ZigBee Technology: Wireless Control that Simply Works

Patrick Kinney
Kinney Consulting LLC
Chair of IEEE 802.15.4 Task Group
Secretary of ZigBee BoD
Chair of ZigBee Building Automation Profile WG

Kinney Consulting LLC
pat.kinney@ieee.org

ZigBee Technology: Wireless Control that Simply Works

Why is ZigBee needed?

– There are a multitude of standards that address mid to high data rates for voice, PC LANs, video, etc. However, up till now there hasn't been a wireless network standard that meets the unique needs of sensors and control devices. Sensors and controls don't need high bandwidth but they do need low latency and very low energy consumption for long battery lives and for large device arrays.
– There are a multitude of proprietary wireless systems manufactured today to solve a multitude of problems that also don't require high data rates but do require low cost and very low current drain.
– These proprietary systems were designed because there were no standards that met their requirements. These legacy systems are creating significant interoperability problems with each other and with newer technologies.

The ZigBee Alliance is not pushing a technology; rather it is providing a standardized base set of solutions for sensor and control systems.

- The physical layer was designed to accommodate the need for a low cost yet allowing for high levels of integration. The use of direct sequence allows the analog circuitry to be very simple and very tolerant towards inexpensive implementations.
- The media access control (MAC) layer was designed to allow multiple topologies without complexity. The power management operation doesn't require multiple modes of operation. The MAC allows a reduced functionality device (RFD) that needn't have flash nor large amounts of ROM or RAM. The MAC was designed to handle large numbers of devices without requiring them to be "parked".
- The network layer has been designed to allow the network to spatially grow without requiring high power transmitters. The network layer also can handle large amounts of nodes with relatively low latencies.

Kinney Consulting LLC
pat.kinney@ieee.org

ZigBee is poised to become the global control/sensor network standard. It has been designed to provide the following features:

- Low power consumption, simply implemented
  - Users expect batteries to last many months to years! Consider that a typical single family house has about 6 smoke/CO detectors. If the batteries for each one only lasted six months, the home owner would be replacing batteries every month!
  - Bluetooth has many different modes and states depending upon your latency and power requirements such as sniff, park, hold, active, etc.; ZigBee/IEEE 802.15.4 has active (transmit/receive) or sleep. Application software needs to focus on the application, not on which power mode is optimum for each aspect of operation.
  - Even mains powered equipment needs to be conscious of energy. Consider a future home with 100 wireless control/sensor devices,
    - Case 1: 802.11 Rx power is 667 mW (always on)@ 100 devices/home & 50,000 homes/city = 3.33 megawatts
    - Case 2: 802.15.4 Rx power is 30 mW (always on)@ 100 devices/home & 50,000 homes/city = 150 kilowatts
    - Case 3: 802.15.4 power cycled at .1% (typical duty cycle) = 150 watts.

    ZigBee devices will be more ecological than its predecessors saving megawatts at it full deployment.
- Low cost (device, installation, maintenance)

  Low cost to the users means low device cost, low installation cost and low maintenance. ZigBee devices allow batteries to last up to years using primary cells (low cost) without any chargers (low cost and easy installation). ZigBee's simplicity allows for inherent configuration and redundancy of network devices provides low maintenance.
- High density of nodes per network

  ZigBee's use of the IEEE 802.15.4 PHY and MAC allows networks to handle any number of devices. This attribute is critical for massive sensor arrays and control networks.
- Simple protocol, global implementation

  ZigBee's protocol code stack is estimated to be about 1/4th of Bluetooth's or 802.11's. Simplicity is essential to cost, interoperability, and maintenance. The IEEE 802.15.4 PHY adopted by ZigBee has been designed for the 868 MHz band in Europe, the 915 MHz band in N America, Australia, etc; and the 2.4 GHz band is now recognized to be a global band accepted in almost all countries.

Kinney Consulting LLC
pat.kinney@ieee.org

ZigBee/IEEE 802.15.4 - General Characteristics

- Dual PHY (2.4GHz and 868/915 MHz)
- Data rates of 250 kbps (@2.4 GHz), 40 kbps (@ 915 MHz), and 20 kbps (@868 MHz)
- Optimized for low duty-cycle applications (<0.1%)
- CSMA-CA channel access
  - Yields high throughput and low latency for low duty cycle devices like sensors and controls
- Low power (battery life multi-month to years)
- Multiple topologies: star, peer-to-peer, mesh
- Addressing space of up to:
  - 18,450,000,000,000,000,000 devices (64 bit IEEE address)
  - 65,535 networks
- Optional guaranteed time slot for applications requiring low latency
- Fully hand-shaked protocol for transfer reliability
- Range: 50m typical (5-500m based on environment)

ZigBee/IEEE802.15.4 - Typical Traffic Types Addressed

- Periodic data
  - Application defined rate (e.g., sensors)
- Intermittent data
  - Application/external stimulus defined rate (e.g., light switch)
- Repetitive low latency data
  - Allocation of time slots (e.g., mouse)

Each of these traffic types mandates different attributes from the MAC. 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 to the network when it needs to communicate saving significant energy.
- Low latency applications may choose to the guaranteed time slot (GTS) option. GTS is a method of QoS in that it allows each device a specific duration of time each Superframe to do whatever it wishes to do without contention or latency.

