

SCADA OVER ZIGBEE™

E. Skafidas¹

I. Mareels¹

ABSTRACT

The Zigbee™ alliance seeks to develop an open standard for reliable, cost-effective, secure wireless interconnectivity of monitoring and control products. The ZigBee™ technology is better suited for control applications, which do not require high data rates, but must have low power, low costs and ease of use. In this paper we investigate the applicability of Zigbee™ to Supervisory Control and Data Acquisition (SCADA) systems and investigate issues relating to: Networking, Security, Reliability and Quality of Service.

INTRODUCTION

The Zigbee™ alliance (see [1]) seeks to develop an open standard for reliable, cost-effective, secure wireless interconnectivity of monitoring and control products. In National ICT² Australia (NICTA), Victoria Research Laboratory a significant group of engineers and researchers is actively engaged with the Zigbee™ alliance and the formulation of the open standards. The team (see [2]) is also developing a wireless Zigbee™ ready communication device, NICTOR™³, able to interface with a variety of sensors and actuators.

Devices, such as NICTOR™, could become the back bone of a new generation of low cost, low maintenance SCADA networks, creating new opportunities and enhanced automation services. In this respect, it is worthwhile to observe that the OnWorld Wireless Sensor Network [3] report identifies agricultural monitoring and agricultural SCADA as one of the top 5 market segments for wireless sensor networks in the near future (home-automation, industrial automation, building-automation). NICTA researchers are exploring ways in which wireless sensor networks can be deployed to great advantage in the irrigation industry.

Wireless communications, in particular in the unlicensed part of the radio spectrum comes with its own set of problems: reliability of interconnectivity and security are perhaps the two most obvious issues. It is our thesis that security and data integrity are well catered for in the Zigbee™ standard and that the more

¹ National ICT Australia, Ltd, & Department of Electrical and Electronic Engineering, The University of Melbourne, Victoria 3010, Australia, e-mail contact: i.mareels@unimelb.edu.au, fax: int+61+3+83 44 74 12 (secure)

² ICT stands for Information & Communication Technology

³ NICTOR refers to a particular communication device, it has no other meaning, it is not an acronym.

interesting problem is that of reliability, and quality of communication service, in particular in those applications where communication power is severely constrained.

In this paper we discuss some of the limits of reliability, and focus on how reliability can be addressed through network topology and network management, in particular adaptive radio technology. The same aspects that address the reliability issue make a Zigbee™ network potentially extremely user friendly and offer the potential to make a Zigbee™ sensor network practically application independent. Moreover, new standards, like sensorML, are being developed to create the middle ware which enables the seamless integration of newly developed data services with standard compliant sensors in a very transparent and flexible manner.

The paper is organized as follows. First we discuss in general terms the Zigbee™ network. Next address the security and data integrity issues. Reliability and adaptive radio mechanisms are discussed next. An application in the area of irrigation is presented.

ZIGBEE™ NETWORKING

ZigBee™ is built on top of the IEEE 802.15.4 (see [4]) low power networking standard. The 802.15.4 standard is a simple packet data protocol for lightweight wireless networks. The standard is capable of operating in the unlicensed 2.4GHz band worldwide, 915Mhz band in the Americas and 868Mhz in Europe. There are 16 channels in the 2.4GHz band, 10 in the 915MHz band, and only 1 in the 868Mhz band. The 2450MHz physical layer employs a 16-ary quasi-orthogonal modulation technique. During each data-symbol period, four information bits are used to select one of 16 nearly orthogonal pseudo-random noise (PN) sequences to be transmitted. The PN sequences for successive data symbols are concatenated, and the aggregate chip sequence is modulated onto the carrier using offset quadrature phase-shift keying (O-QPSK). Essentially, this modulation format can be thought of as coded O-QPSK and is typically implemented with a table look-up for generating channel symbols which reduces transceiver cost.

The IEEE 802.15.4 standard has a simple packet frame structure and is specifically designed to enable devices to join and leave the network (Association and Disassociation) as required by the applications on the network (without human intervention). The 802.15.4 standard supports 64-bit IEEE-assigned addresses (assigned in the same manner as an Ethernet addresses, i.e. one preconfigured address per device) and a 16-bit short address, for network-specific addressing (similar to an IP address which is assigned to a device on a local network). This addressing supports a maximum network size of 2^{64} nodes with several thousand devices per node.

