

# Semantic Internetworking of Sensor Systems<sup>1</sup>

Ke Wang, Salma Abu Ayyash, Thomas D.C. Little

Department of Electrical and Computer Engineering  
Boston University, Boston, Massachusetts 02215, USA  
(617) 353-9877  
*{ke,saayyash,tdcl}@bu.edu*

MCL Technical Report No. 08-01-2004

**Abstract**– Deployed wireless sensor networks are usually tightly coupled with specific tasks and therefore present unique communication needs. With technological advances that are making easy deployment of wireless sensor networks a reality, mechanisms that allow data exchange between sensor systems deployed at different times and initially for different purposes are necessary. Otherwise, sensors which are co-located and can collaborate with one another to increase capacity and/or performance may be prevented from doing so due to the inability to exchange data with one another.

To this end we propose a semantic-based routing architecture that enables internetworking of sensor systems. In this architecture each sensor system can define its own address names that reflect application-level needs, but which follows a specific naming structure. Semantic routers are elected in the network and process packets based on both the address name and the naming structure. Processing is based on rules that can be modified according to changes in application-level requirements or changes in the naming structure. In this paper we describe this architecture and show that it enables easy internetworking of sensor systems deployed at different times and for different applications, enables data forwarding to only relevant regions in the network (as opposed to network-wide flooding), supports application-level communication needs, and supports hierarchical routing schemes.

---

<sup>1</sup>*Proc. 1st IEEE Intl. Conf. on Mobile Ad hoc and Sensor Systems*, Fort Lauderdale, FL, October 2004.

# 1 Introduction

Current advances in technology increasingly put more processing power and multiple sensing circuits into smaller devices that are then coupled with wireless communication interfaces [9]. Such devices can then be deployed to obtain hitherto unavailable data. Some examples of applications that have been enabled through sensor networks include: soil condition monitoring for agricultural purposes, habitat monitoring and surveillance of the structural conditions of buildings [8].

With the increasing facility experienced in deploying sensor networks, soon multiple users will be able to deploy sensor networks for their own applications over the same geographic region. Communication primitives then are required to enable these networks to interoperate if resource sharing is desired. One readily available communication primitive is flooding. While flooding enables data-centric routing paradigms [11] and allows two sensor networks to exchange data, it does not scale well with increasing network size due to its flat architecture. Given that communication costs are among the highest experienced by sensor devices [25, 6, 24], it is well known that flooding has the disadvantage of spreading packets indiscriminately to all nodes in the network, irrespective of their intended destination.

To address these issues we propose a semantic-based routing architecture. Routing in networks has two essential components: an address naming system and rules that dictate how incoming packets to a host may be forwarded. Such rules are usually based on the knowledge the current host has of the rest of the network and of the intended destination of the packet. The goal of our work is to provide routing that achieves dynamic modification to both such structures and to the rules that describe the routing paradigms. Rules in this architecture process packets based on the underlying structure of the address naming system instead of any instance of a naming system. Thus multiple overlaid addressing systems can be simultaneously supported. With dynamic rule modification, if conditions in the network change, routing rules are updated to reflect application level communication needs. Because the focus is on the meaning of terms of the address system (and not on the value of the terms themselves) and on the relationships among the terms, our architecture is essentially a semantics-based routing architecture.

The advantages of semantic routing is that the address system may be specified in terms of application-level data attributes and the sensor network's routing structure can be established according to attributes most suited to data exchange requirements. In this way data are forwarded only to relevant parts of the network, therefore increasing bandwidth

utilization efficiency and conserving energy. In addition, our work enables the simultaneous resolution of multiple overlaid addressing systems while maintaining a single routing process. Thus location-based addressing schemes, application specific based addressing schemes or IP-based addressing schemes may co-exist in the same physical network simultaneously each successfully delivering packets to their destinations. With semantic routing, parts of the network can be communicating in order to decide where a tracked object's current location is, while simultaneously another subset of sensors is sharing collected data to obtain the average of the variable under observation (e.g., humidity) at a local scope.

We establish the context of our work in Sec. 2, explain our motivation in Sec. 3, describe the components of our architecture in Sec. 4 and conclude in Sec. 7. This paper is essentially a description of our architecture. Currently efforts are being directed to implement the concepts presented in a sensor network testbed. Evaluation results will be published in the future.

