A Virtual Global Bus Active Messaging Protocol for Sensor Webs

David Andrews, Joe Evans
University of Kansas

dandrews@ittc.ukans.edu, evans@ittc.ukans.edu

Mixed signal, micro-electro-mechanical, and wireless communication technologies have accelerated our ability to define and build existing and newly emerging hardware platforms into globally and geographically distributed virtual computing systems. The application domain of these new sensor webs is broad, ranging from bio-medical applications, through remote environmental analysis and sensing, to bio-terrorism. Operationally, these future sensor web systems will be comprised of hundreds of thousands of small, autonomous devices that dynamically engage or disengage in the generation of data, and processing of the data into knowledge. Nodes will be deployed in an ad-hoc fashion, with no a-priori knowledge of network and sensor/actuator connectivity. The physical structure of each node will be designed to minimize power, prolong useful deployment time, increase reliability through large numbers, and efficiently process the low level operational mode dictated by the sensors. Behaviorally, node selection, data generation, and signal pre-processing will occur dynamically in order to allow unknown numbers and placement of sensors to be controlled in a system centric coordinated fashion, and increased system reliability.

A challenge exists in developing appropriate network centric models that tightly integrate the computing and communications requirements to "enable (a) accurate distributed sensing, (b) multimode data fusion, (c) transformation from one domain to another, (d) extraction of key information, and (e) detection and circumvention of faulty sensors in ultra large arrays"[1]. To support this challenge we have been exploring new models that

1) Define a new integrated communications/computation machine model for sensor web processors in terms of a Meta instruction set architecture
2) Define a new link layer protocol as a part of the computation/communication fabric that forms a virtual global bus independent of exact sensor locations, types, or numbers. We are investigating this approach as it takes advantage of the law of large numbers in increasing system connectivity, while eliminating power intensive multi-hop route explorations and maintenance of routing tables in sensor node memories.

First, we present background on sensor webs in order to understand their operational domain. After presenting this background, we survey the current state of the art in developing wireless sensor web networks and propose a radically different computation/communication model along with a new supporting link layer protocol.

Requirements for Sensor Web Systems

Conceptually, evolving sensor web systems can be viewed as functional info-spheres of computation, where hundreds of thousands of sensors are data producers, distributed autonomous agent processing algorithms are knowledge generators, timeliness responses are specified on dynamic, non-stochastic knowledge states, and distributed actuators co-ordinate to achieve an enhanced system response. Culler [2] characterizes this new genre of embedded software as being agile, self-organizing, critically resource constrained, and communication-centric on numerous small devices operating as a collective. The operational mode is intensely concurrent, with environmentally bursty activity. The burst rate is projected to be very low, in the 1-10 hertz range. The application space is ubiquitous, spanning numerous devices that interact in a context aware manner. Additionally, power considerations will be a driving factor in the realization of these systems. Current goals include 5 mW/MIP power factors, orders of magnitude lower than current techniques and technology can provide.

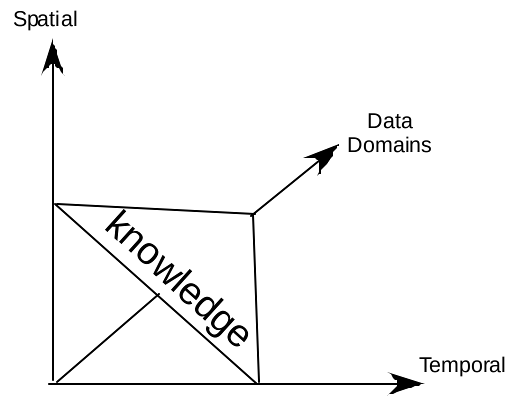

Figure 1. Course Decomposition of Embedded Systems Functional Domains

The quality of a response is both a function of the timeliness, and the quality of knowledge at an instant in time. We can evaluate this model by a course decomposition of the abstract domains that comprise this model as shown in Figure 1.

Computational nodes perform demand driven pass through conditioning of sensor data, and engage in dynamic knowledge formation via data sharing and processing across distributed platforms. This represents a data centric model, with focus on applying operations to data fields and sets. The formation of knowledge is based on retrieval and combination of data across the network in a spatial domain, and timeliness of data updates. Primitive operations, such as domain data selection, reduction and combination are required to support processing within this new paradigm. This is fundamentally opposite to non-embedded models that have "flat" spatial locality and no temporal or domain specific planes. Temporal requirements result from sensor input rates as well as timeliness constraints on the knowledge production and response. The computational model should elevate timeliness information up into the application domain where the relationship between time, precision, and quality of service are best understood.