The standard supported data rates are 250, 40 and 20 kb/s. The standard supports contention based channel access based on Carrier Sense Multiple Access / Collision Avoidance and contention free Time Division Multiple Access which maintains quality of service (QoS) and bandwidth guarantees by allocating guaranteed time slots. The physical layer medium access control (MAC) protocol is a fully acknowledged protocol for reliable data transfer. It caters for low power consumption, more so than most other wireless standards, by enabling each device to have Energy Detection and Link Quality Indication. The 802.15.4 standard supports Periodic data transfer for applications or devices with an Application defined rate such as sensors that are required to produce periodically a sample. It equally supports an Intermittent data mode to support an Application where an external stimulus defines the need for data transfer. Examples of such applications are a sensor from which data are only required if the monitored variable has changed by more than a preset threshold, or an actuator node, such as an irrigation gate, where the actuation depends on an external event, or Repetitive low latency data. Each of these traffic types mandates different attributes from the Medium Access Controller. The IEEE802.15.4 MAC is flexible enough to handle each of these types. Periodic data can be handled using the beaconing system whereby the sensor will wake up for the beacon, check for any messages and then go back to sleep. Intermittent data can be handled either in a beaconless system or in a disconnected fashion. In a disconnected operation the device will only attach itself to the network when it needs to communicate saving significant energy. Low latency applications may choose the guaranteed time slot (GTS) option. GTS is a method of QoS in that it allows each device a specific duration of time in each Superframe to do whatever it wishes to do (including nothing at all) without contention or latency.

The 802.15.4 standard supports star or peer-to-peer operation. In any 802.15.4 network there are three types of devices

- Network Coordinator: Hands out addresses and coordinates the forming of the network as devices join and leave the network.
- Full Function Device (FFD): Could be a sensor or control node but also provides packet routing and forwarding.
- Reduced Function Device: Does not provide packet routing or forwarding, intended as a low cost node dedicated to a specific purpose such as sensing a condition or controlling a device.

A ZigBee™ network built on the IEEE 802.15.4 standard can support up to 254 client nodes plus one FFD, as master. The protocol is optimized for very long battery life measured in months to years (see [1] for the sales pitch). It has been conceived for building automation, including such diverse items as control for lighting, air conditioning, smoke and fire alarms, and other security devices. The responsibilities of the ZigBee™ network layer include several mechanisms used to join and leave a network, and to route frames to their intended destinations. The routing may involve using multiple intermediate relay devices within the network.

The discovery and maintenance of routes (there is no preset routing table) between devices devolves to the network layer. Also the discovery of one-hop neighbors and the storing of pertinent neighbor information are done at the network layer.

The default routing in a Zigbee™ network is based on a tree-topology. This has the significant advantage that Routers do not have to maintain extensive tables but more importantly do not have to perform route discovery. Notwithstanding this, Zigbee™ Routers have the capability to discover shortcuts, and they can maintain a routing table for these shortcuts of the form (D,N) – (destination, next device). The particular mechanism to discover a shortcut is based on the Request/Response part of the Ad-hoc On Demand Distance Vector (AODV) protocol. The actual network protocol in Zigbee™ is a combination of this AODV, Motorola's Cluster-Tree algorithm and Ember Corporations Gradient Routing in ad hoc networks. Zigbee™ allows for ad-hoc network formation, which makes the network robust to the failure of any one node.

ZIGBEE™ SECURITY

In Zigbee™ networks security is performed at two levels, at the physical layer as implemented in the 802.15.4 MAC and at Zigbee™ network layer. It provides for state of the art time varying, and application dependent security facilities.

The MAC layer security is responsible for single-hop MAC command, beacon and acknowledgement frames. The 802.15.4 medium access controller is based on the Advanced Encryption Standard (AES) as the core cryptographic system. When the MAC layer transmits a frame with security enabled, it Retrieves first a key associated with the destination Processes frame and uses this key according to security mode associated with that key. The Upper layer selects and manages all the various keys and security levels more generally. The MAC frame header has a bit indicating whether security is enabled or not.

In order to protect against other types of attacks such as replay attacks, where an un-authenticated user may re-transmit one or more sensor or actuation messages or repeats intercepted messages or commands, the IEEE 802.15.4 protocol has a message integrity system which detects such messages and ignores these. The MAC header creates a Message Integrity Code (MIC) that consists of 4, 8 or 16 octets and is right-appended to the MAC payload. If confidentiality is required Frame and Sequence counts (nonce) are left-appended. Nonce prevents replay attacks. Upon receipt of a package, the MIC is verified first, after which the payload is decrypted. Sending devices increment the frame count and receiving devices keep track of the last frame received.

The 802.15.4 standard supports security at three levels:

- Encryption at the MAC layer using AES in Counter (CTR) mode, (without authentication)
- Authentication or Integrity AES in Cipher Block Chaining (CBC-MAC) mode (without encryption)
- Combined Encryption and Integrity where the AES implements both CTR and CBC-MAC. This is called CCM.