## 2 Related Work

In the design of our semantic routing framework we draw on a rich literature on semantic web, web services, peer-to-peer (P2P) networks and content-based routing. In P2P lookup systems such as Tapestry, Pastry, CAN, Chord [28, 19, 17, 22] among others, a data object is associated with a key and information about this association is stored in the network. Name-based routing algorithms allow nodes to locate these storage points for any key. Routing tables become effectively hash tables for various attribute values. Simple Semantic Routing Trees for sensor networks are proposed in many places such as [13, 11], in which queries are forwarded based on the attributes of children nodes. Our work, while enabling the deployment of semantic-based data routing as the ones proposed above, seeks to introduce an architecture that can support a general semantic-based routing mechanism. In addition, we also focus on enabling dynamic modification to the routing rules in order to meet application level communication needs.

Ontologies are at the core of semantic information processing and we use ontologies to define the actions supported by our routing process. There is a vast software technology to support systems that require ontologies. For example, ontologies can be represented using the DARPA Agent Markup Language (DAML) [4]. Sensors can be marked-up with these ontologies using RDF Instance Creator (RIC) [18] for discovery by routers.

Reference [7] presents ideas most similar to ours. In this work, the authors implement the sensor network as sensor web services that are dynamically integrated with other network components. DAML describes the sensor services (ontology of services). Fuselets are lightweight data fusion algorithms that are hosted on various nodes. Every node distributes its fuselet indexing information to other nodes and content-based routing tables are created based on this indexing information. A content-based routing module is responsible for discovering new peers and updating routing tables. The prototype system can dynamically integrate various sensors and fusion services into new applications. Exposing sensor systems as services available to other nodes in the network can be enabled in our framework, but our main focus is on establishing routing architectures within the sensor network that can be used to route data within the sensor network itself and across sensor networks deployed for different applications.

Reference [26] is a recent work which describes a region based communication programming primitive that allows programmers to treat regions as single abstractions. These regions maybe defined by connectivity, location or other properties of the nodes (i.e., they are marked by an ontology of attributes). The routing process in our architecture can benefit from the expressiveness of the abstractions above and implement more efficient routing rules.

### 3 Motivation

Semantic routing implies the use of labeled or attributed data from within a data stream as a basis of routing decisions. By attributed data we mean [attribute, value] pairings. For example, [Dest IP address, 128.157.51.13] represents an association between the destination IP address attribute and a value. Similarly we can have [pressure, 1100]. Complex associations require multiple attribute-value pairs and thus the need to define a framework for their representation. Ontologies are used to describe the essential concepts of our routing architectures while XML-Schema like documents are used to define the data structures [12].

With a representation of attributes, it is then possible to define routing tables and associated routing rules to complete our semantic routing approach. Conventional routing tables are indexed by destination and define next hop outgoing links (or nodes) and accessory such as weighting, sequence number (for updates), etc. For our proposed attribute-based routing we organize the table based on frequently used attributes. Data arriving from neighboring nodes are routed based on the delivered attributes, the routing rules, and any

other application-specific processing (such as embedded functions or consensus rules). Thus, for example, a message arriving with [Dest IP address, 128.157.51.13] can be successfully routed based on any matching entry in the routing table on the corresponding attribute. Similarly, [GPS\_lat,xxxx, GPS\_lon,yyyy] can be routed based on a destination-routing rule.

Consider the following scenario: sensor networks are being deployed over a forest for habitat monitoring. Various research groups expressed interest in evaluating the impact pollution to a lake would have on the drinking habits of the animals living in the forest. For this application scenario, most requests for data would be directed to sensors in the vicinity of the lake, to report the intensity and frequency of animal activity close to the lake as the quality of the water changes.




Figure 1: Flat communication architecture for habitat monitoring sensor networks

Suppose flooding is the underlying communication primitive in the sensor network and the forest and lake are as depicted in Fig. 1. Flooding in these conditions successfully delivers data requests to all the relevant sensors, but every sensor in the network is affected each time a new type of data request is issued. This happens despite the fact that only sensors in the vicinity of the lake need to receive the request.

A more efficient way of directing such requests is to partition the sensors into logical regions, and forward requests that map into regions covering the lake. Such location attributes could be part of the sensors’ capabilities (e.g., sensors with GPS), tagged at deployment time, or updated via restricted flooding started at strategically selected sensors (e.g., *all sensors within 5 hops update location attribute to “near Lake”*).

