature readings streaming in from 6 sensor nodes deployed at the AIMS campus.

lifetime of the network can be maximised to run for a few years without battery replacement. As mentioned earlier, data collected by sensor nodes is usually transmitted to the base station through intermediate sensor nodes.

This results in a *parent-child* relationship such that communication is performed through a tree structure as shown (Fig. 4). The problem with such a structure is that as the collected data propagates towards the base station, the intermediate nodes that are closer to the base station are required to transmit far greater amounts of data than their counterparts who are further away. This results in two problems. Firstly, due to the larger amount of data that needs to be transmitted, nodes closer to the base station tend to exhaust their energy reserves earlier. This causes the nodes lower down in the tree to be "cut off" thus resulting in a disconnected or partitioned network. Secondly, as the sensor nodes typically have a low bandwidth, nodes closer to the base station are unable to relay all the received data to their parent nodes.

This results in a substantial amount of lost messages which in turn has an adverse effect on the quality of data collected. Processing data within a Wireless Sensor Network In order to alleviate the problems mentioned above, we use a distributed data aggregation algorithm that reduces the amount of data that needs to be transmitted by taking advantage of spatial and temporal correlations that exist in sensor readings between neighbouring sensor nodes. As illustrated in Figure 5, instead of having every node transmit its sensor reading, we assign specific roles to the various sensor nodes – a node can either be an *aggregating node* or a *non-aggregating node*.

Aggregating nodes are put in charge of figuring out whether a correlation exists between itself and its neighbours which are one hop away. For example, the neighbouring node N2 might always be 1.7°C above the aggregating node N1 while N4 might always be 0.2°C below node N1. In such a scenario, the aggregating node N1 would first send the correlation information describing how its readings are related to its correlated neighbours to the base station and subsequently, this would be followed by its own readings.

Thus if for instance the temperature reading of N1 increases, the base station would be able to compute the readings of all the other neighbours of N1 that have correlated readings. Nodes that do not have correlated readings, continue to transmit raw sensor readings. As correlations may not be constant through out the day, our algorithms are able to adapt to varying correlations. Since the network will be deployed in a harsh environment, certain nodes in the network could always fail. However, our aggregation algorithm is designed to work seamlessly even when certain nodes fail or if new nodes are added to the network.

This also ensures that the network is easily scalable. Thus our sensor nodes are capable of making completely autonomous decisions in order to ensure that the network continues to operate properly even with a dynamic network topology.

As we mentioned earlier, our algorithm does not only reduce energy consumption by reducing the operation of the transceiver by minimising the amount of data that needs to be transmitted, but also decreases the number of sensor sampling operations. The environmental parameters that we monitor do not always fluctuate very rapidly. However, there may be certain periods when such fluctuations occur. We take advantage of this pattern by reducing the sampling rates of sensors when sensor readings change gradually (and can be predicted fairly accurately) and increase the sampling rate at other times. Our algorithm also uses a distributed mechanism to "wake up" neighbouring sensors if a particular sensor detects a sudden spike so as not to miss out on any unusual events.

A large-scale sensor network of around 100 nodes is currently being tested around the campus of AIMS and also in the waters around the field operations jetty at Cape Ferguson to fix both software bugs to guarantee that communication is carried out reliably and also to ensure that the hardware (e.g. weather proof canister) is able to withstand the harsh climate that can be experienced out in the open sea. Figure 6 is an extract from the preliminary sensor readings that have been streaming in so far. The initial large scale deployment on the GBR has been planned for the later part of this year.

Supriyo Chatterjea *is located at the Faculty of Electrical Engineerings, Mathematics and Computer Science, University of Twente, P.O. Box 217, 7500AE Enschede, The Netherlands. Email: supriyo@ewi.utwente.nl*

Stuart Kininmonth *is located at the Australian Institute of Marine Science, Townsville, QLD 4810, Australia. Email: s.kininmonth@aims.gov.au*

Paul Havinga *is located at the Faculty of Electrical Engineerings, Mathematics and Computer Science, University of Twente, P.O. Box 217, 7500AE Enschede, The Netherlands. Email: P.J.M.Havinga@ewi.utwente.nl*