A Zigbee™ network uses the MAC layer to do the security processing, but it uses the upper layers, to set up the keys, determine the security levels and control the processing. As described previously when the MAC layer transmits (receives) a frame with security enabled, it looks at the destination (source) of the frame, retrieves the key associated with that destination (source), and then uses this key to process the frame according to the security suite or mode designated for the key being used.

The network layer uses a variant of the AES based on the CCM* mode of operation. The CCM* mode of operation is very similar to the CCM mode used by the MAC layer. It includes an encryption and integrity mode of operation, and additionally offers encryption-only and integrity-only capabilities and eliminates CTR and CBC-MAC modes. Also, the use of CCM* in all security modes allows the use of a single key for all different security modes. Since a key is not strictly bound to a single security mode a Zigbee™ application now has the flexibility to specify the actual security mode to be applied to each Network frame, not just whether security is enabled or disabled.

When the Network Layer transmits (receives) a frame using a particular security mode it uses the Security Services Provider (SSP) to process that frame. The SSP looks at the destination (source) of the frame, retrieves the key associated with that destination (source), and then applies the appropriate security mechanisms to the frame. The Network Layer is responsible for the security processing, but the upper layers control this processing as it sets up the keys and determines which CCM* security mode is to be used for each frame.

In conclusion, the Zigbee™ network is secure and immune to most security type of attacks provided that the key distribution integrity is appropriately considered.

ADAPTIVE RADIO AND COEXISTENCE

Where security is not a real issue for a Zigbee™ network, interference is a real problem. This problem is inherent in all wide band communication services over radio. A Zigbee™ network is a radio based network and hence subject to jamming, the more so as low power consumption in the nodes is an important design feature of the Zigbee™ network.

How does one then ensure that the system is robust to jamming and denial of service attacks given the potentially critical nature (e.g. fire alarming and protection in a building) of Zigbee™ networks?

Appropriate enhancements need to be added here to ensure that a SCADA network based on Zigbee™ is robust, both in terms of interference caused by other radio networks that may operate in the same environment or malicious interference such as jamming. First it must be understood that it is not possible to safeguard against a continuous very powerful (as in transmit power) radio jamming signal (but of course such jamming contravenes the use of the radio spectrum). Accepting that the Zigbee™ network is operating within a zone where the use of the radio spectrum is subject to the general rules, which are law enforced, the Zigbee™ network makes use of adaptive radio [5,6], that is a radio that

- monitors its own performance (errors of transmission),
- monitors the path quality through sounding or polling (evaluates the environment for jamming and interference),
- varies its general operating characteristics, such as frequency, power, coding and data rate, or spatial diversity and
- uses closed-loop control action to optimize its performance by automatically adapting from time to time any of the four mechanisms available to combat interference: transmit power, transmit frequency or channel, code and/or spatial diversity.

In this manner a Zigbee™ network ensures that it utilizes the available radio spectrum in the best possible way, without itself causing undue interference to other users of the same spectrum. Within the physical limitations of the radio medium (frequency channels available), the available power and energy, this is the best one can do to alleviate the problems posed by interference caused by other users of the same radio spectrum, be it malicious or otherwise. The precise algorithms underpinning the adaptation of its radio transmission form part of the particular realization of the Zigbee™ network, and may be application dependent. These algorithms are not necessarily open source, nor is it necessary that these are open source, as they are transparent to the user and application layer.

Another important aspect of interference the Zigbee™ network must cope with is the issue created by having potentially in the same radio spectrum licensed and unlicensed users (like Zigbee™ low power nodes). Here the unlicensed users have to cope with any interference of the licensed users, that typically transmit at higher power and yet themselves cannot impede the licensed users in any way. The same ideas of adaptive radio overcome this issue too. These aspects are often referred to as cognitive radio, a radio which is aware of who are the different users of the same spectrum, and what their legitimate claims are on the spectrum, and takes appropriate action to ensure that at any time the appropriate radio frequency usage rules are observed. For further information about cognitive radio, and its potential and limitations we refer to [9].

SENSOR ML AND SENSOR NETWORK PLUG AND PLAY

A critical component of any SCADA system is how easy it is to maintain and develop. Scaling, or expanding a SCADA network is a natural requirement and the incorporation of new sensors and/or actuators in a plug and play fashion are essential to the development of sensor networks in general. It is also important that both sensor/actuator and network management data are readily accessible for archiving and processing. Any closed loop operation of the network will depend on the data being readily available.

The importance of the long-term monitoring of an irrigation system (on time-scales compatible with the physical hardware, i.e. over a number of decades, if not a century) or environment monitoring raises the need for the preservation of the low-level sensor data as well as the information required for processing the raw sensor data over such time scales. To a lesser extent the same applies to actuator nodes, where interface information is needed to determine how information sent to the actuator leads to a physical action. Unfortunately, such information is often lost or difficult to find after the system has initially been deployed. The SensorML (see [7,8]) standard is one step toward preserving part of the vital information required for processing of sensor data for both real-time and archival observations.