Kinney Consulting LLC
pat.kinney@ieee.org

The IEEE 802.15.4 PHY and MAC along with ZigBee's Network and Application Support Layer provide:

- Extremely low cost
- Ease of implementation
- Reliable data transfer
- Short range operation
- Very low power consumption
- Appropriate levels of security

There are two physical device types for the lowest system cost

To allow vendors to supply the lowest possible cost devices the IEEE standard defines two types of devices: full function devices and reduced function devices

– Full function device (FFD)
  - Can function in any topology
  - Capable of being the Network coordinator
  - Capable of being a coordinator
  - Can talk to any other device

– Reduced function device (RFD)
  - Limited to star topology
  - Cannot become a network coordinator
  - Talks only to a network coordinator
  - Very simple implementation

An IEEE 802.15.4/ZigBee network requires at least one full function device as a network coordinator, but endpoint devices may be reduced functionality devices to reduce system cost.

– All devices must have 64 bit IEEE addresses
– Short (16 bit) addresses can be allocated to reduce packet size
– Addressing modes:
  – Network + device identifier (star)
  – Source/destination identifier (peer-peer)

Kinney Consulting LLC
pat.kinney@ieee.org

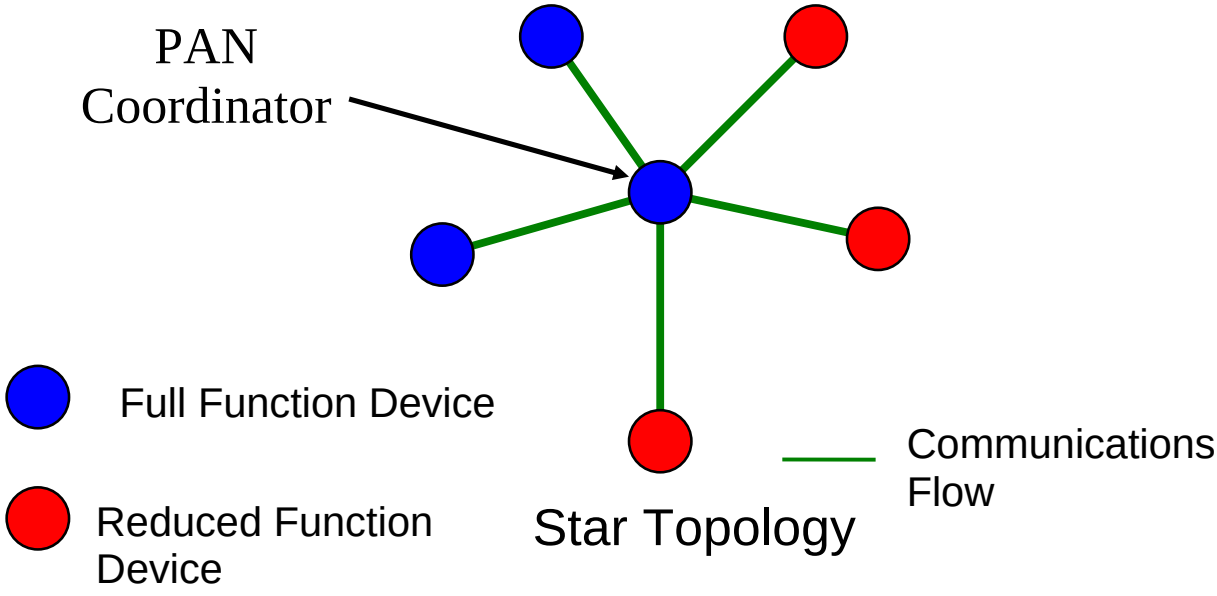

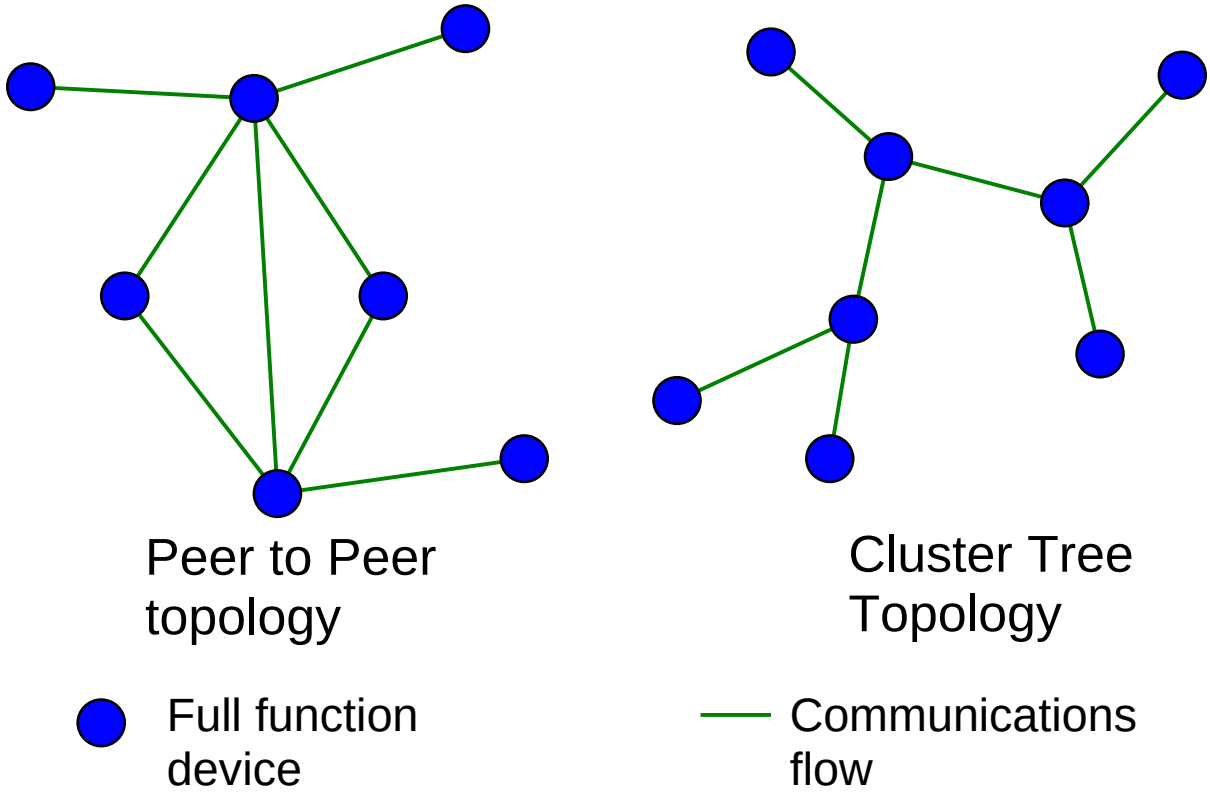

Kinney Consulting LLC
pat.kinney@ieee.org

Frame Structure

The frame structures have been designed to keep the complexity to a minimum while at the same time making them sufficiently robust for transmission on a noisy channel. Each successive protocol layer adds to the structure with layer-specific headers and footers.

The IEEE 802.15.4 MAC defines four frame structures:

- A beacon frame, used by a coordinator to transmit beacons.
- A data frame, used for all transfers of data.
- An acknowledgment frame, used for confirming successful frame reception.
- A MAC command frame, used for handling all MAC peer entity control transfers.

The data frame is illustrated below:

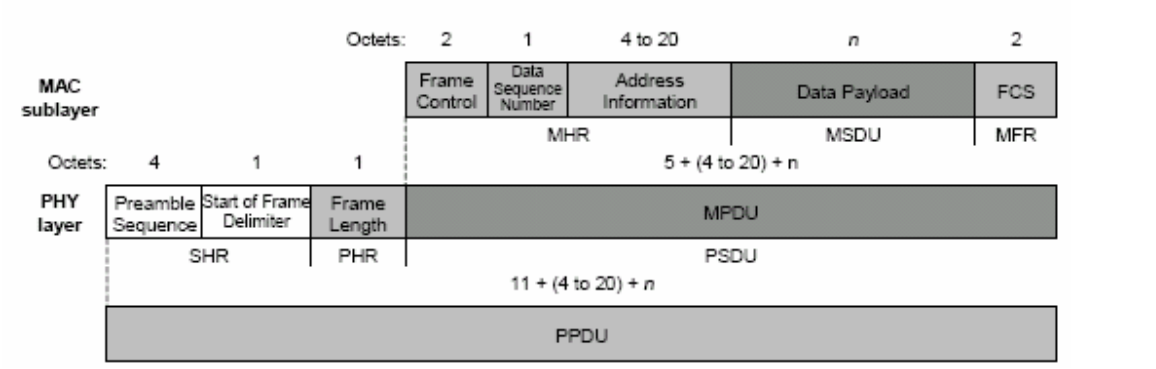

The Physical Protocol Data Unit is the total information sent over the air. As shown in the illustration above the Physical layer adds the following overhead:

Preamble Sequence	4 Octets
Start of Frame Delimiter	1 Octet
Frame Length	1 Octet

The MAC adds the following overhead:

Frame Control	2 Octets
Data Sequence Number	1 Octet
Address Information	4 – 20 Octets
Frame Check Sequence	2 Octets

In summary the total overhead for a single packet is therefore 15 -31 octets (120 bits); depending upon the addressing scheme used (short or 64 bit addresses). Please note that these numbers do not include any security overhead.

Kinney Consulting LLC
pat.kinney@ieee.org

Super Frame Structure

The LR-WPAN standard allows the optional use of a superframe structure. The format of the superframe is defined by the coordinator. The superframe is bounded by network beacons, is sent by the coordinator (See Figure 4) and is divided into 16 equally sized slots. The beacon frame is transmitted in the first slot of each superframe. If a coordinator does not wish to use a superframe structure it may turn off the beacon transmissions. The beacons are used to synchronize the attached devices, to identify the PAN, and to describe the structure of the superframes. Any device wishing to communicate during the contention access period (CAP) between two beacons shall compete with other devices using a slotted CSMA-CA mechanism. All transactions shall be completed by the time of the next network beacon.

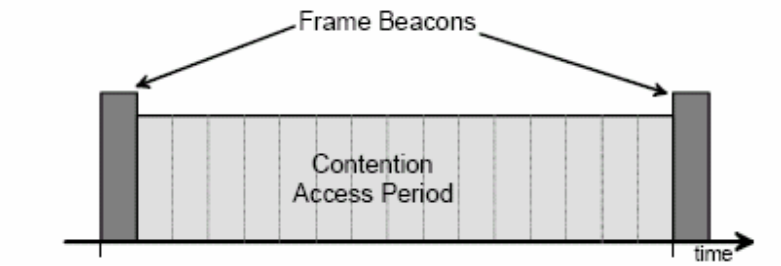

For low latency applications or applications requiring specific data bandwidth, the PAN coordinator may dedicate portions of the active superframe to that application. These portions are called guaranteed time slots (GTSs). The guaranteed time slots comprise the contention free period (CFP), which always appears at the end of the active superframe starting at a slot boundary immediately following the CAP, as shown in Figure 5. The PAN coordinator may allocate up to seven of these GTSs and a GTS may occupy more than one slot period. However, a sufficient portion of the CAP shall remain for contention based access of other networked devices or new devices wishing to join the network. All contention based transactions shall be complete before the CFP begins. Also each device transmitting in a GTS shall ensure that its transaction is complete before the time of the next GTS or the end of the CFP.

