Plug-and-Play Power Resources and Agent-Based Coordination for Energy-Aware Wireless Sensor Nodes

A. S. Weddell, W. T. L. Teacy, N. J. Grabham, G. V. Merrett,
N. R. Harris, A. Rogers, N. M. White and N. R. Jennings
School of Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
{asw,wtlt,njg,gvm,nrh,acr,nmw,nrj}@ecs.soton.ac.uk

***Abstract*– A major constraint of sensor network deployments is their power supply: batteries have a limited lifetime and must be replaced when depleted. Recent advances in the field of energy harvesting mean that sensor nodes can now be powered by environmental energy such as light, vibration, or temperature differences; however, the variety of environments that sensor nodes are deployed into, and their varying levels of power consumption which is dependent on their operation, dictates the type of power supply which must be fitted to the node. This demonstration includes the work done at the University of Southampton in developing a plug-and-play energy architecture for sensor nodes that can accommodate a range of power sources and stores, and agent-based coordination which allows sensor nodes to negotiate between one another to allocate sensing tasks. These capabilities allow the sensor node to be energy-aware, with a flexible energy subsystem, to make best use of their available power. The demonstration is presented in two parts: (i) a plug-and-play energy architecture which is used to power a wireless sensor node, and (ii) a decentralized negotiation algorithm that is deployed on resource-constrained sensor nodes.**

I. Introduction

A. Wireless Sensing

By definition wireless sensors must not rely on a wired power supply, and they are most often powered by non-rechargeable batteries. Batteries are popular as they are cheap and have a high energy density, but they provide a limited amount of energy and must be replaced when depleted. The constrained nature of energy supplies for wireless sensor nodes, along with recent developments in sensor technology, mean that energy harvesting (the generation of electrical energy from other forms of energy in the sensor's environment – such as light, vibration, or temperature difference) is becoming an attractive way of powering sensors for long-term deployments or in situations where changing a battery is expensive or impractical.

Modern wireless sensor nodes utilize low-power microcontrollers such as the MSP430 or 8051. For example the CC2430 system-on-chip device (which is based on an extended 8051 microcontroller) from Texas Instruments typically draws below 1µA when asleep. The device incorporates an IEEE 802.15.4-compliant transceiver and, when it is active and transmitting, draws a maximum of 27mA. It is capable of operating at supply voltages between 2.0V and 3.6V [1]. These properties mean that it is feasible for the operation of the device to be sustained by the power obtained from energy harvesting devices. Examples of devices operating from harvested energy include Prometheus from the University of California, Berkeley, which operates from outdoor solar energy [2]; and the VIBES demonstrator developed by the University of Southampton, which harvests energy from the vibration of machinery [3].

In addition to the immediate concerns of developing low energy sensor hardware, it is also essential that energy is used efficiently across the sensor network. Thus, individual sensors within a network must typically coordinate their sensing actions with nearby sensors to achieve system-wide 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 typically autonomously adapt its responses in a dynamically changing environment such that it can achieve the long-term system-wide goals without the need for direct human intervention.

B. Plug-and-Play Energy Architecture

Very few projects have incorporated multiple energy resources onto a single node. AmbiMax, developed by the University of California, Irvine, is a notable example which combines energy harvesting from wind and light, and stores it in supercapacitors and lithium rechargeable batteries [4]. An advantage of the AmbiMax power module is that it is entirely analogue and autonomous. However, the system design must be adapted to accommodate changes of energy resource. Furthermore, the sensor node powered by the module has no means of finding out the levels of production or availability of energy as the output voltage of the module is fixed at 4.1V. An alternative system is MPWiNodeX, which is capable of using up to three energy sources to recharge a NiMH battery pack [5]. However, the type of energy store cannot be changed, and the energy sources only give a coarse indication of their status (they cannot be actively managed).

Here, we present a demonstrator which allows the energy resources on a sensor node to be connected and configured (at the time of system deployment) without the need to re-program the embedded software on the node. The energy subsystem of the sensor node is split into a number of modules including the multiplexer (which facilitates the scheme). A common hardware interface is defined, which permits the sensor node to communicate with each module individually, and a preliminary electronic datasheet format has been developed which stores device operating parameters in memories on the modules. The scheme as a whole allows the energy hardware of the sensor node to be configured and re-configured in-situ, with the sensor node being able to interrogate the electronic datasheet on each module to determine its operational parameters and learn how to interpret the obtained data (thus achieving system energy-awareness by using these device models to estimate the power generated or energy stored on each module).

