SERVICE DISCOVERY PROTOCOLS AND ROUTING GATEWAYS FOR SDR APPLICATIONS IN SENSOR NETWORKS

Andreas Yankopolus, James Thomson, Peter Sholander and Susan Brookins
Scientific Research Corporation
Atlanta, GA
{ayank, jthomson, psholander, sbrookins}@scires.com

ABSTRACT

While future military networks will rely heavily on IP capable Software Defined Radios (SDRs), such as those being developed for the Joint Tactical Radio System (JTRS) program, microsensors are likely to use lightweight non-IP technologies into the foreseeable future. As such, SDR radio users must be able to seamlessly retrieve readings from both IP and non-IP sensors. This paper describes the use of the Internet Engineering Task Force's (IETF's) Service Location Protocol (SLP) for bridging IP and non-IP sensor networks and the Simple Network Management Protocol (SNMP) for managing and receiving reports from sensors.

1. INTRODUCTION

There has been extensive research on networking techniques for Unattended Ground-based Sensor (UGS) networks [1]. Examples include Adaptive Power Control (APC), reactive routing, clustering, sleep-mode operation, and timer-based Media Access Control (MAC) protocols. However, additional research is needed on the "systems aspects" of combining those point technologies into large-scale fieldable systems. Of particular interest is the problem illustrated in Figure 1. A user in a command center wishes to ascertain the sensor readings for a given sensor type(s) within a given region.

In military networks, the maneuver units may use IP-based software-defined networked radios (e.g, JTRS handheld or vehicle mount radios). However, many future sensor networks will continue to use non-IP nodes for cost and bandwidth reasons. This leads to a requirement that the service discovery protocols and sensor-to-user messaging protocols must span the boundary between those disparate IP and non-IP networks. This requirement occurs whether the users reside in command posts or are tactical users equipped with handheld JTRS radios. As such, the paper examines the use of two Internet Engineering Task Force (IETF) standards for service discovery (Service Location Protocol [3]) and network management (Simple Network Management Protocol [2]) in wireless sensor networks. Of those two protocols, SNMP has been widely proposed for managing individual Software Defined Radios (SDRs). Similarly, SDRs could use SLP to automatically locate various services that exist in their networked environment.

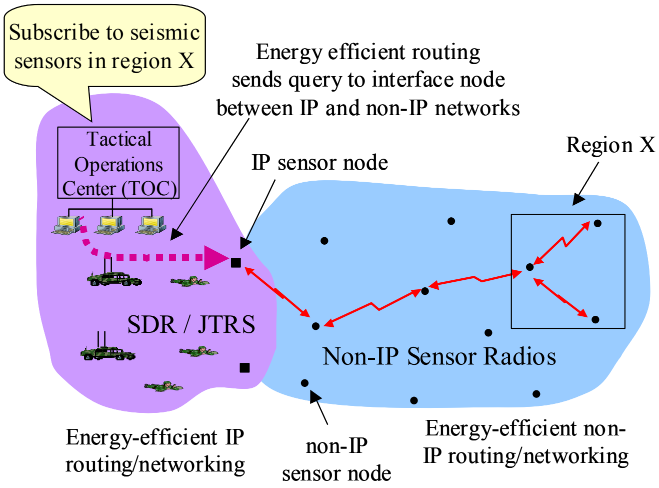
Subscribe to seismic sensors in region X
Energy efficient routing sends query to interface node between IP and non-IP networks
Tactical Operations Center (TOC)
IP sensor node
Region X
SDR / JTRS
Non-IP Sensor Radios
Energy-efficient IP routing/networking
non-IP sensor node
Energy-efficient non-IP routing/networking

Figure 1. UGS network concept of operations

2. SERVICE DISCOVERY PROTOCOLS

2.1 Service Discovery Overview

Service discovery in computer networks has traditionally involved finding the address of a device (e.g., printer, file server, or scanner) and typing it into a configuration file or Graphical User Interface (GUI) window. Service discovery protocols eliminate the need to physically track down devices and manually configure the applications utilizing them. Instead, the user (or application) selects the desired service and any additional attributes, and the service discovery protocol then automatically queries the network for devices matching this description. The devices return both an address and additional descriptive information (e.g, supported paper sizes and physical location for office-based printers).

