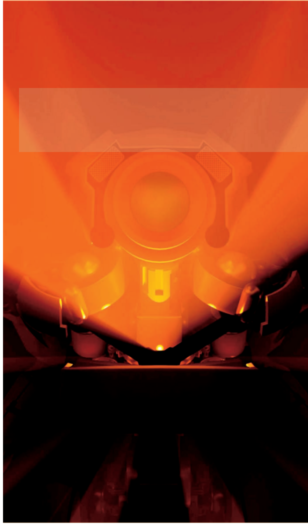

Editor: **Terry R. Payne,** University of Southhampton, trp@ecs.soton.ac.uk

Agent Technologies for Sensor Networks

Alex Rogers and Nicholas R. Jennings, *University of Southampton*
Daniel D. Corkill, *University of Massachusetts Amherst*

Wireless sensor networks are increasingly seen as a solution to the problem of performing continuous wide-area monitoring in many environmental, security, and military scenarios. Such networks consist of small, battery-powered devices that are physically distributed over a wide area and connected through a wireless communication network. Since these networks often must collect data over extended periods of time and are deployed in inhospitable environments where replacing batteries is inconvenient or impossible, much of the research in this domain addresses the challenge of minimizing each sensor's energy needs. To this end, researchers have developed a wide range of energy-efficient sensor nodes and wireless communication protocols and demonstrated them in varied applications.

In addition to the immediate concerns of energy efficiency and reliable wireless communication, which are normally the concerns of electronics engineers, the distributed nature of such networks and the autonomous behavior expected of them present several more generic challenges. Specifically, the individual sensors in these networks must typically coordinate their sensing actions with nearby sensors to achieve systemwide goals (for example, varying their sense/sleep duty cycles to maximize battery life while reducing the redundant sensing of overlapping areas). Furthermore, the network must autonomously adapt its responses in a dynamically changing environment such that it can achieve the long-term systemwide goals without direct human intervention. Such problems have long been the domain of computer scientists, particularly those researching multiagent systems. As such, the multiagent systems community would seem to have an extensive set of formalisms, algorithms, and methodologies in place to address these challenges.

However, the mapping from sensor to agent isn't trivial. Research in the multiagent systems domain typically doesn't address the constrained computational and communication resources of low-powered sensor nodes. Moreover, research often fails to consider that communication might be slow and intermittent, hardware might be unreliable and failure prone, and environments might be highly dynamic. As such, although existing agent technologies are extremely valuable, they cannot be used directly. Rather, to address the specific constraints and challenges posed by this application setting, we need a new synthesis that adapts and extends traditional technologies using approaches from other disciplines such as electronic engineering.

Against this background, this article describes three examples where this synthesis has succeeded. One showcases the development of communication and computationally efficient decentralized coordination algorithms to coordinate the behavior of physically distributed sensors. The second addresses the real-world challenges of using sensor-agent platforms in the field. Finally, we describe intelligent agents that can autonomously acquire data from these networks and perform information-processing tasks such as fusion, inference, and prediction. In each case, researchers have demonstrated the work in the wild, implemented it on real sensor hardware, deployed it in real, hostile environments, and evaluated it on real sensor data.

Agent-Based Decentralized Coordination

As we mentioned, coordinating the activities of physically distributed devices to achieve good systemwide performance is a fundamental challenge. Such coordination might include routing data through the network, choosing the appropriate sampling rates of sensors that exhibit spatial correlations, or determining