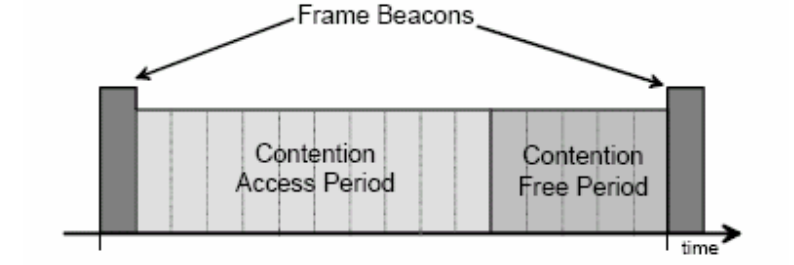

Kinney Consulting LLC
pat.kinney@ieee.org

MAC Data Service Diagrams

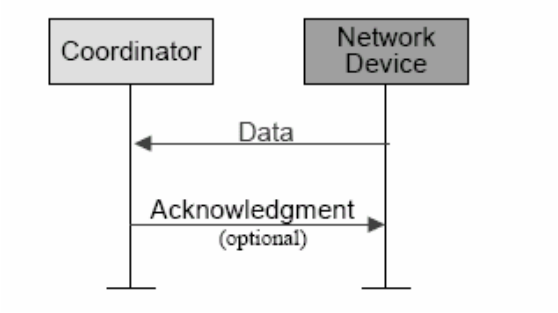

Non-beacon network communication

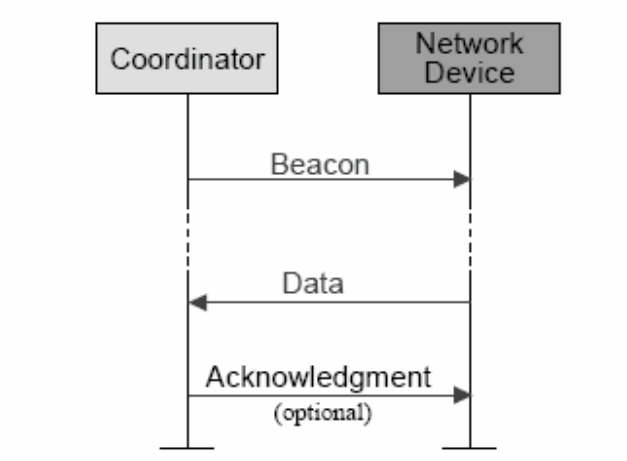

Beacon network communication

MAC Primitives

MAC Data Service

- MCPS-DATA – exchange data packets between MAC and PHY
- MCPS-PURGE – purge an MSDU from the transaction queue

MAC Management Service

- MLME-ASSOCIATE/DISASSOCIATE – network association
- MLME-SYNC / SYNC-LOSS - device synchronization
- MLME-SCAN - scan radio channels
- MLME- COMM-STATUS – communication status
- MLME-GET / -SET– retrieve/set MAC PIB parameters
- MLME-START / BEACON-NOTIFY – beacon management
- MLME-POLL - beaconless synchronization
- MLME-GTS - GTS management
- MLME-RESET – request for MLME to perform reset
- MLME-ORPHAN - orphan device management
- MLME-RX-ENABLE - enabling/disabling of radio system

Kinney Consulting LLC
pat.kinney@ieee.org

Security

When security of MAC layer frames is desired, ZigBee uses MAC layer security to secure MAC command, beacon, and acknowledgement frames. ZigBee may secure messages transmitted over a single hop using secured MAC data frames, but for multi-hop messaging ZigBee relies upon upper layers (such as the NWK layer) for security. The MAC layer uses the Advanced Encryption Standard (AES) [10] as its core cryptographic algorithm and describes a variety of security suites that use the AES algorithm. These suites can protect the confidentiality, integrity, and authenticity of MAC frames. The MAC layer does the security processing, but the upper layers, which set up the keys and determine the security levels to use, control this processing. 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 designated for the key being used. Each key is associated with a single security suite and the MAC frame header has a bit that specifies whether security for a frame is enabled or disabled.

When transmitting a frame, if integrity is required, the MAC header and payload data are used in calculations to create a Message Integrity Code (MIC) consisting of 4, 8, or 16 octets. The MIC is right appended to the MAC payload. If confidentiality is required, the MAC frame payload is also left appended with frame and sequence counts (data used to form a nonce). The nonce is used when encrypting the payload and also ensures freshness to prevent replay attacks. Upon receipt of a frame, if a MIC is present, it is verified and if the payload is encrypted, it is decrypted. Sending devices will increase the frame count with every message sent and receiving devices will keep track of the last received count from each sending device. If a message with an old count is detected, it is flagged with a security error. The MAC layer security suites are based on three modes of operation. Encryption at the MAC layer is done using AES in Counter (CTR) mode and integrity is done using AES in Cipher Block Chaining (CBC-MAC) mode [16]. A combination of encryption and integrity is done using a mixture of CTR and CBC- MAC modes called the CCM mode.