Sensor Web Architectures

Researchers are investigating clustering and domain selection approaches that allow individual sensors to engage and disengage from the network in a power aware fashion. The sensor information networking architecture (SINA) [3] forms hierarchical clusters of autonomous groups. This clustering process is applied recursively to form the hierarchy. Information is accessed via attribute-based naming instead of explicit addressing. Nodes are queried for information based on attributes, where complex queries can be formed with little overhead within the network. In contrast to SINA, a self-organizing medium access control for sensor networks (SMACS) [8] has been proposed that enables nodes to discover neighbors without the need for master nodes. The SMACS architecture builds a flat topology with no clusters or cluster heads. The eavesdrop and register (EAR) [8] architecture has been developed for communication between mobile nodes and stationary nodes on the ground. To conserve energy, mobile nodes keep a registry of all sensors within a neighborhood and make handoff decisions whenever the SNR drops below a pre-determined threshold value. During a bootup period, invitation messages are broadcast as a trigger. Each mobile node eavesdrops and forms a registry of all stationary nodes within hearing range.

Routing Protocols

Most proposed protocols fall into two main categories: flat routing protocols, and hierarchical protocols. The objective of all routing protocols is to limit node to node communications between pairs of near nodes to reduce power.

Flat Routing Protocols

The first flat routing protocol, sequential assignment routing (SAR) [3], builds multiple routes between a sink and source. This is to minimize the time and cost of computing new routes during failures. A routing tree is built outwards from the sink nodes that attempt to minimize the use of low QoS and energy reserves. Each node belongs to multiple paths and each sensor can control which one-hop neighbor of the sink to use for messaging. The SAR algorithm uses an adaptive QoS metric and a measure of energy resources to arrive at an additive QoS metric and a weight coefficient associated with a packet priority level. During system operation, the SAR algorithm attempts to minimize the average weighted QoS metric. This algorithm requires re-computing paths to account for changes in network topology.

Directed diffusion is a flat routing protocol proposed by Estrin et. al. [4] that allows sensor data to be named. Sink nodes query the sensor web based on data names, and sensor nodes may selectively respond. Intermediate nodes may route data from sensors towards the sink node. Ye et. al. [9] proposed a minimum cost forwarding algorithm for large sensor networks. In this approach, each node contains a least cost estimate between itself and the sink node. Each message in the system is broadcast, and intermediate nodes check to see if they are on the least cost path. If so, the node re-broadcasts the message. Kaulik et. al. [3] proposed the sensor protocols for information via negotiation (SPIN). These protocols disseminate individual sensor information to all sensor nodes under the assumption that all are potential sinks. The solutions use negotiation and information descriptors to overcome the potential information implosion that can be caused by flooding the network with messages sent to all nodes.

Hierarchical Routing Protocols

Chandrakasan et. al. [6] proposed a low energy adaptive clustering hierarchy (LEACH) routing protocol as an energy efficient communication protocol for wireless sensor networks. In LEACH, self-elected cluster heads collect data from all sensor nodes in a cluster, and use data fusion methods to aggregate and transmit the data directly to the base station. The appointment of a cluster head is made periodically with the self appointed cluster head announcing it's role to it's neighbors.

Proposed Link Layer Protocol

Our system approach shares some behavioral commonality with flat routing protocols such as those used in directed diffusion and SPIN. However, we choose to investigate a slightly different approach that seeks higher reliability and greater coverage at what will hopefully result in modest power requirements. Our approach extends the basic principle of Active Messages in TinyOS into a network centric perspective by integrating handler identifiers into the routing of messages at the link layer level. We form communications across the network by utilizing coherent fusion of transmission energies of the sensor nodes to form a virtual global bus. We chose this approach in order to address two unsolved issues. First, we believe that exploring multi hop routes and storing routing information is inherently time and power consuming, just the opposite of what is desired. With the current approaches, receivers are allocated time slots and are allowed to sleep during the intra-periods between allocated slots. This can result in missed messages and coverage. Second, multi-hop routes introduce single point of failure potentials. A multi-hop organization requires the system to constantly perform self-diagnosis and exploration of healthy nodes for new routes. Figure 6 shows the approach we chose to investigate. In this approach, a sender node arbitrates for the bus (as discussed below) and once obtained, continues with transmission of data. All other nodes that are monitoring then assume relay status, and repeat the transmission sequence from the sender. This approach will impose longer data width sizes on the bus. However, prior research indicates that the amount of data being transmitted is relatively modest, allowing longer data cell times. We plan on investigating this tradeoff during the proposal. Finally, we believe our approach is a continued step in the direction of forming a tight computation/communication model. This model allows attributed based access of unknown numbers of sensor nodes, providing flexibility in ad-hoc deployment and operation.