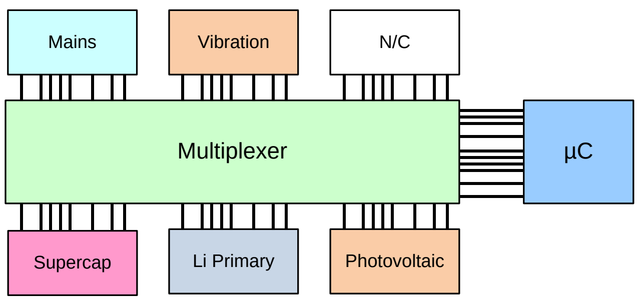

Fig. 1. Example of the type of energy subsystem enabled by our scheme. This incorporates a number of energy devices, connected through a 'multiplexer' to provide power to the microcontroller. As shown, the scheme is flexible and some ports can be left unoccupied with no negative implications to the operation of the system.

This prototype system has been developed, facilitating the connection of a range of energy devices to the sensor node. The scheme allows the sensor node to be powered by a selection of energy resources. Modules have been produced, and are demonstrated, which supply energy to the system from light, temperature difference, air flow, vibration, non-rechargeable batteries, and a mains power adapter. The energy harvested may also be buffered in rechargeable batteries or supercapacitors. The motivation behind this scheme is to allow the energy sources available to be exploited through the connection of appropriate modules. The operational requirements of the node will typically dictate the type and size of energy store that is required. A typical connection scheme is shown in Fig. 1. The major advantage of the plug-and-play functionality is that the system can be installed and configured by a person without an in-depth understanding of the operation of the sensor node. The interface allows up to six energy modules to be used to power the sensor node, and the voltages and type of interface used means that the system can be utilized to power sensor nodes based on a range of low-power microcontrollers.

C. Agent-Based Coordination

As we mentioned earlier, coordinating the activities of physically distributed devices to achieve good system-wide 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 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:

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

We present a demonstrator which implements an agent-based decentralised coordination algorithm on the same resource-constrained sensor nodes used in the other parts of this demonstration. The decentralized algorithm solves a global optimisation problem through local computation and communication, is robust to lossy communication, and exhibits a computational load that scales linearly with the scale of the network (in contrast to the exponential increase observed in alternative approaches). In this demonstration, the sensor nodes are tasked with solving a decentralised graph-colouring problem (a canonical coordination problem which is representative of the type of agent-based negotiation that would take place to distribute sensing or communication operations between sensors in an energy-constrained situation. Each node is equipped with red, green, and yellow LEDs and the algorithm works to ensure that adjacent nodes have different coloured LEDs illuminated.

III. Plug-and-Play Architecture

A demonstration of the plug-and-play energy architecture, which includes a multiplexer module and four energy modules (supercapacitor, battery, photovoltaic and mains), is shown in Fig. 3. A further two energy module sockets on the multiplexer module are available. Energy modules can be connected to any RJ45 socket on the multiplexer module, and are connected here by standard 300mm RJ45 patch leads. Patch leads are used here for ease of use. The energy subsystem shown is connected to a single port on a TI CC2430 evaluation module (EM) via a 10-way IDC cable. The interface with the CC2430 EM is via its interface battery board (used without its batteries installed).

The default behaviour of the system on first installation is to allow the energy harvesting device(s) to charge up the supercapacitor module. Once this store reaches approximately 2.1V, the system connects the power supply to the CC2430, which then starts up and tests its voltage. This voltage is periodically tested until a suitable level has been reached to allow a useful period of operation (nominally 2.7V). At this time the microcontroller is allowed to perform the first energy-intensive tasks such as scanning its energy subsystem. To deliver a near-instant start-up to the system, the 'On' button may be pressed on the primary battery module, which will cause the system to receive power from the battery. Once the microcontroller has taken control of the energy subsystem, the microcontroller can disconnect the primary battery in order to conserve the charge level on the cell. Alternatively, the mains adapter may be turned on, which would also act to rapidly charge up the supercapacitor module; however, the microcontroller cannot act to turn off this supply as it is assumed to be a zero-cost resource which should be taken advantage of whenever it is available.

The first scan of the energy subsystem by the microcontroller is used to ascertain which sockets are occupied, what types of device are present, and their operating parameters (read from their on-board data-sheets). From this initial scan the microcontroller can reach an estimate of the amount of energy stored by the system. This data is stored in the microcontroller memory, so the electronic datasheet on each module only needs to be read once. The microcontroller keeps a record of which sockets on the multiplexer module are occupied. The microcontroller will periodically (at least once per minute when active) re-scan the sockets on the multiplexer and detect any changes. Newly-connected modules can be interrogated for their full datasheet, while disconnected modules are removed from memory and excluded from future calculations.