The NWK layer also makes use of the Advanced Encryption Standard (AES). However, unlike the MAC layer, the security suites are all based on the CCM* mode of operation. The CCM* mode of operation is a minor modification of the CCM mode used by the MAC layer. It includes all of the capabilities of CCM and additionally offers encryption-only and integrity-only capabilities. These extra capabilities simplify the NWK layer security by eliminating the need for CTR and

Kinney Consulting LLC
pat.kinney@ieee.org

CBC-MAC modes. Also, the use of CCM* in all security suites allows a single key to be used for different suites. Since a key is not strictly bound to a single security suite, an application has the flexibility to specify the actual security suite to apply to each NWK frame, not just whether security is enabled or disabled

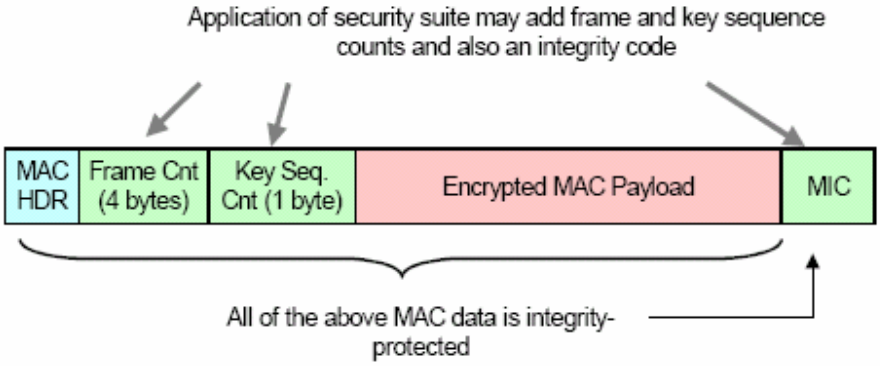

When the NWK layer transmits (receives) a frame using a particular security suite it uses the Security Services Provider (SSP) to process the frame. The SSP looks at the destination (source) of the frame, retrieves the key associated with that destination (source), and then applies the security suite to the frame. The SSP provides the NWK layer with a primitive to apply security to outgoing frames and a primitive to verify and remove security from incoming frames. The NWK layer is responsible for the security processing, but the upper layers control the processing by setting up the keys and determining which CCM* security suite to use for each frame.

Similar to the MAC layer frame format, a frame sequence count and MIC may be added to secure a NWK frame.