There are many different service discovery protocols in existence. They include Jini, JiniME, JXTA, MOCA, Rendezvous, Salutation, Service Location Protocol (SLPv2) and Universal Plug and Play (UPnP). This paper focuses on Service Location Protocol Version 2 (SLPv2), which is an Internet Engineering Task Force (IETF) standard. As a brief comparison with the other service discovery protocols, SLP runs on any language -- as opposed to Jini, JiniME, JTXA and MOCA that require Java at each node. Second, SLP has good capabilities for querying the capabilities of sensor nodes. In contrast, Rendezvous, which uses standard DNS packets to exchange service information, lacks the ability to issue detailed queries. Instead, the querying machine would have to sort through numerous replies. Third, SLP has a smaller code footprint than UPnP. Finally, sample code for SLPv2 is available for the BSD operating system. As such, this paper examines the suitability of SLP for sensor networks that contain both Internet Protocol (IP) and non-IP nodes.

2.2 Service Location Protocol (SLP) Overview

SLP [3] relies on three types of software agents to conduct the service discovery process:

- The User Agent (UA) issues service requests and typically runs on the host looking for the service.
- The Service Agent (SA) responds to service requests if it represents a matching service and typically runs on a host that provides a service. SAs do not advertise services but wait until they receive a request from a UA. (This helps avoid period advertisements that can drain the batteries of wireless sensor nodes.)
- The Directory Agent (DA) is a repository of service information and responds to service requests if one of the services registered with it (by an SA) matches the request criteria. The DA is optional and may be located on any machine in the network. (Note: the DA concept can be extended to bridge service discovery between IP and non-IP networks.)

The UA discovers services by multicasting a service request (SrvRqst) message that specifies attributes of the desired service on the network. These attributes are specified using the Lightweight Directory Access Protocol version 3 (LDAPv3) [4] and may contain multiple filters such as Boolean operators and wildcards. In an office network, attributes may include items such as a type (printer, copier, or fax), features (color, duplex, and/or stapler), and floor (1, 2, 3, …). In sensor networks, the attributes could include type (seismic, magnetic, and/or passive IR) and location (latitude, longitude, and altitude). LDAPv3 enables the creation of complex queries using a variety of Boolean operators and the specification of ranges for numeric attributes.

Services return their network location to the UAs by sending a Uniform Resource Locator (URL) that includes the address, port number, and protocol for accessing the service. The URL has the form "service:"<srvtype>"://"<addrspec>. The <addrspec> field is the hostname or dotted quad followed by an optional colon and port number. The <srvtype> field indicates the offered service, such as "service:printer:lpr" for network printing via the lpr protocol.

An example URL for a networked thermometer device (taken from RFC2609) is as follows:

URL = service:net-transducer:thermometer: //v33.test/ports=3211

Attributes =

> *(location-description=Missile bay 32),*
>
> *(operator=Joe Agent),*
>
> *(sample-units=C),*
>
> *(sample-resolution=10e-1),*
>
> *(sample-rate=10),*
>
> *(template-type=service:net-transducer:thermometer),*
>
> *(template-version=0.0),*
>
> *(template-description=The Thermometer is a Net-Transducer capable of reading temperature. The data is read by opening a TCP connection to one of the ports in the service URL and reading an ASCII string until an NULL character is encountered. The client may continue reading data at no faster than the sample-rate, or close the connection.),*
>
> *(template-url-syntax= \0D "ports=" port-list \OD port-list = port / port "," ports \OD port = 1*DIGIT \OD; See the Service URL <port> production rule. \OD; These are the ports connections can be made on.\OD)*

This example is obviously too verbose for use in wireless sensor networks – especially energy-constrained non-IP sensor networks. However, it does illustrate the relative simplicity of the service descriptions. Since SLP runs at the application layer, a binary encoding format could be used to reduce the number of bits actually transmitted over the wire. Section 4.3 discusses this issue in more detail. Section 4 also discusses how to bridge SLP messages into non-IP sensor clouds.