1541-1672/09/$25.00 © 2009 IEEE
Published by the IEEE Computer Society

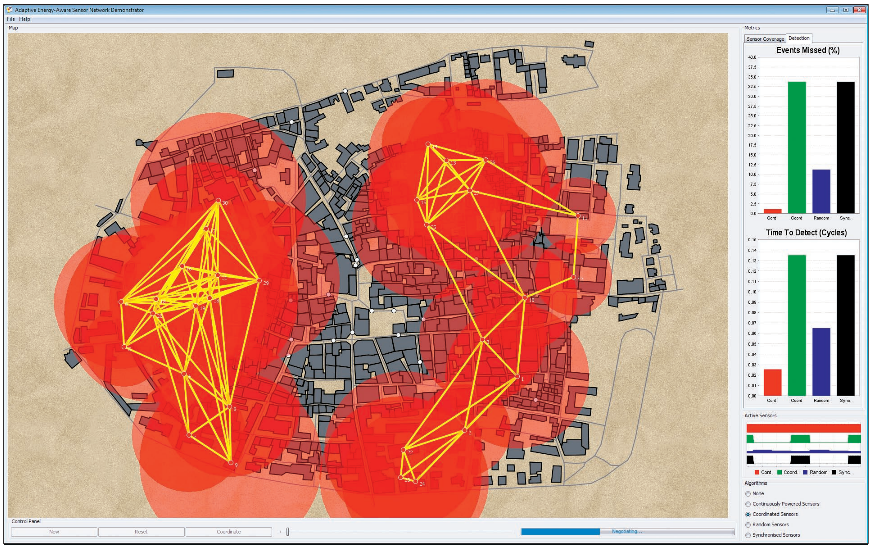

Figure 1. Agent-based decentralized coordination algorithm implemented in a simulated sensor network for wide-area surveillance. The max-sum algorithm enables the coordination of the sense/sleep cycles of energy-constrained sensors.

the scheduling of each sensor's sleep/sense cycle. In each case, we must consider the specific constraints of each device (its limited power, communication, and computational resources) and the fact that each device typically can communicate only with a few other local devices. Furthermore, we should perform this coordination in a decentralized manner so that

- no central point of failure or communication bottleneck exists,
- the computation required for coordination is shared over the distributed resources, and
- the solution scales well as the number of devices in the network increases.

The multiagent systems literature often represents such challenges as *distributed constraint optimization problems* (DCOPs), and researchers have proposed several algorithms that generate optimal solutions to them. Examples include Adopt (Asynchronous Distributed Constraint Optimization), DPOP (Dynamic Programming Optimality Principle), and OptAPO (Optimal Asynchronous Partial Overlay). However, optimality demands that some aspect of these algorithms (either the computational cost or the number or size of messages exchanged) must increase exponentially with the problem size. So, such algorithms are generally unsuitable for sensors that exhibit constrained computational and communication resources. In addition to these optimal algorithms, numerous approximate stochastic algorithms have been proposed for solving DCOPs. These algorithms are typically based on entirely local computation. They maximize a global utility function by having each agent update its state on the basis of the communicated (or observed) states of local neighbors that influence its individual utility. These approaches scale well and are thus well suited to large-scale distributed applications, but they often converge to poor-quality solutions because agents typically communicate only their preferred state, failing to explicitly communicate utility information.

To address this shortcoming, the University of Southampton's Adaptive Energy-Aware Sensor Networks project recently proposed an approximate, decentralized solution that can maximize the social welfare of a group of agents (maximizing the sum of each agent's utilities) when any individual agent's utility depends on its own state and the state of a small number of interacting neighbors.[1] This solution is based on the max-sum algorithm, a message-passing technique that's often used to decompose complex computations on single processors but had never previously been used for multiagent coordination. In particular, this approach exploits extensive empirical evidence that the max-sum algorithm generates good approximate solutions when applied to cyclic graphs. It operates by representing agents' interactions as a factor graph in which each agent—represented by a variable node (representing its state) and a function node (representing its utility)—iteratively passes messages between connected nodes.

An empirical evaluation on a suite of graph-coloring problems (a canonical coordination problem used to evaluate many such algorithms) indicates that this algorithm produces better solutions than approximate stochastic algorithms (such as the Distributed Stochastic Algorithm), that it requires significantly less computational and communication resources than complete algorithms (such as DPOP), and that it's robust to message loss.[1] The researchers at the University of Southampton have implemented the algorithm in a simulated sensor network for wide-area surveillance in an urban environment (see Figure 1 for a screenshot and http://users.ecs.soton.ac.uk/acr/wideareasurveillancedemo for a video of this in operation).