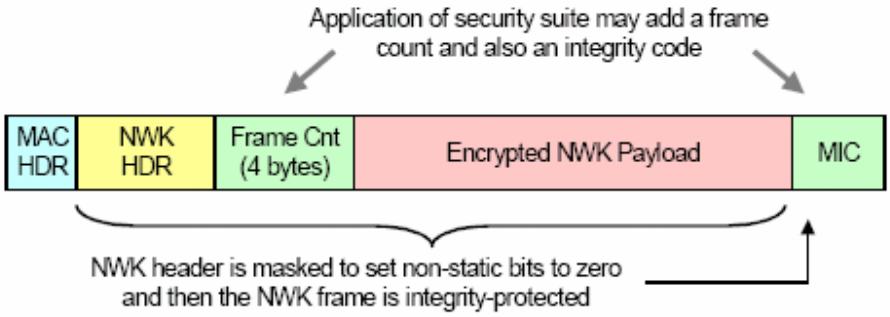

Kinney Consulting LLC
pat.kinney@ieee.org

ZigBee Network Model

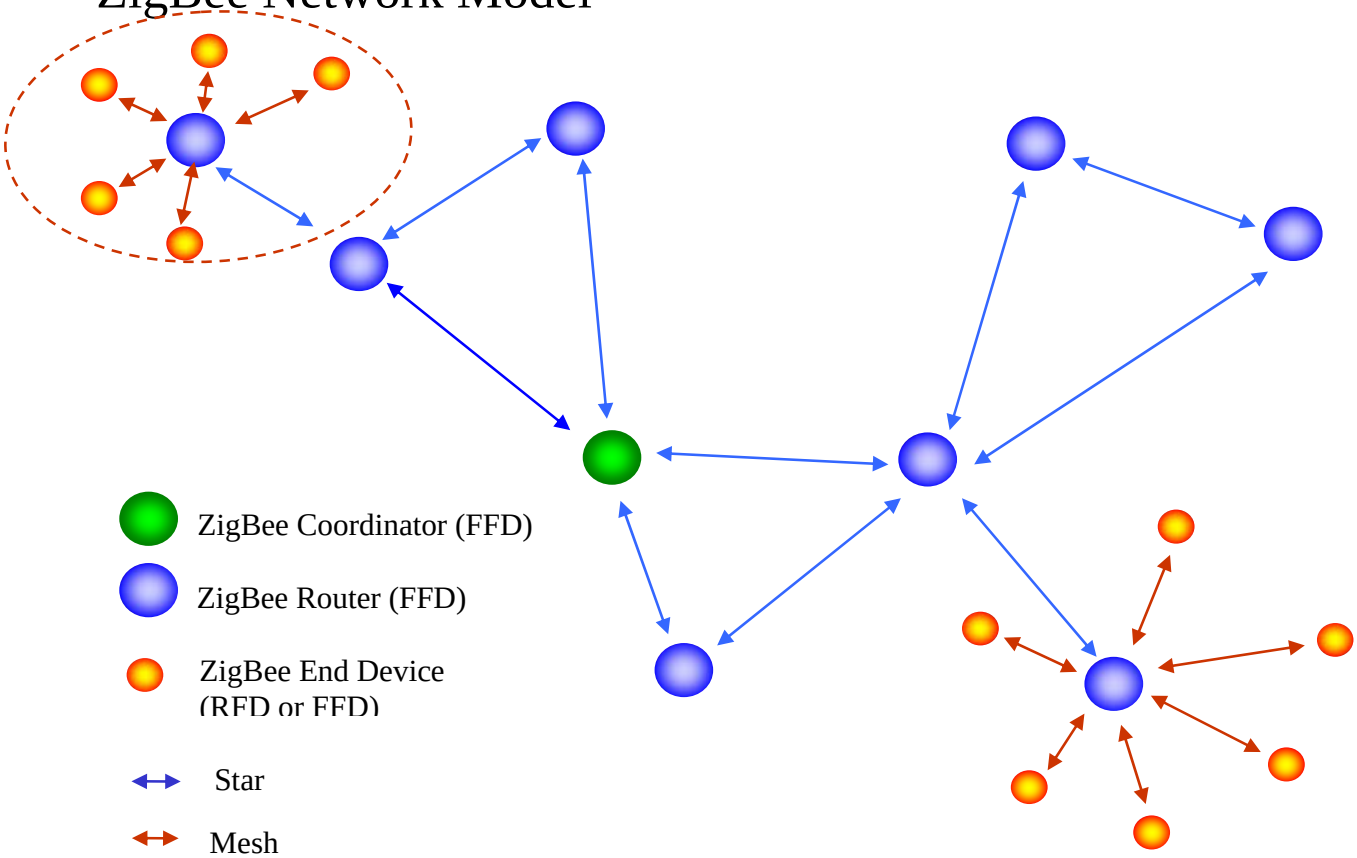

The ZigBee Network Coordinator

- Sets up a network
- Transmits network beacons
- Manages network nodes
- Stores network node information
- Routes messages between paired nodes
- Typically operates in the receive state

The ZigBee Network Node

- Designed for battery powered or high energy savings
- Searches for available networks
- Transfers data from its application as necessary
- Determines whether data is pending
- Requests data from the network coordinator
- Can sleep for extended periods

Kinney Consulting LLC
pat.kinney@ieee.org

ZigBee Stack

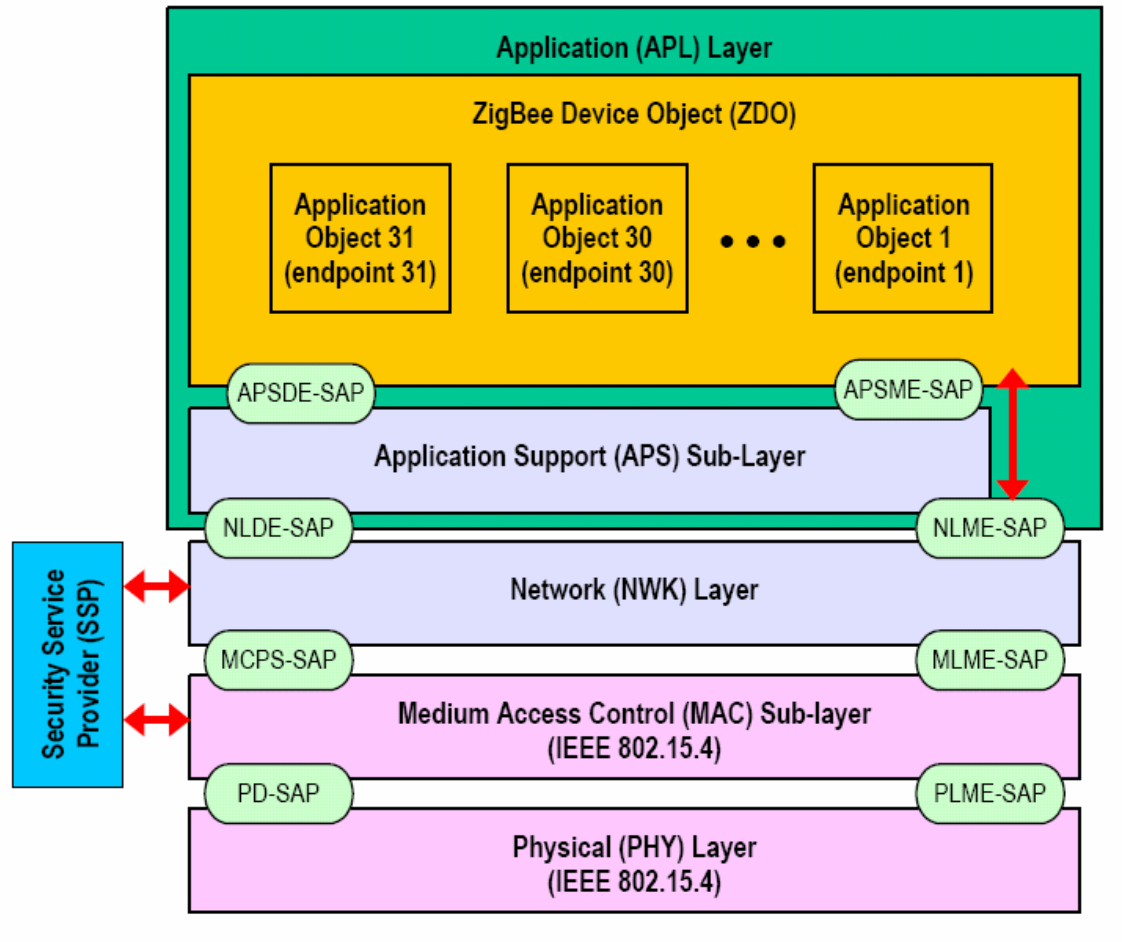

ZigBee Stack System Requirements

- 8-bit μC, e.g., 80c51
- Full protocol stack <32k
- Simple node only stack ~6k
- Coordinators require extra RAM
  - node device database
  - transaction table
  - pairing table

Kinney Consulting LLC
pat.kinney@ieee.org

Network Layer

The responsibilities of the ZigBee NWK layer include:

- **Starting a network:** The ability to successfully establish a new network.
- **Joining and leaving a network**: The ability to gain membership (join) or relinquish membership (leave) a network.
- **Configuring a new device:** The ability to sufficiently configure the stack for operation as required.
- **Addressing:** The ability of a ZigBee coordinator to assign addresses to devices joining the network.
- **Synchronization within a network:** The ability for a device to achieve synchronization with another device either through tracking beacons or by polling.
- **Security:** applying security to outgoing frames and removing security to terminating frames
- **Routing:** routing frames to their intended destinations.