As an example, consider splitting the forest into quadrants and each quadrant being

split into four sub-quadrants. We then effectively have established a two level quadtree-like partition of the forest. If we name the upper half part of the partitioned region as “T” (top), the bottom half “B” (bottom), the left half as “L” (left) and the right half as “R,” then we have specified an address naming scheme for the sensors in the forest. “Lake” can be mapped to sensors that are defined by [Q1,BL] and [Q2,B], in which Q1 denote the quadrant specification and Q2 the sub-quadrant. Each sub-quadrant would elect one sensor to be the semantic router for the region, and the same would be true for each quadrant. Sub-quadrant semantic routers establish paths to their quadrant semantic routers, while these latter establish paths among themselves. Data requests are forwarded to the closest sub-quadrant semantic router that has the same {Q1,Q2} attributes as the sender. This is illustrated in Fig. 2.




Figure 2: Partitioning of the forest into logical quadrant and sub-quadrant.

In this context, the attributes “L,” “R,” “T,” and “B” are merely address names. The underlying structure is the quadtree, which is reflected in the address naming structure. Each leaf node in the quadtree represents one sub-quadrant specified uniquely by the combination of {Q1,Q2} attributes. Q1 and Q2 represent the two levels of the quadtree, while sensors with the same Q1 attribute belong to the same “parent.”

Given this direct mapping between the structure found in the address names and the topology of the known paths interconnecting the semantic routers, very simple routing rules can be established for data forwarding (such as up\*/down\* [5] algorithms). It is important to note that the names themselves (“L,” “R,” etc.) employed in this address systems serve only to indicate to the routing process whether a particular region has been reached or not. In other words, as long as the naming structure remains the same, and the names are properly identified, i.e., either through explicit pairings (e.g., [Q1, ...], [Q2, ...]) or a pre-specified

order of appearance (e.g. the name for Q1 precedes Q2), the specific names themselves are arbitrary. That is, we can choose to have  $\{TL = 1, TR = 2, BL = 3, BR = 4\}$  and still be able to apply the same routing rules. Thus, an application that uses “.” as name separator and has provision for four levels in the quadtree may well resemble the structure shown in Fig. 3.




Figure 3: Using “1,” “2,” “3” and “4” as names for logical regions. The whole network is named “1.0.0.0”.

Since the routing rules are the same, semantic routers that are aware of both naming systems can function as gateways and forward data requests from one sensor system to the other, as depicted in Fig. 4. While Fig. 4 depicts two sensor systems side-by-side, nothing prevents the two address naming systems to co-exist in the same set of sensor devices. In this way, semantic routers in each sub-quadrant would be aware of two naming systems and applications may choose one addressing name over the other.




Figure 4: Internetworking of two sensor systems.

While this example is straightforward, it is representative of a broad range of semantic routing possibilities. Because the logical regions are formed based on the attributes existent at the sensors, different routing structures based on the same regions can be formed. Thus, instead of the example above (two hierarchical quadtree structures co-existing in the same network with the only difference being the address names), suppose that the two routing structures that co-exist in the same network are: (A) the hierarchical quadtree structure