The datasheet read from each energy module is used by the embedded software on the microcontroller to assess the overall energy status of the system. For example, with the system interconnection shown in Fig. 3, the amount of energy stored in the supercapacitor is estimated by monitoring its voltage and using the capacitance value extracted from the datasheet. The energy remaining in the battery is estimated using a 200Ω impulse load, the results of which are compared against the discharge curve provided in the datasheet for the cell. The status of the mains module is ascertained simply by querying a digital output from the module. The photovoltaic module's nominal power is determined by disconnecting the photovoltaic cell from its load and analyzing its open-circuit voltage. Its nominal power can then be estimated from the cell parameters given in the electronic datasheet for the module.

The microcontroller classifies the energy status of the node as a discrete power priority level. This is shown in Table I. The priority level provides an input to task-oriented algorithms on the node. They permit applications to make decisions about activity levels without those applications having to have a detailed

Fig. 3. Our hardware demonstrator, incorporating a CC2430 sensor node which is powered by a range of energy devices. This scheme is enabled by the multiplexer module, which facilitates the connection of energy sources (such as the photovoltaic and mains module) and stores (such as the supercapacitors). The scheme allows each module to be individually managed by the microcontroller to deliver a truly reconfigurable and energy-aware system.

TABLE I. POWER PRIORITY LEVELS

Priority	Max %	Description
PP_Mains	–	Operating from mains power
PP_5	–	Plentiful energy
PP_4	80	Intermediate energy levels
PP_3	60	
PP_2	40	
PP_1	20	Very limited energy
PP_Empty	2	Cannot sustain activity
PP_Err	–	Error calculating status/unknown

knowledge of the details of the energy subsystem on the sensor node. The percentage values given in the table are the default values used in the demonstrator – these are modified as required.

IV. AGENT-BASED COORDINATION

From the perspective of the -agent systems literature, many of the sensor coordination problems described in the introduction can be naturally represented as *distributed constraint optimization problems* (DCOPs), in which a global optimisation problem must be solved by physically distributed entities with only a limited local view of the entire system. Many decentralised algorithms have been proposed to solve such problems, and several such algorithms are guaranteed to generate optimal solutions. Examples include Adopt

(Asynchronous Distributed Constraint Optimization) [7], DPOP (Dynamic Programming Optimality Principle) [8], and OptAPO (Optimal Asynchronous Partial Overlay) [9]. 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 neighbours that influence its individual utility [10]. 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, we 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 neighbours [6,11]. 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 multi-agent 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. For example, Fig. 4 shows a simple example of three sensors whose states and utilities are decomposed into a factor graph. In this case, the max-sum algorithm effectively solves the following expression:

$$\mathbf{x}^* = \arg\max_{\mathbf{x}} \sum_{m=1}^{M} U_m(\mathbf{x}_m)$$

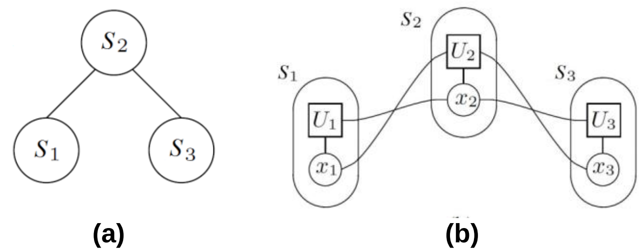

Fig. 4. Example showing (a) 3 interacting sensors and (b) the resulting factor graph representation on which the max-sum algorithm operates.

and thus, effectively finds the states of each sensor in order to maximise the total utility of each sensor, through local message passing and computation.

An empirical evaluation on a suite of graph-colouring 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 [6].

This decentralised algorithm is extremely general and can be applied to any coordination problem. It generates solutions very close to the global optimum, exhibits a low communication overhead, and is robust to lossy communication. To illustrate the practicality of this algorithm we demonstrate its deployment on multiple Texas Instruments CC2430 nodes where it is used to solve the graph colouring problem (a canonical coordination problem that has been widely studied and maps directly to many sensor coordination problems). Each node must choose a colour for itself (indicated by multi-coloured LEDs) to minimise the number of nodes in communication range that share the same colour (see Fig. 5). Furthermore, we have implemented the algorithm in a simulated sensor network for wide-area surveillance in an urban environment, where the algorithm is used to coordinate the sleep/sense cycle of neighbouring sensors to ensure that the effectiveness of an energy constrained sensor network is maximised [12] (see Fig. 6 for a screenshot and www.ecs.soton.ac.uk/~acr/wideareasurveillancedemo for a video of this in operation).