To prove the algorithm's practical applicability, researchers have also implemented it in hardware using the Texas Instruments CC2430 system-on-chip to solve the graph-coloring benchmark problem. The CC2430, an extremely low-power device incorporating a 32-MHz 8-bit 8051 microcontroller and 8 Kbytes of RAM, is intended to form the core of future low-power sensor nodes. Figure 2 shows a simple example in which three graph-coloring sensors have successfully coordinated to avoid a color clash (a video of the sensor nodes in operation is available at http://users.ecs.soton.ac.uk/acr/graphcolouringdemo).

Deploying Sensor Agents in the Field

Developing these results further into sensor-agent platforms for field deployment presents significant additional challenges in adapting to hardware, communication, and environmental limitations. In addressing these challenges, the Collaborative Network for Atmospheric Sensing (CNAS) demonstration project at the University of Massachusetts Amherst has created a rapidly deployable, agent-based, power-aware sensor network for ground-level atmospheric monitoring.[2] Each CNAS sensor agent (see Figure 3 on the next page) uses DARPA's Pasta (Power-Aware Sensing, Tracking, and Analysis) microsensor platform (see http://pasta.east.isi.edu). Despite this platform's relatively low-powered Intel PXA255 processor, specialized Linux operating system, and 64-Mbyte total address-space limit, each sensor agent can run GBBopen (http://GBBopen.org), an open source blackboard system that provides opportunistic AI reasoning at each CNAS sensor agent.

In addition to the Pasta platform, each CNAS sensor agent contains a weather sensor providing temperature, relative humidity, barometric pressure, and wind speed and direction measurements every second. A wireless adapter, connected to the Pasta USB interface, provides standard IEEE 802.11b Wi-Fi communication (which is turned off much of the time to conserve battery power). This hardware and a 12 V battery are packaged in a PVC housing that positions the Wi-Fi antenna and sensors 1.3 m above ground level. The housing is intentionally oversized to enable easy access during testing and field deployment.

Effective CNAS operation requires the sensor agents to make intelligent, proactive decisions regarding how to best use their limited power and communication resources to achieve mission objectives. Specifically, each agent decides what activities it should perform and when to perform them in the context of its own current and projected situation as well as the overall status of the network (the agent organization). For example, the agents must coordinate when they activate their Wi-Fi adapters in order that the high-priority, mission-critical data collected by one agent can be relayed to the base station by a number of other agents. This contextual application and organizational awareness is developed and maintained as part of the GBBopen-based opportunistic reasoning system, and can even be used to improve the network-routing protocols' performance by allowing agents to determine whether they should reinitiate the network discovery process or use stored routing tables, depending on the observed dynamism of the network and the priority and quantity of data to be transmitted. Similarly, this application and organizational awareness allows detailed information regarding network-level communication characteristics such as packet loss and power expenditure to inform high-level operational decision making (by identifying periods of poor communication and delaying the transmission of nonessential data).

This sharing of information, both between sensor agents and at different levels of decision making within a single sensor agent (from the low-level network routing protocol to the highest strategic operational and organizational decisions), is critical to the effectiveness of CNAS and illustrates how these novel approaches depart from more conventional protocols where the application and networking layers are distinct.

To date, CNAS deployments in the field have included one at the 2006 Patriot Exercise at Fort McCoy, Wisconsin,[2] and two (drop zone and urban) at the 2007 Talisman Saber Combined Exercise in Queensland, Australia.[3] Such operational deployments present