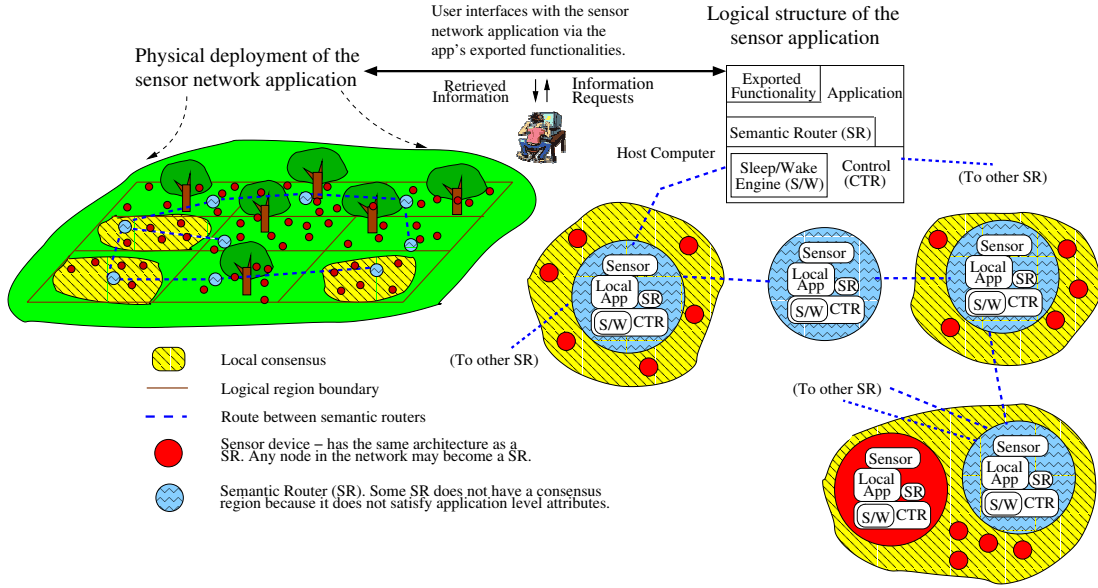


Figure 5: Physical Deployment of the Sensor Network Application

depicted above, and (B) a “mesh” like structure, in which sub-quadrant semantic routers are aware of paths only to adjacent sub-quadrant semantic routers. Routing rules in this condition may take the form of *route first in “T-B,” then in “L-R”* (similar to XY wormhole routing schemes [1]). Applications that are interested in obtaining results based on processing data from the various regions (such as average temperature or humidity) can choose to employ addressing/routing scheme (A), which facilitates in-network data fusion/aggregation, whereas applications such as target tracking may choose to employ addressing/routing scheme (B), which facilitates communication with adjacent sub-quadrants (e.g., requesting them to increase monitoring frequency as the target moves closer to the boundaries of the current sub-quadrant). Hybrid combinations of the two schemes are also possible.

From the example above, different routing schemes can benefit different applications, and a routing architecture that can be adaptive to application-level needs will enhance the performance of the application executed by the sensor network.

## 4 Semantic Routing

Our proposed routing framework is illustrated in Fig. 5. Users interface with the sensor network through their own control points. Sensor networks can be tasked to retrieve information regarding the event they are monitoring through the exported functionality of the application



level. Once an information request reaches the sensor network, application-level collaborations are formed that can solve the problem. Data exchanges occur according to the routing requirements of the applications. The logical structure of a sensor network application is the presence of information flow from the user to the sensors mediated through semantic routers. Once the information request reaches the sensors, application specific routing structures are formed to solve the problem satisfactorily.

In Fig. 5, an application requiring estimation of a variable (e.g., soil moisture) within a local scope is needed. The addressing system employed is a location based scheme, similar to the one in Sec. 3, but in which the deployed area of the sensor network is subdivided into 9 regions. Regions in a column receive attributes of  $\{N$  (North),  $C$  (Center),  $S$  (South) $\}$  while regions in a row receive attributes of  $\{W$  (West),  $C$  (Center),  $E$  (East) $\}$ . One device (that has more resources available) in each region is elected as the semantic router for that region. Regions that have attributes matching those specified by the application can form localized group that can collaborate among themselves to solve the estimation problem. Once an estimated result is obtained the information gathered is sent back to the user. In the following sections we explain in detail the components of our solution and their functionality.

In the case of application-specific attributes (e.g. for intruder detection or estimate local soil conditions), we apply semantic routing to intercept the incoming data and pass them to a locally-executing component of the appropriate distributed algorithm. For example, for a problem requiring building a consensus among sensors in a group (“Is the intruder present?”), the value destined for distribution within a local quadrant of sensor nodes is routed to the application (Fig. 5) where the consensus algorithm is executed. Additionally, the value is propagated (routed) to other nodes in the local quadrant unchanged, or based on the consensus computation, updated and propagated. These routing rules are established and maintained in the routing rules base. More details of the components of our solution follow (Fig. 6).

## 4.1 Semantic Router

Our proposed Semantic Router (Fig. 6) is comprised of a list-based routing table, a routing rules base, a rules engine, and a set of base routing functions provided by the node’s transport function. The routing table (Fig. 6) contains entries associated with each next hop. The illustration of Fig. 6 depicts only one such logical routing table that is currently being supported by the router. The illustrated example follows the example of Fig. 5. Nine square