Open source software for SLPv2 is available at http://www.openslp.org. The OpenSLP developer community has verified proper operation of the code on Linux, Win32, FreeBSD, Solaris, Mac OS X, and several other operating systems.

When using that open-source SLP code, an SLP process (slpd) runs in the background to handle SLP network messaging and registration and deregistration of applications. Services register and deregister with slpd by calling SLPReg() or SLPDeReg(), respectively. Users or applications search for services by calling SLPFindSrvs() or SLPFindSrvTypes(). Both registration and searching include support for LDAPv3 service attributes.

3. MANAGEMENT AND REPORTING IN WIRELESS SENSOR NETWORKS

Managing sensors requires a method for reading and setting configurable variables on the sensor node. Reporting requires a way for the sensor node to send an alarm message to a user when it has been tripped. Sensor networks typically use custom protocols for bandwidth efficiency reasons. This paper examines the use of a standards-based ITEF protocol for managing sensors and reporting their readings in networks that contain mixtures of IP and non-IP nodes. This section focuses on IP networks. Section 4 describes the issues associated with transmitting sensor readings from non-IP networks to users at IP-based terminals in a command center and tactical users equipped with handheld, IP-capable SDRs.

The Simple Network Management Protocol (SNMP) provides messages for these management and reporting operations along with a Management Information Base (MIB) for storing configurable values on the managed device. An SNMP agent runs on the device while the manager runs on another computer in the network and allows the user to administer the device.

A sensor system could use SNMPv3 to configure, manage, and receive reports from the sensor nodes. This involves developing a custom enterprise MIB for the sensors that includes entries for the:

- Node itself (e.g., latitude, longitude, and altitude).
- Fused sensor readings (e.g., target azimuth and range).
- Attached sensors (e.g., number and type of sensors along with current readings).
- Proxied sensor nodes (e.g., alarms returned by these nodes).
- Configuration information (e.g., SLP Directory Agent discovery information).
- Registered users (e.g., addresses of users that wish to receive alarm messages).

The SNMP manager can read from and write to the MIB by sending a "get-request" and "set-request" message, respectively, to the agent at the sensor. The "get-request" includes the Object IDs (OIDs) of the MIB variables while the "set-request" includes both the OIDs and the new values. In addition, SNMP provides a "get-next-request" message that is used for stepping through a MIB. In each case, the managed device returns the OIDs and values in a "get-response" message. An SNMP agent informs a manager of an event by sending a "trap" message that includes an event code along with relevant OIDs and their values.

After discovering a sensor node, a user may add herself to the list of registered users, thereby receiving traps when the node determines that one of its sensors has been tripped. (Note: user registrations might be set up to time out if not refreshed in order to avoid sending traps to non-active users.) For example, tripping the seismic sensor attached to a node would cause it to send a trap message to each user registered for seismic alarms. The trap message would include the OID for the seismic sensor along with any data from the sensor itself.

Sensors in the IP portion of the network are administered directly by a standard SNMP manager application. However, nodes in the non-IP portion of the network are neither directly addressable nor powerful enough to run SNMP. The IP nodes must learn about nearby non-IP nodes and add them to their list of proxied nodes. This is explained further in the next section of this paper.

4. SERVICE DISCOVERY OPERATION ACROSS IP AND NON-IP NETWORKS

This section describes how SLP and SNMP can be used to discover services within non-IP sensor networks, and subsequently receive information from those non-IP nodes. The proposed architecture is applicable to both users residing at workstations in a command post and also tactical users equipped with handheld JTRS radios.