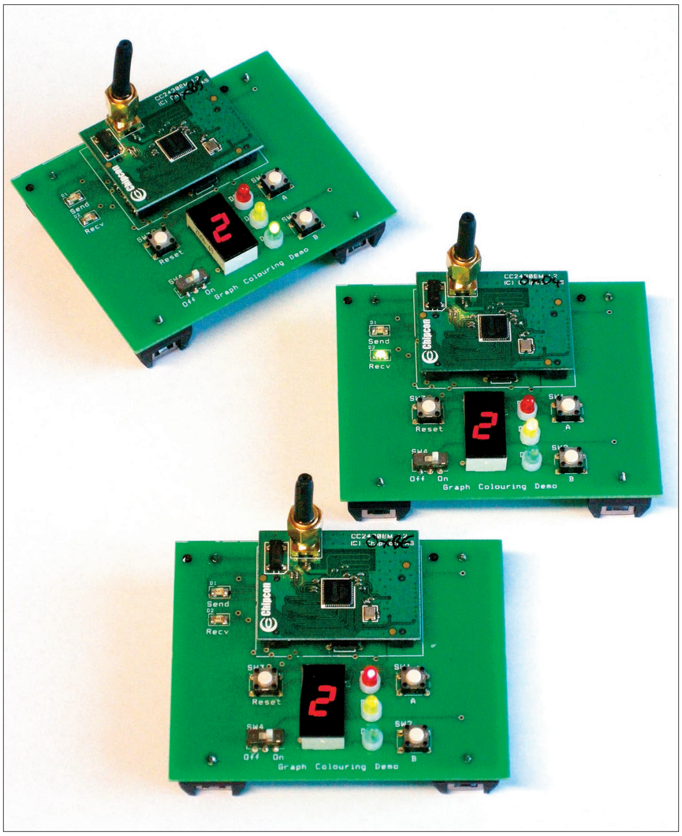

Figure 2. Hardware implementation of the max-sum algorithm and the graph-coloring benchmark problem using the Texas Instruments CC2430 System-on-Chip. The seven-segment display indicates the number of neighbors that each sensor has located, and the three LEDs indicate their respective sensor's chosen color.

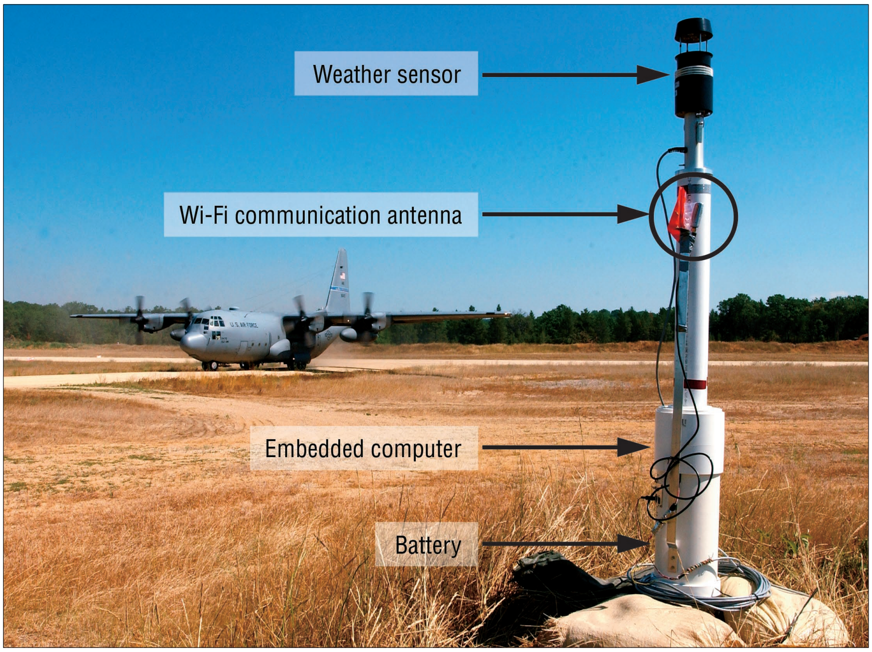

Figure 3. A CNAS sensor agent at the 2006 Patriot Exercise at Fort McCoy, Wisconsin, deployed to collect real-time weather data at a landing strip. (photo courtesy of the US Air Force)

their own challenges. At the Patriot exercise, unseasonably high air temperatures and humidity levels produced transient hardware failures in nearly half of the sensors. These sensor agents returned to full functionality when the temperature dropped, demonstrating how the CNAS sensor agent organization was able to autonomously adapt to dynamic network conditions. The sensor agents encountered extremely poor Wi-Fi communication conditions during the Talisman Exercise because of tall native vegetation at the drop-zone deployment and reflection and interference in the urban deployment. However, as before, no direct intervention was required, because the CNAS sensor agent software was sufficiently robust to recover gracefully during interference-free periods with no data loss, and CNAS was able to provide hourly ground-level weather observations to the Air Force Weather Agency and the Australian Bureau of Meteorology.