SensorML provides a standard means by which sensor and platform capabilities and properties can be published and discovered. SensorML also provides information that allows for processing of the sensor observations without a priori knowledge of the sensor's properties. This provides a significant advantage in that it reduces the time lag between making measurements and applying those measurements in decision-making. Time savings are particularly noticeable in the management of time critical events such as emergency response, advanced warning systems, and forecasting.

Traditionally, low-level sensor data processing has required writing or utilizing software specifically designed for that sensor system. The availability of a standard model language for describing platform position and rotation, as well as instrument geometry and dynamics, allows for the development of generic multi-purpose software that can provide geolocation for potentially all remotely sensed data. The availability of such software, herein referred to as an Observation Dynamics Model, in turn provides a simple, single Application Programming Interface (API) for tool developers to incorporate sensor data and processing into their application software. The SensorML standard allows the development of software libraries that can process these files and data and also for a number of derivative standards concerned with storage and transmission of sensor dynamics, platform location and rotation in order to ensure that such formats are also maintained, available, and readable by similar APIs in the future.

Typically, sensors fall into one of two basic types. In-situ sensors measure a physical property within the area immediately surrounding the sensor, while remote sensors measure physical properties at some distance from the sensor, generally by measuring radiation reflected or emitted from an observed object. Regardless, any geometric properties described within the SensorML schema are defined within the sensor's local coordinate frame and are only related to the geospatial domain through its frame's association with the platform, mount, and their association with some geospatial reference frame.

A SensorML document can be considered a "living" description of a sensor. The SensorML document can begin as a template document, which is initially created using the sensor model design and is then appended or altered during the manufacturing, calibration, deployment, maintenance, and ultimately the removal of the sensor from service. Much of the specification of a sensor is shared by all sensor instances of the same model-number from the same manufacturer. This will typically include a description of measurement properties (with fields for such items as accuracy, scale, raw data format, linearity, repeatability, response rate), sample geometry, and the geometry and dynamics of any internal sampling arrays (such as scan patterns or frame camera properties). This initial template may include in addition some calibration parameters.

However, a SensorML document describing a particular sensor instance will acquire additional information that will distinguish it from other instances of the same model. In particular it may acquire unique identifiers such as ID and serial number. It will further be attached to some platform that will provide it with location and orientation within a known geospatial-temporal frame. In many cases the sensor instance will also have (for example) additional calibration information specific to its deployment (e.g. a datum level for water level measurements). As a sensor progresses through these stages, the SensorML document will not only gain additional property information, but it will also record the changes to the sensor and the document itself through the inclusion of a history description. SensorML will permit new sensors and devices to be easily plugged into a sensor network with minimal effort and maximal reliability.

CONCLUSIONS

Zigbee™ networks are well suited to SCADA applications. Moreover, it is our thesis that mere sensor networks will not provide the real driver for Zigbee™ networks, but rather true SCADA networks where closed loop control is enabled to provide innovative services.

Zigbee™ provides excellent levels of security and data integrity; is extremely end-user friendly as the network auto-configures itself and can be deployed in stages. Unavoidably, because of their low power features, they are subject to jamming and denial of service from a powerful jammer or interference source.

Nevertheless, the adaptive and cognitive radio features, which form an integral part of any Zigbee™ network implementation, provide the best protection against any interference within the physical power/energy limitations of the network devices.

The standardization which comes natural with the communication network is being developed further to encompass also sensors and actuators. (The Zigbee™ specifications are publicly available.) The development of open-source interfaces for a new generation of sensors (like SensorML) and actuators is essential to achieve the true potential of this new technology. Amongst the manufacturers embracing Zigbee™ and SensorML standards are multinationals such as Honeywell International Inc., Philips Electronics NV, Samsung Electronics Co. and Motorola Inc.

We are in the process (see [2]) of implementing a Zigbee™ based SCADA network based on the NICTOR™ platform for low-cost on-farm irrigation and general on-farm automation.

REFERENCES

- [1] <http://www.zigbee.org>
- [2] <http://nicta.com.au/director/research/programs/sn.cfm>
- [3] <http://www.onworld.com/>
- [4] <http://www.ieee802.org>
- [5] <http://www.atis.org>
- [6] <http://au.itpapers.zdnet.com>
- [7] <http://mail.opengeospatial.org/mailman/listinfo/sensorml>
- [8] <http://vast.nsstc.uah.edu/SensorML/>
- [9] <http://www.fcc.gov/oet/cognitiveradio/>