Network Routing Overview

Perhaps the most straightforward way to think of the ZigBee routing algorithm is as a hierarchical routing strategy with table-driven optimizations applied where possible.

– NWK uses an algorithm that allows stack implementers and application developers to balance unit cost, battery drain, and complexity in producing ZigBee solutions to meet the specific cost-performance profile of their application.

– Started with the well-studied public-domain algorithm AODV and Motorola's Cluster-Tree algorithm and folding in ideas from Ember Corporation's GRAd.

Network Summary

The network layer builds upon the IEEE 802.15.4 MAC's features to allow extensibility of coverage. Additional clusters can be added; networks can be consolidated or split up.

Kinney Consulting LLC
pat.kinney@ieee.org

Application layer

The ZigBee application layer consists of the APS sub-layer, the ZDO and the manufacturer-defined application objects. The responsibilities of the APS sub-layer include maintaining tables for binding, which is the ability to match two devices together based on their services and their needs, and forwarding messages between bound devices. Another responsibility of the APS sub-layer is discovery, which is the ability to determine which other devices are operating in the personal operating space of a device. The responsibilities of the ZDO include defining the role of the device within the network (e.g., ZigBee coordinator or end device), initiating and/or responding to binding requests and establishing a secure relationship between network devices. The manufacturer-defined application objects implement the actual applications according to the ZigBee-defined application descriptions

ZigBee Device Object

- Defines the role of the device within the network (e.g., ZigBee coordinator or end device)
- Initiates and/or responds to binding requests
- Establishes a secure relationship between network devices selecting one of ZigBee's security methods such as public key, symmetric key, etc.

Application Support Layer

This layer provides the following services:

- **Discovery:** The ability to determine which other devices are operating in the personal operating space of a device.
- **Binding:** The ability to match two or more devices together based on their services and their needs and forwarding messages between bound devices

Kinney Consulting LLC
pat.kinney@ieee.org

The Inevitable Question is whether ZigBee and Bluetooth are competitors or complements?

Bluetooth seems best suited for:

- Synchronization of cell phone to PDA
- Hands-free audio
- PDA to printer

While ZigBee is better suited for:

- Controls
- Sensors
- Lots of devices
- Low duty cycle
- Small data packets
- Long battery life is critical

Air Interface comparison:

ZigBee	Bluetooth
DSSS	FHSS
11 chips/ symbol	1600 hops / second
62.5 K symbols/s	1 M Symbol / second
4 Bits/ symbol	1 bit/symbol
Peak Information Rate ~128 Kbit/second	Peak Information Rate ~108-723 kbit/second

Battery Drain comparison to Bluetooth

Packet length can affect battery drain. Typically the shorter the packet the quicker the device can go to sleep. Bluetooth is a slotted protocol. Communication can occur in either: 625 µS, 1875 µS, or 3125 µS slots. The following graph showing effective data rate was based upon the transmissions speeds stated in Bluetooth v1.1 and IEEE 802.15.4 draft 18, using the 250 kb/s rate. The general trend is that at larger packet sizes the effective data rate approaches the raw data rate.
The peaks for the Bluetooth rate are a result of the three slot sizes, when a packet becomes too big for one slot it must increment to the next slot even

Kinney Consulting LLC
pat.kinney@ieee.org

though it doesn't fill the whole slot allocation.
IEEE 802.15.4 was designed for small packets so it is no surprise it is more efficient at those small packets resulting in a higher effective rate despite its lower raw data rate.
From this graph we can see that for packets less than 75 bytes ZigBee has a higher effective data rate than Bluetooth. Having a lower rate for small packets means that BT needs longer transmit and receive times and therefore current drain is higher for small data packets.
Although these numbers do not represent retransmissions or multiple devices requesting the bandwidth; the author believes that the same traits will be exhibited in these other cases.

Effective Data Rate

(based upon theoretical values with no retransmissions)