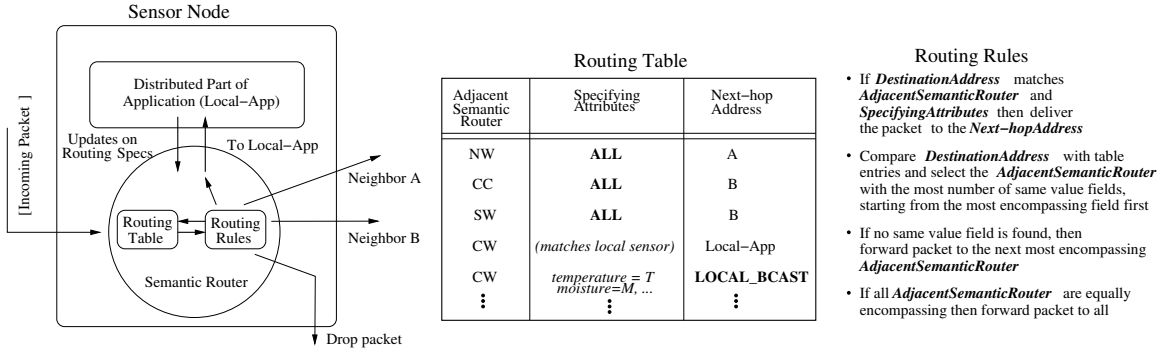


Figure 6: Relationship Among Sensor Node, Routing Table, and Routing Rules

regions partition logically the area on which the sensor network is deployed, and location attributes are primarily used to route data. The first field (most encompassing one) is arbitrarily chosen to be the field for the North-South direction. The routing table belongs to the semantic router for the square “Center + West” (CW). Thus if a packet arrives with “destination” square “CW” and specifies enough attributes for the semantic router to determine it was addressed to itself, it will pass the packet up to the local application. Otherwise, if the list of attributes is not enough to single out a specific device the packet will be flooded to all members which share the “CW” location attribute. Packets that are addressed to other squares are routed first along “NORTH-SOUTH” direction and then along “WEST-EAST,” that is, forwarded to all adjacent semantic routers, since the “NORTH-SOUTH” field does not contain (i.e., is not more encompassing than) the “WEST-EAST” field.

The routing rules base is intended to capture relatively static routing rules associated with the identified attributes. These rules are captured in a tabular format and interpreted and executed by the rules engine. Fig. 6 shows a few entries in the rules base. The rules engine is intended to interpret (thus permitting the rules base to be modified at run time) each rule as it applies to an incoming packet. The rules engine is binary whereas the rules are not. The rules engine calls the base networking functions to achieve handshaking with existing point to point communication protocols available in the sensor hardware.

## 4.2 Routing Extensibility

A major goal of the Semantic Routing is to enable multiple overlaid routing paradigms and the adaptability to changing routing semantics. To permit this flexibility, both the routing attributes and the routing rules must be dynamic and extensible. For this purpose

we choose XML for attribute definition which when used properly allows the introduction of new attributes. New attribute fields may be merely a new instance in the ontology, or a new class if it is introducing new paradigms. These changes can be disseminated by a control point. With respect to the routing rules, we anticipate on-the-fly changes to routing rules in reaction to the localized execution of the consensus algorithms (but not the attributes).

As an example, terms that need be specified and implemented in the ontology include but are not limited to (using an OIL [10] based language):

```
begin-ontology
  ontology-container
    title "Semantic Routing"
  ontology-definitions
    class-def level
    class-def name
    class-def value
    class-def address
      slot-constraint number-of-levels
        has-value (min 1)
    slot-def compose-address
      domain ((address AND address) OR
        ((name AND value)))
      range address
    slot-def number-of-levels
      domain address
      range integer
    slot-def find-address-at
      domain ((address AND integer))
      range (address)
    slot-def address-name
      domain address
      range name
    slot-def address-value
      domain address
      range value
    slot-def oper
      range (one-of True False)
    slot-def binary-value-oper
      subslot-of oper
      domain ((value AND value))
    slot-def binary-address-oper
      subslot-of oper
      domain ((address AND address))
    slot-def adjacent-to
      subslot-of binary-address-oper
      properties symmetric
```

```

slot-def equal-to
  subslot-of binary-address-oper
  properties functional, symmetric
slot-def greater-than-or-equal
  subslot-of binary-value-oper
  properties transitive
end-ontology

```

In the example ontology above (by no means exhaustive), we defined concepts such as *name*, *value*, *address*. We also introduced slots (also may be called *roles* or *attributes* [10]) such as *number-of-levels*, *address-name* and *adjacent-to*, among others. These terms compose the vocabulary upon which routing rules may be built. The routing process fully supports all the terms in the ontology, so if an address A.B.C.D is given to the routing process, it can find out how many levels A.B.C.D contains by calling *number-of-levels*, or whether A.B.C.D is *adjacent-to* A.B.C.E. The ontology supported by the routing process will define the expressiveness of the routing rules. A routing ontology that is not concerned about hierarchies in the structure may well dispense with *level* definitions.