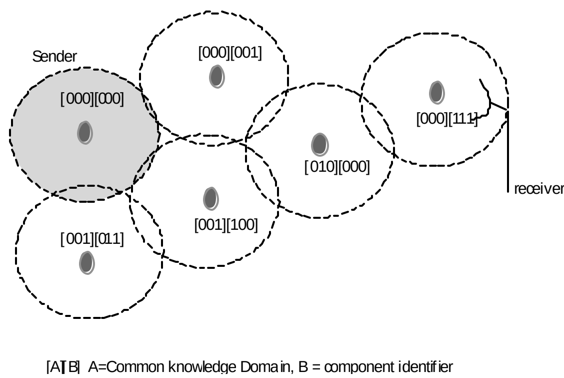

Figure 6. Coherent Domain Coverage

Arbitration/Protocol

Our protocol is similar to the CAN [7] bus in arbitration and protocol but with identifiers serving as handler identifiers ala Active Messaging. This allows fully distributed arbitration between nodes attempting to become bus master. In the CAN protocol, an approach similar to open collector circuits is used to control the logic level on the bus. Multiple masters attempt to drive their identifier onto the bus while also monitoring the success of each bit from the identifier. A logic "0" is dominant, and all nodes attempting to drive a 1 when a logic 0 is present immediately transition to slaves. Within our approach, when a device determines that it is not the bus master, it transitions to a "repeater" device that coherently transmits the data bits that follow on the bus. Communications across the bus result from explicit remote transfer requests on a particular identifier, or may also be initiated by a node when desiring to broadcast data to other listening devices.

Assignment of Identifiers

Identifiers are associated with knowledge domain handlers. Meta (ISA) instructions are defined within a specific domain. The protocol controller matches handler identifiers streaming on the bus to a domain handler match register. The match register determines if a handler is resident on the node. The advantage of this protocol is that it allows instructions, queries, and data to be broadcast within a knowledge domain. A query can easily be issued on an identifier soliciting sensor nodes to transmit data within the knowledge domain.

We have defined Meta Instructions within our ISA that are network operations, such as gather, broadcast and reduction. The gather operation will perform bundling to reduce the overall bandwidth requirements, and the reduction operator will specify an arithmetic or Boolean operation to be applied to the data within the network. We are currently investigating implementation techniques for supporting these instructions across the network. Included in our investigation will be adopting current techniques such as frequency and time multiplexing in order to increase the available bandwidth. On advantage of our approach is that logically it can be implemented across a wide range of physical networks and protocols. Additionally, we will investigate the power tradeoffs between forming coherent bursts, and updating and maintaining routing tables.

Conclusion

In this paper, we have presented the operational mode of emerging sensor web systems. Sensor web nodes will be required to operate autonomously, adapting to ad hoc deployment of thousands of nodes in unpredictable deployment patterns. Of concern within these systems are the network organization and link layer protocol. The link layer protocol must support power conserving operation while still providing reliable communications. Current approaches adopt TDMA based techniques that minimize power draw by shutting down nodes for intra-transmission periods. While this approach minimizes power, it does not encourage higher reliability of connections between multi-hop nodes, and requires power dissipation in path exploration and forming of clusters. We propose a new approach based on the CAN protocol, which offers the promise of higher reliability and greater connectivity. Our link layer protocol is a portion of a new computation/communication integrated machine model for sensor web systems.

Bibliography

[1] http://www.nsf.gov/pubs/2002/nsf02039/nsf02039.html

[2] David Culler et. al, A Network-Centric Approach to Embedded Software for Tiny Devices, Proceedings of the First International Workshop, EMSOFT 2001, Tahoe City, Ca, Oct. 2001

[3] Rentala, Praveen, Musunuri, Ravi, Gandham, Shashidhar, Saxena, Udit, "Survey on Sensor Networks"

[4] Chalermek Intanagonwiwat, Ramesh Govinda, and Deborah Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks" Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (Mobicom'2000), Boston, Massachusetts, August 2000.

[5] Ye F. Chen, A., Liu, S. Zhang L. "A scalable solution to minimum cost forwarding in large sensor networks" Proceedings of the Tenth International Conference on Computer Communications and Networks, pp. 304-309, 2001

[6] Wang, A. Chandrakasan, A., "Energy efficient system partitioning for distributed wireless sensor networks", IEEE International Conference on Acoustics, Speech, and Signal Processing, Vol. 2. pp. 905-908, 2001

[7] http://www.can.bosch.com/index.html

[8] Sohrabi, K., Gao, J. Ailawadhi, V. Pottie, G.J."Protocols for self-organization of a wireless sensor network" IEEE Personal Communications, Vol. 7., Issue 5, Pages 16-27, 2000

[9] Ye F. Chen, A., Liu, S. Zhang L. "A scalable solution to minimum cost forwarding in large sensor networks" Proceedings of the Tenth International Conference on Computer Communications and Networks, pp. 304-309, 2001