Several of the embedded JTRS radios will handle multiple over-the-air waveforms simultaneously, making them candidates for use in the IP-capable sensors. Such radios would allow them to bridge the high-speed waveform carrying IP datagrams and the low-speed waveform carrying DBR interest and data messages to and from non-IP sensor nodes.

4.1 Routing Protocol for Non-IP Networks

Diffusion Based Routing (DBR) [5] is a routing technique for wireless sensor networks that provides "data-centric" routing as opposed to the "address-centric" routing commonly used in IP networking. DBR forwards data between nodes based on data attributes as opposed to a destination address. For example, a user could ask a sensor network for "all acoustic sensor readings from Zone X" as opposed to an "acoustic sensor reading from 192.168.32.114 and 192.168.59.212". Recent academic work has shown that DBR is sometimes more efficient in sensor networks than traditional IP routing and offers significantly shorter messages, which may improve battery life.

Unlike "address-centric" routing schemes like IP where each node requires a globally unique address, nodes running DBR only need a mechanism for distinguishing among their one-hop neighbors. This differentiation could be based on a Media Access Control (MAC) layer address, radio ID, spreading code, frequency, or other unique item associated with a given wireless sensor node.

DBR operates by having one or more nodes ("sinks") propagate an "Interest Messages" across the network. These Interest Messages describe the sensor readings that the sink(s) is interested in receiving. An example Interest Message appears in Table 1. The *Type* field specifies the type of sensor (or threat ID) that the user is interested in. The I*nterval* field (Int.) is the refresh interval for the Interest Messages. The *Duration* field (Dur.) is the time for which the sensor node should retain that Interest Message. The *Coordinates* field is the user's specified region of interest. (Note: SLP's underlying LDAP syntax allows for approximate matches to the coordinates fields.)

Table 1. Example DBR Interest Message

Header				**Payload**
Type	**Int.**	**Dur.**	**Coordinates**	**Data**
Acoustic	60s	1hour	[lat_min, lon_min, lat_max, lon_max]	Varies

Each sensor resends received Interest Messages to all or a subset of its neighbors. This Interest Message appears to come from the transmitting node, even if the original sender was many hops away. Sensors cache recently received Interest Messages for the time specified in the *Duration* field along with the neighbor from which they received them. This information is used to configure initial gradients between nodes along which data may flow back to the sink. The authors in [5] recommend that the metric associated with the gradient indicate the rate at which updates are sent across the link. Alternately, the gradients may be "energy-aware" and seek to form a network between the sources and sinks that has the maximum residual energy-capacity. A sensor may drop an Interest Message that matches one that was recently forwarded; a match occurs if all the fields are equal.

A sensor checks its cache when it receives a sensor alert. If the alert matches a cached Interest Message, the sensor sends the alert data along the stored gradients associated with that message. Sensors next in line along the gradient continue to forward the message until it arrives at an IP node.

Sensors may positively or negatively reinforce gradients associated with incoming messages (by sending the original Interest Message with a greater or smaller interval). A sensor may choose to negatively reinforce a gradient if it receives the same message from two neighbors or positively reinforce a gradient if it appears to provide the highest quality link to the message source.

This process is shown in Figure 2, where the initial interest propagation, gradient set up, and data delivery are shown for an example network.