### 4.3 Implementation Issues

In order for the semantic routing scheme to provide energy savings in communication, the underlying routing infrastructure must provide different abstractions that incur in different energy costs when delivering packets, e.g., unicast delivery over multi-hop vs. hop-restricted flooding. It is in the proper matching of the routing abstractions with the intended destinations that energy savings can be achieved. Potential types of destinations within an address system include (1) one single host (destination), (2) multiple co-located hosts or (3) multiple non-co-located hosts.

We chose a clustered approach, in which clusters are formed based on location attributes, as the underlying routing infra-structure on which we build the semantic routing scheme. A hierarchy of attributes and corresponding clusters ensures scalability, which is essential when considering the potential size of sensor networks. Clusters naturally establish localized control boundaries, and cluster leaders are the semantic routers that can coordinate communication among sensors within the cluster. For unicast communications across clusters, if the unicast route is not known, data packets can be forwarded through the semantic routers. Group communication as exemplified by (2) above can also be mapped to clusters, and individual hosts can be tracked by cluster leaders, resulting in support for the case (3) listed in the preceding paragraph.

Maintenance of the hierarchy of clusters requires a fixed cost in terms of message exchanges, which does not depend on the amount of incoming data requests. In fact, by spending energy in maintaining the hierarchical clusters a high volume of incoming requests for different types of data can be better supported, as data requests can be directed to relevant regions. A higher number of levels in the hierarchy requires a corresponding higher maintenance cost, as more clusters are involved. However, this also means that data packets can be more accurately delivered to relevant parts of the network. Dynamic changes in the number of levels or the cluster attributes can effectively support changes in the type of incoming data requests. As described in Sec. 4.2, an ontology specifies the relationship between terms in the routing scheme. Such terms define data structures or functions that can manipulate the data structures used to obtain the information necessary to apply the routing rules correctly. Applications may update the rules base to fit their communication needs. This can be done on-the-fly as long as the terms modified are present in the ontology understood by the semantic router. Sharing different ontologies will be part of our future work.

## 5 Applications and Semantic Routing

In our architecture, the semantic routing layer will export its services directly to the application layer. In addition to support for communication needs that are related to locating the desired sensors, we also propose exporting functionalities that deal with the performance of the network, such as controlling the sensors' Sleep/Wake cycles to improve or degrade metrics such as throughput and/or delay as needed (Sec. 5.1).

The main interface an application has with sensor networks will be a single computer (e.g., PC laptop) bridging the sensor network and the scientist. From this component the sensor system can be tasked with experiments, schedules, routing rules changes and other parameters that affect performance and execution. In addition, through this interface the application is capable of distributing components of itself that are required to execute locally at the sensor nodes (local-apps). These distributed parts of the application enable in-network processing and event triggering and resolution.

### 5.1 Accommodation of Sleep-Wake Schedules of Sensors

The challenge to intelligently manage sleep wake cycles is to both maximize the inactive (asleep) periods of sensor nodes while preserving sufficient active (awake) time to accomplish

sensor tasks of actual sensing, transmission, reception and most relevant to this paper, routing.

With respect to energy conservation, research has shown that wireless communication accounts for the largest portion of energy consumption [23, 25, 6, 24] in a typical system, and maintaining the radios in idle state (not awake or asleep) can be as expensive as the data reception mode. In a complex sensing environment with many sensors and one or more sensing missions, scheduling must be performed in an intelligent and adaptive manner which reflects the needs of the application at hand or *modes* of operation (for example, slow/fast sampling, tracking, localization etc ...).

Within a sensor subnet (defined by say a clustering mechanism) we propose the use of a coordinator (a designated local-app) which can manage the SW schedules for communication purposes such that collision and overhearing are minimized. Simple schedules based on the periodicity of transmissions and priorities can be devised and distributed to all sensors within a subnet. Such schedules address a subset of real-time scheduling concerns. In a more complex system, the task communication model for each mode of operation maybe known in advance permitting a coordinator to centrally devise an initial schedule. Collaboration among a subnet of sensors might be required however, to determine the best schedules given unknown sensed data frequencies or patterns or any local behavior unknown a priori. Therefore, we are working on developing a set of SW schedules based on ontologies and conditions associated with known modes of operation with tunable parameters (such as self-awake frequency and duty-cycle). We envision that this vocabulary will be tied in to the semantic routing ontology thus enabling an efficient SW schedule which addresses the routing needs of the tasks and visa versa (accommodating SW schedules in the selected routes). Such a coupling reduces delays which can be incurred if nodes needed for routing do not have SW schedules which coincide temporally or spatially with the routing requirements (for e.g. nodes required for routing should not be asleep at crucial times/ locations). Building a framework where the SW schedules are directly related to an application and its routing structure is an important aspect of our work.