Information Agents for Pervasive Sensor Networks

Finally, researchers are also applying agent technology to handle the increasing quantity of real-time sensor data that the pervasive deployment of such sensor networks will provide. In this context, we need *information agents* that can support operational decision making by autonomously acquiring and presenting relevant information to the task at hand. While the notion of such agents isn't new, dealing with real-time sensor network data introduces novel challenges. In particular, these agents must be able to autonomously handle missing or delayed data (perhaps due to network outages), detect faulty sensors, fuse noisy measurements from several sensors, and efficiently manage bandwidth by deciding how often sensor readings must be acquired. Because additional sensors can be deployed at any time, and existing sensors can fail or be repositioned, information agents must perform these processing tasks with minimal prior domain knowledge and, as much as possible, infer details such as a sensor's reliability and accuracy from the data itself.

Against this background, the Aladdin (Autonomous Learning Agents for Decentralized Data and Information Networks) project (www.aladdinproject.org) has demonstrated an information agent that can perform a number of such information-processing tasks.[4] The agent uses a novel iterative formulation of a multi-output Gaussian process to build a probabilistic model of the environmental parameters the sensors are measuring.[5] A Gaussian process enables the agent to apply principled Bayesian inference to functions (in this case, the changing environmental parameters over time), and the probabilistic model that the agent builds lets it infer the sensor readings' accuracy. In this way, the prototype agent can predict both the value of missing sensor readings and how the monitored environmental parameters will evolve in the near future. It also autonomously performs active sampling by automatically determining which sensor to acquire readings from and when.

Researchers at the University of Southampton have evaluated this prototype agent on a permanent network of weather sensors located on England's southern coast. These sensors measure a range of environmental parameters such as wind speed and direction, air temperature, sea temperature, and tide height. While some of the sensors form part of an oceanographic network used to gather long-term coastal data (www.channelcoast.org), other sensors form part of an operational weather sensor network (see Figure 4) used by port authorities for planning shipping movements in the Port of Southampton (see www.southamptonvts.co.uk). Such weather sensors are attractive because scientists can verify the resulting probabilistic model against known meteorological and hydrological phenomena. Furthermore, the sensors are subject to network outages because of extreme weather events, so real instances of missing sensor readings occur, against which the researchers can evaluate the agent's prediction capabilities. A live implementation of this prototype agent is currently available online at www.aladdinproject.org/

Figure 4. The Bramble Bank weather station, located in the Solent. The sensor measures a range of environmental parameters including tide height, wind direction and speed, and air and water temperature.