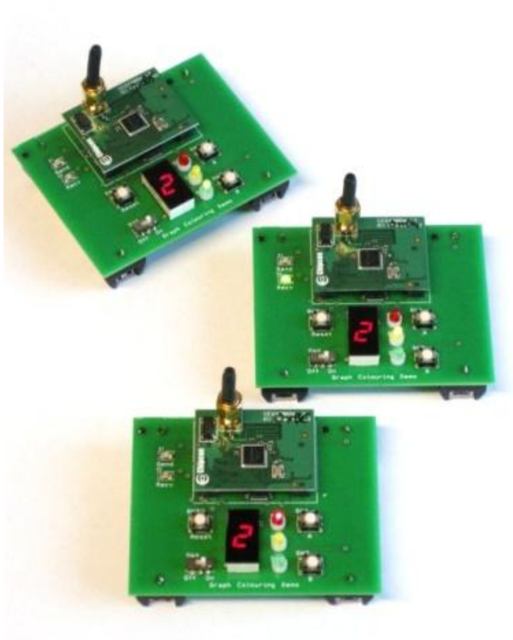

Fig. 5. Agent-based decentralized coordination algorithm implemented in hardware to solve a decentralised graph colouring problem in which the sensors coordinate to chose one of three colours in order that no two neighbouring sensors have made the same choice.

Fig. 6. Agent-based decentralized coordination algorithm implemented in a simulated sensor network for wide-area surveillance in which the max-sum algorithm enables the coordination of the sense/sleep cycles of energy-constrained sensors.

ACKNOWLEDGMENT

Part of the work reported was carried out under the ‘Adaptive Energy-Aware Sensor Network’ project, which was funded by the Data Information Fusion Defence Technology Centre (DIF DTC) and General Dynamics UK Ltd.

REFERENCES

[1] *A True System-on-Chip solution for 2.4 GHz IEEE 802.15.4 / ZigBee™* [Online], Texas Instruments Inc, Dallas, TX, 2008. Available: http://www.ti.com/cc2430/

[2] X. Jiang, J. Polastre, and D. Culler, “Perpetual environmentally powered sensor networks,” *Proc. IPSN 2005*, Los Angeles, CA, 2005, pp. 463-468.

[3] R. Torah, P. Glynne-Jones, J. Tudor, T. O’Donnell, S. Roy, and S. Beeby, “Self-powered autonomous wireless sensor node using vibration energy harvesting,” *Measurement Scence and Technology,* vol. 19, 8 pp., Oct. 2008.

[4] C. Park, and P. Chou, “AmbiMax: autonomous indoor energy harvesting platform for multi-supply wireless sensor nodes,” *Proc. SECON 2006*, Reston, VA, 2006, pp. 168-177.

[5] R. Morais *et al.,* “Sun, wind and water flow as energy supply for small stationary data acquisition platforms,” *Computers and Electronics in Agriculture,* vol. 64, no. 2, pp. 120-132, Dec. 2008.

[6] A. Farinelli, A. Rogers, A. Petcu, and N. R. Jennings, “Decentralised Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm,” *Proc. AAMAS-08*, Estoril, Portugal. pp. 639-646, May 2008.

[7] P. J. Modi, W. Shen, M. Tambe, and M. Yokoo. “ADOPT: Asynchronous distributed constraint optimization with quality guarantees,” *Artificial Intelligence Journal*, (161):149–180, 2005.

[8] A. Petcu and B. Faltings. “DPOP: A scalable method for multiagent constraint optimization,” *Proc. IJCAI’05*, pp. 266–271, 2005.

[9] R. Mailler and V. Lesser, “Solving distributed constraint optimization problems using cooperative mediation,” *Proc. AAMAS’04*, pp. 438–445, 2004.

[10] S. Fitzpatrick and L. Meetrens, *Distributed Sensor Networks: A Multiagent Perspective,* pp. 257–293, Kluwer Academic, 2003.

[11] A. Rogers, D. D. Corkill, and N. R. Jennings, “Agent Technologies for Sensor Networks,” *IEEE Intelligent Systems*, 24 (2). pp. 13-17, 2009.

[12] A. Farinelli, A. Rogers, and N. Jennings, “Maximising Sensor Network Efficiency Through Agent-Based Coordination of Sense/Sleep Schedules,” *Proc. DCOSS 2008*, Santorini, Greece. IV-43-IV-56, June 2008.