**Related Work** In general, although the topic of sleep-wake cycle scheduling has received some attention in the literature, we believe that more coupling of application requirements is necessary. Most of the techniques considered focus on link- or MAC-layer mechanisms. Works such as [14, 15, 6] explored ideas of staggered transmission along a query tree in order to reduce contention among nodes. Similarly, [25, 20] adapts 802.11 MAC protocol's

power save mode to sensor networks and try to avoid overhearing to save energy. Slot based power management techniques include works such as [24, 16]. A distinct TDMA approach is presented in [21].

Some SW scheduling techniques take into account the real-time needs of applications. Solutions provided previously include offering packet prioritization via admission control protocols [3], velocity-monotonic scheduling to minimize deadline miss ratios [2] and algorithmic solutions that explore energy-latency tradeoffs [27]. Such techniques must be explored in order to meet the various requirements of our tasks.

## 6 Future Work

Our future work will focus on a general paradigm and framework for semantic routing using extensible data representations for attributes. This will include building routing table data structures based on XML attribute representations and building a semantics-based attribute handling rules for driving a routing engine. The routing engine will also provide common function calls and interpretation of rules-base for processing attributed data under sleep wake schedules.

## 7 Conclusion

In this paper we presented our proposed architecture of semantics based routing to internetwork sensor systems. In semantic routing network addresses are data attributes that are grouped together according to a pre-specified structure. Packets are forwarded by semantic routers according to rules that process the packets based on the structure of their source/destination addresses. Thus new attribute values may be used concurrently as long as the addressing structure is well defined. We also support dynamic modification of the rules base that governs the routing process and enable changes in the Sleep/Wake cycle of sensors to improve or degrade network performance. In this way we established a layer that is expressive enough to support different communication paradigms needed by complex sensor network applications.

The significance of the proposed semantic routing approach is threefold: (1) by limiting the scope of routing within specific subsets of the sensor network there is an opportunity to limit excess energy use otherwise expended on general purpose data routing. In essence, we can limit routing to involve only those nodes relevant for the data flow of the sensing

application, (2) the use of attribute-based routing permits the efficient reconciliation of multiple overlaid routing paradigms and (3) supporting dynamic modification of routing rules enables establishing application specific communication needs. For example, it is possible, using our attribute-based approach, to bridge a location-based routing paradigm with name-based hierarchical one (e.g., IP). In the context of the broader sensor network research, this facilitates routing models corresponding to sensed attributes (e.g., temperature, moisture, etc.) or on derived but orthogonal attributes (e.g., consensus beliefs or other data aggregations and boundary detection). The technique is intended to be flexible and extensible as new attributes are required.

## References

- [1] K. M. Al-Tawil, M. Abd-El-Barr, and F. Ashraf. A survey and comparison of wormhole routing techniques in a mesh networks. *IEEE Network*, 11(2), Mar Apr 1997.
- [2] J. S. C. Lu and et al. “RAP: A Real-Time Communication Architecture for Large-Scale Wireless Sensor networks”. In *Proc. of the IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2002, San Jose, CA, September 2002*. <http://citeseer.nj.nec.com/lu02rap.html>.
- [3] M. Caccamo, L. Y. Zhang, L. Sha, and G. Buttazzo. “An Implicit Prioritized Access Protocol for Wireless Sensor Networks”. In *Proc. of the IEEE Real-Time Systems Symposium*, December 2002. <http://citeseer.nj.nec.com/caccamo02implicit.html>.
- [4] The DARPA Agent Markup Language. URL. <http://www.daml.org/>.
- [5] M. D. S. et. al. Autonet: a high-speed, self-configuring local area network using point-to-point links. SRC Research Report 59, Apr 1990.
- [6] B. K. Gang Lu and C. Raghavendra. “An Adaptive Energy-Efficient and Low-Latency MAC for Data Gathering in Sensor Networks”. In *Proc. of the 4th International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN 04)*, April 2004. To appear.
- [7] W. C. George V. Cybenko, Guofei Jiang. Semantic agent technologies for tactical sensor networks. In *Proceedings of SPIE Conference on AeroSense*, 2003.