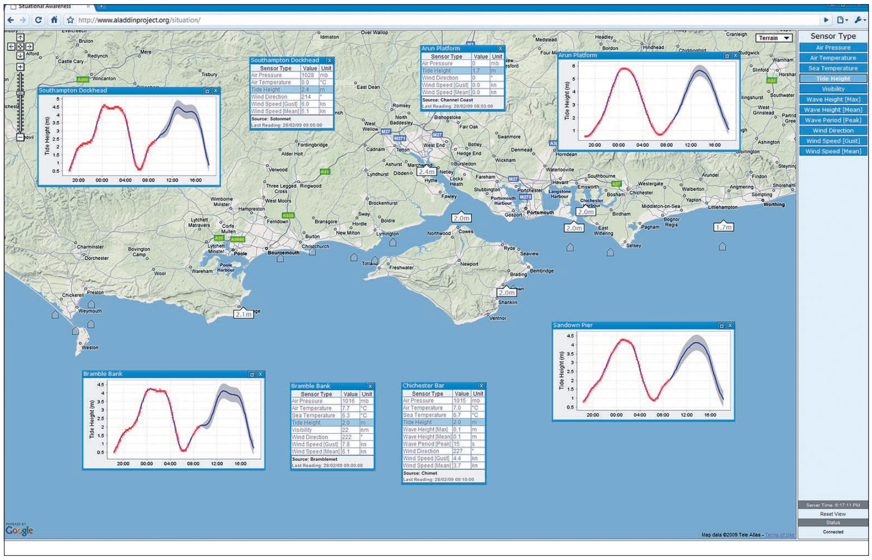

Figure 5. Screenshot of an information agent. A live implementation is available at www.aladdinproject.org/situation.

situation. This implementation applies the Gaussian process predictions to several measured environmental parameters; it also makes available the data collected from the various sensors through an interactive Web-based map (see Figure 5).

The examples described here illustrate that even experimental sensor agent technology has become sufficiently reliable for operational use in the field. Ongoing developments in solar harvesting, coupled with new low-cost processing and sensing hardware, will soon allow the permanent deployment of sensor-agent networks in areas where the economic cost of losing a sensor agent due to damage or theft becomes tolerable. In these cases, many of the information-processing tasks we've outlined can also be delegated to the sensor agents themselves. Doing so will no doubt introduce novel challenges; however, as the CNAS work shows, providing system-level information to individual sensor agents will allow them to make more informed decisions and lead to sensor networks that exhibit more flexibility, robustness, and autonomy.

Acknowledgments

The US Air Force Material Command's Advanced Computing Architecture program supported the Collaborative Network for Atmospheric Sensing (CNAS) work at the University of Massachusetts Amherst under contract FA8750-05-1-0039. The Aladdin and Adaptive Energy-Aware Sensor Network projects are run at the University of Southampton. The first is jointly funded by a BAE Systems and EPSRC (Engineering and Physical Research Council) strategic partnership (EP/C548051/1), and the second was undertaken as part of Phase II of the Data Information Fusion Defence Technology Centre (DIF DTC), funded jointly by the UK Ministry of Defense and General Dynamics UK.

References

1. A. Farinelli et al., "Decentralized Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm," *Proc. 7th Int'l Conf. Autonomous Agents and Multiagent Systems* (AAMAS 08), ACM Press, 2008, pp. 639–646.
2. D.D. Corkill, D. Holzhauer, and W. Koziarz, "Turn Off Your Radios! Environmental Monitoring Using Power-Constrained Sensor Agents," *Proc. 1st Int'l Workshop Agent Technology for Sensor Networks* (ATSN 07), May 2007, pp. 31–38; www.atsm07.org.
3. D.D. Corkill, "Reporting Down Under: A CNAS (Collaborative Network for Atmospheric Sensing) Update," *Proc. 2nd Int'l Workshop Agent Technology for Sensor Networks* (ATSN 08), 2008, pp. 25–32; www.atsm08.org
4. A. Rogers et al., "Information Agents for Pervasive Sensor Networks," *Proc. 6th Ann. IEEE Int'l Conf. Pervasive Computing and Communications* (PerCom 08), IEEE CS Press, 2008, pp. 294–299.
5. M.A. Osborne et al., "Towards Real-Time Information Processing of Sensor Network Data Using Computationally Efficient Multi-Output Gaussian Processes," *Proc. 8th Int'l Conf. Information Processing in Sensor Networks* (IPSN 08), IEEE Press, 2008, pp. 109–120.

Alex Rogers is a reader in the Intelligence, Agents, Multimedia research group in the School of Electronics and Computer Science at the University of Southampton. Contact him at acr@ecs.soton.ac.uk.

Daniel D. Corkill is a senior research scientist in the Department of Computer Science at the University of Massachusetts Amherst. Contact him at corkill@cs.umass.edu.

Nicholas R. Jennings is a professor of computer science in the School of Electronics and Computer Science at the University of Southampton. Contact him at nrj@ecs.soton.ac.uk.