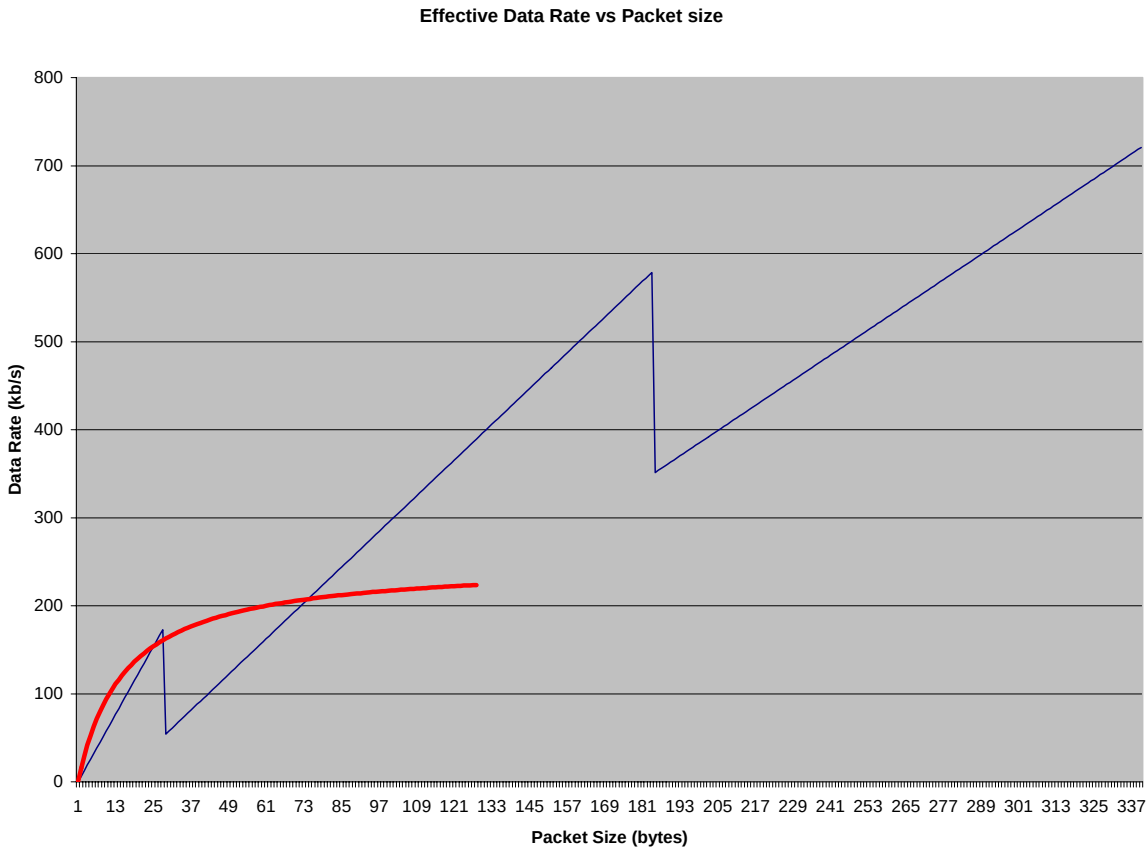

Kinney Consulting LLC
pat.kinney@ieee.org

Timing Considerations

ZigBee

- New slave enumeration = 30ms typically
- Sleeping slave changing to active = 15ms typically
- Active slave channel access time = 15ms typically

Bluetooth

- New slave enumeration = >3s, typically 20s
- Sleeping slave changing to active = 3s typically
- Active slave channel access time = 2ms typically

Conclusion:

ZigBee devices can quickly attach, exchange information, detach, and then go to deep sleep to achieve a very long battery life. Bluetooth devices require about ~100X the energy for this operation.

Power Considerations

ZigBee

– 2+ years from 'normal' batteries
– Designed to optimize slave power requirements

Bluetooth

– Power model as a mobile phone (regular daily charging)
– Designed to maximize ad-hoc functionality

Since IEEE 802.15.4 uses a CSMA-CA protocol the end nodes only talk when they have data to send with the following benefits:

– No waiting for polling (however they must wait for a clear channel which shouldn't be a problem in low duty cycle networks such as with sensor and control devices)
– Current drain is substantially reduced over a polling protocol that must poll to maintain latencies even though the majority of the time the device needed be polled
– IEEE 802.15.4 protocol was designed to yield 6 months to 2 yrs on alkaline cell

Kinney Consulting LLC
pat.kinney@ieee.org

ZigBee Battery Drain

In this section we'll look at different aspects of a networked device's battery drain.

A typical scenario for sensors and control devices is to remain "connected" to the network. We use connected to mean that the device periodically listens for incoming packets. In this manner the device's behavior may be altered or at least checked to verify correctness.

Scenario 1: ZigBee Battery Drain, network connection

Let's review a couple of aspects for ZigBee devices:

Goal: Two year battery life

Assumptions:

- AAA cell = 1.15 Ahr (Duracell alkaline)
- 2 yrs = 17,532 hrs

Partial result: Average current drain $\leq$ 65 µA (capacity/time)

- Tx/Rx current drain ~ 15 mA and sleep current = 1 µA

Partial result: Maximum duty cycle $\leq$ .43% (Avg. current drain-sleep current)/current drain

- Beacon duration of 3 mS (longer beacons containing more information would drain more current)
- Beacon rate of 1/s (beacon rates can be as slow as .03/s)

Partial result: beacon use in this case requires a .3% duty cycle

Final result: 22.8 hours (0.13%) of transmission time would be allowed for data transmission or reception

Scenario 2: Battery Drain when the unit is not connected to the network

This mode can be used to maximize battery life. The device will only connect to the network when it needs to send data. A disadvantage of this technique is that the device cannot be sent data, so for the most part it is seldom part of the network.

Assumptions:

- Device will connect only when necessary to send data
- Acquisition time
  - Bluetooth requires about 20 – 30 seconds (~98% confidence) for an Inquiry (first time) and about 3 seconds for a Page (subsequent times)
  - IEEE 802.15.4 acquisition time is about 30 mS
- Using maximum duty cycle of .43% and 40 byte packet

Kinney Consulting LLC
pat.kinney@ieee.org

Result:
– ~ 45,140 data transmissions for Bluetooth
– ~ 4,269,670 data transmissions for ZigBee

Battery drain conclusion: ZigBee has an inherent advantage for these modes of operation due to its short attach time and/or its ability to remain in the sleep mode for long periods.

Comparison Summary

- ZigBee and Bluetooth are two solutions for two different application areas.
  - The differences are from their approach to their desired application. Bluetooth has addressed a voice application by embodying a fast frequency hopping system with a master slave protocol. ZigBee has addressed sensors, controls, and other short message applications by embodying a direct sequence system with a star or peer to peer protocols.
  - Minor changes to Bluetooth or ZigBee won't change their inherent behavior or characteristics. The different behaviors come from architectural differences.

Kinney Consulting LLC
pat.kinney@ieee.org