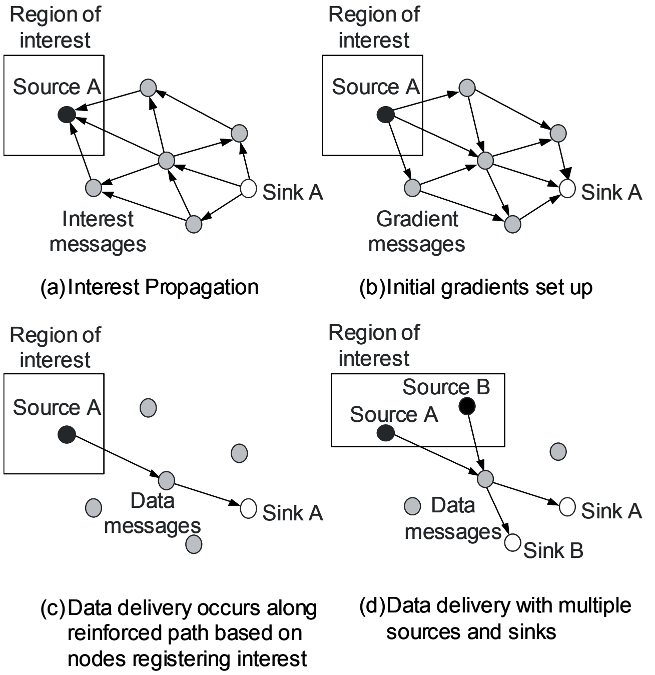

(a) Interest Propagation

(b) Initial gradients set up

(c) Data delivery occurs along reinforced path based on nodes registering interest

(d) Data delivery with multiple sources and sinks

Figure 2. Diffusion Based Routing

4.2 Application-Layer Gateway Between IP and non-IP Sensor Networks

The DBR-based network described in the previous subsection does not share IP's concept of globally unique node identifiers; nodes can only differentiate among their neighbors. Instead, data sources are defined as nodes that match the attributes in an Interest Message, while data sinks as nodes that emit Interest Messages. Nodes do not know the eventual destination for either the Interest Messages or data packets they transmit.

This routing paradigm has similarities with service discovery in IP networks. A node locates services in a network by multicasting a service discovery message. The distribution of the service discovery message throughout the multicast tree is analogous to the initial diffusion of DBR Interest Messages. Similarly, nodes providing the requested service return a service URL to the requestor, which is analogous to the return of a data message to a sink along the gradients set up during the initial diffusion of the DBR Interest Messages.

The sensor network under investigation consists of:

- Non-IP nodes that have a short-range radio and use DBR as their networking protocol.
- IP Nodes that have an SDR that provides both short-range and long-range operation running DBR and IP-based routing respectively.
- SDR users who operate the IP-capable SDR and are end-users of the sensor information.

The IP nodes therefore function as bridges between the IP network (such as the tactical internet) and the DBR inter-sensor network. An example of this architecture appears in Figure 3 where a user within a tactical operations center (TOC), or a tactical user equipped with a handheld SDR, registers for sensor readings within the highlighted region of interest.

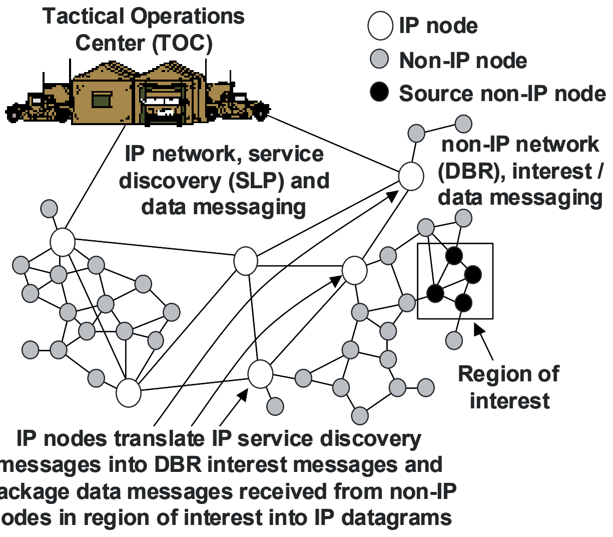
Tactical Operations Center (TOC)
○ IP node
● Non-IP node
● Source non-IP node
IP network, service discovery (SLP) and data messaging
non-IP network (DBR), interest / data messaging
Region of interest
IP nodes translate IP service discovery messages into DBR interest messages and package data messages received from non-IP nodes in region of interest into IP datagrams

Figure 3. Service discovery across IP and non-IP networks

Translating a service discovery packet from the IP network to the non-IP network requires the IP node to convert the contents of the service discovery packet into the DBR protocol's Interest Messages. In addition, the IP node must maintain a table that associates client IP addresses with the parameters in service discovery / interest messages. This table allows the IP node to send data readings received from the non-IP sensor nodes to the correct IP-based client(s). Table entries are cleared after the interval given in the *Duration* field (see Table 1) has elapsed.