- [8] M. Hamilton, B. Reinert, and J. Wallace. Emerging sensor net applications. Invited Panel First ACM International Conference on Embedded Networked Sensor Systems, Nov 2003.
- [9] J. Hill and D. Culler. “MICA: A Wireless Platform For Deeply Embedded Networks”. *IEEE Micro*, 22(6):12–24, Nov/Dec 2002.
- [10] I. Horrocks, D. Fensel, J. Broekstra, S. Decker, M. Erdmann, C. Goble, F. van Harmelen, M. Klein, S. Staab, R. Studer, and E. Motta. The ontology inference layer oil. <http://www.ontoknowledge.org/oil/>.
- [11] C. Intanagonwiwat, R. Govindan, and D. Estrin. “Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks”. In *Proc. International Conference on Mobile Computing and Networking (MobiCom)*, Boston, MA, August 2000.
- [12] M. Klein, D. Fensel, F. van Harmelen, and I. Horrocks. The relation between ontologies and schema-languages: Translating oil-specifications in xml-schema. In *Proceedings of the ECAI’00 workshop on applications of ontologies and problem-solving methods*, Berlin, August 2000.
- [13] S. Madden, M. Franklin, J. Hellerstein, and W. Hong. “The Design of an Acquisitional Query Processor for Sensor Networks”. In *Proc. ACM SIGMOD*, San Diego, CA, June. 2003.
- [14] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. “TAG: a Tiny AGgregation Service for Ad-Hoc Sensor Networks”. In *Proc. Symposium on Operating Systems Design and Implementation (OSDI)*, December 2002.
- [15] S. Madden, R. Szewczyk, M. Franklin, and D. Culler. “Supporting Aggregate Queries Over Ad-Hoc Wireless Sensor Networks”. In *Proceedings of 4th IEEE Workshop on Mobile Computing and Systems Applications*, 2002. <http://citeseer.nj.nec.com/madden02supporting.html>.
- [16] J. C. H. R. Zheng and L. Sha. “Asynchronous Wakeup For Ad Hoc Networks”. In *ACM MobiHoc03*, 2003. <http://www.sigmobile.org/mobihoc/2003/papers/p35-zheng.pdf>.
- [17] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content addressable network. In *Proceedings of ACM SIGCOMM 2001*, 2001.

- [18] RDF Instance Creator. URL. <http://www.daml.org/tools/#RIC>.
- [19] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. *Lecture Notes in Computer Science*, 2218:329, 2001.
- [20] S. Singh and C. S. Raghavendra. “PAMAS: Power Aware Multi-Access protocol with Signaling for Ad Hoc Networks”. *ACM Computer Communications Review*, 28(3):5–26, July 1998.
- [21] K. Sohrabi and G. J. Pottie. “Performance of a novel self-organization protocol for wireless ad hoc sensor networks”. In *Proceedings of the IEEE 50th Vehicular Technology Conference*, pages 1222–1226, 1999.
- [22] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 conference on applications, technologies, architectures, and protocols for computer communications*, pages 149–160. ACM Press, 2001.
- [23] S. P. V. Raghunathan, C. Schurgers and M. Srivastava. “Energy-Aware Wireless Microsensor Networks”. *IEEE Signal Processing Magazine*, pages 40–50, March 2002.
- [24] T. van Dam and K. Langendoen. “An Adaptive Energy-Efficient MAC Protocol for Wireless Sensor Networks”. In *Proc. of ACM Sensys03*, Los Angeles, CA, November 2003.
- [25] J. H. Wei Ye and D. Estrin. “Medium Access Control with Coordinated, Adaptive Sleeping for Wireless Sensor Networks”. In *ACM/IEEE Transactions on Networking*, 2004. to appear.
- [26] M. Welsh. “Exposing Resource Tradeoffs in Region-Based Communication Abstractions for Sensor Networks”. In *Proceedings of the 2nd ACM Workshop on Hot Topics in Networks (HotNets-II)*, November 2003.
- [27] B. K. Y. Yu and V. K. Prasanna. “Energy-Latency Tradeoffs for Data Gathering in Wireless Sensor Networks”. In *Proc. of the 4The 23rd Conference of the IEEE Communications Society (INFOCOM 04)*, Hong Kong, March 2004. To appear.
- [28] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical Report UCB/CSD-01-1141, UC Berkeley, Apr. 2001.