In order for service discovery to function properly, the non-IP nodes in the region of interest must send an acknowledgement (for the Interest Messages) along the initial gradient to inform the IP node of their presence. The IP node aggregates those acknowledgements, resolves the DBR header into an IP address and unicasts a service URL (containing the IP address of that IP node) to the TOC or tactical user equipped with a handheld SDR. At no time do the non-IP nodes see the actual SLP message. The user must add himself to the IP node's list of registered users via SNMP in order to receive future messages from the IP node.

Data from the source non-IP nodes may reach multiple IP nodes on its way to the TOC. It is the responsibility of the TOC to identify duplicate sensor readings and possibly deregister from the IP nodes transmitting the duplicate packets. This message causes those IP nodes to suppress messages from the source non-IP sensor nodes by sending negative reinforcement messages along those gradients, thereby eliminating the duplicate messages from the non-IP sensor nodes.

The IP nodes also serve as SNMP proxies for the non-IP nodes. Upon receiving an Interest Message acknowledgement from a non-IP node, the IP node adds a new row to the proxied nodes table in its SNMP MIB. Alarms from a non-IP node in response to the initial DBR protocol's Interest Message are then received by an IP node, where they are forwarded to registered users as SNMP traps using OIDs from this table.

4.3 Data Compression Issues in Wireless Sensor Networks

Both SLP and SNMP are very verbose protocols "over the wire" in that the messages trade size for ease of programming and debugging. This tradeoff results in unacceptably large messages for Military applications.

The compression technique under development uses a data dictionary that associates a 2-octet value with common SLP and SNMP messages needed for sensor networking. For example, an SNMP-get for the sensor node ID includes the OID of the node ID MIB variable (1.3.6.1.4.1.1274.3.1.1.1.0), which takes up 16 octets in the SNMP packet. Furthermore, many values in the packet, such as the Version and Community fields, do not change.

Each 2-octet value in the data dictionary specifies a packet type along with a combination of fixed and variable values. Hence, upon reading the first two octets of a packet, the receiver knows the packet type, the value of several of the fields, and the manner in which to parse subsequent octets. For example, an SNMP-get for a node ID is translated into a 2-octet descriptor followed by a 2-octet truncated request ID. The version, community, PDU type, error status, and error index never change and are therefore not transmitted. Request IDs, which serve to associate requests with responses, do change and must therefore be transmitted. However the full six octets are not needed. Upon receiving the packet, the node looks up the first two octets in the data dictionary and learns that it contains an SNMP-get, the values of fixed fields, and that the following two octets contain a truncated request ID. In this example, an SMNP message with more than 40 octets has been reduced to four octets.

5. CONCLUSIONS

Future military sensor networks will contain mixtures of IP and non-IP nodes. This paper presented a network-layer and application-layer architecture for service discovery and sensor management in that application. The proposed standards-based architecture applies to both IP-based workstations in command posts and tactical users equipped with handheld versions of the Joint Tactical Radio System. Future papers will present results from field trials of the proposed architecture.

6. REFERENCES

[1] "Special Issue on Energy-Aware Ad Hoc Wireless Networks", IEEE Wireless Communications, August 2002.

[2] RFC 3411, An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks. D. Harrington, R. Presuhn, B. Wijnen. December 2002

[3] RFC 2608, Service Location Protocol, Version 2, E. Guttman, C. Perkins, J. Veizades, M. Day. June 1999.

[4] Lightweight Directory Access Protocol (v3): Technical Specification, J. Hodges, R. Morgan. September 2002.

[5] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication for Sensor Networks", "Mobile Computing and Networking", Proc. IEEE MOBICOM, Boston, MA, August 2000.