Untangling the Mesh

The Ins and Outs of Mesh Networking Technologies

White Paper

Joel K. Young
Senior Vice President & CTO, Digi International Inc.

Abstract

This paper is based on a presentation titled "What a Mesh!" given by Joel Young at Embedded Systems Conference Silicon Valley 2008. The paper discusses discusses mesh networks, including wireless network basics, the criteria for evaluating different wireless mesh networking technologies, an overview of mesh related technologies, and an evaluation of the technologies.

Introduction

Over the past few years, mesh networks have become ever more popular, following the trend to create more wireless things. As with other technology trends, there have been a plethora of different mesh networking technologies and architectures. This paper is intended to bring order and understanding to the diversity associated with mesh networking. First, the network basics will be discussed – focusing on the specifics of wireless. Second, criteria for evaluating different wireless mesh networking technologies will be presented. This will be followed by an overview of five different mesh related technologies – including key characteristics, the network architecture, strengths and limitations. Finally, all of this information will be aggregated to create an evaluation of these different approaches, including when they should be applied.

Networking Basics

To begin we start with the basics of networking and different topologies. It is important to note that network topologies describe the interaction and interconnection of the participants. This means how they communicate with each other and how they establish paths between each other. Networks topologies are not always what they seem to be. In the wired days, they generally followed the path of the wires – very simple. If devices were wired in a ring, then the network topology is a ring. The journey to wireless complicates everything because we all share the same air space so the path and access method is not always obvious. For example, is a Wi-Fi® access point a star topology or a bus?

Before we go further in looking at these questions, it is important to first understand some common terminology which will be used throughout this paper.

- **DSSS – Direct Sequence Spread Spectrum.** This is a method of encoding a signal which distributes information over a wide path of spectrum using a pseudo random code. Because of the wide spreading, the signal appears to be noise for those without the spreading code.
- **FHSS – Frequency Hopping Spread Spectrum.** Similar to DSSS, the big difference is that is uses a more constrained spreading algorithm and changes channels as a function of time, theoretically making the transmission more immune to interference.
- **TSMP – Time Synchronized Mesh Protocol.** This is a mesh protocol that uses time slots to allocate spectrum for communication between two nodes. Because time slots differ over pairs, interference is minimized because access to the channel is controlled by timeslot.
- **AODV – Ad-hoc On-demand Distance Vector routing algorithm.** AODV is a pure on-demand route acquisition algorithm: nodes that do not lie on active paths do not maintain any routing information nor participate in any periodic routing table exchanges. Further, a node does not have to discover and maintain a route to another node until the two need to communicate.
- **Cluster-Tree – Region based mesh network routing algorithm.** In this algorithm, routes are formed and maintained between clusters of nodes. Route discovery is completed and maintained between the clusters – providing access to the children of each cluster.
- **ISM – Industrial Scientific and Medical band.** This describes the frequency bands that can be used license-free. Generally we refer to the 2.4 GHz band, but it also covers spectrum in 900 MHz, 5.8 GHz, 433 MHz in North America. 2.4 GHz is used worldwide.
- **IPv6 – Internet Protocol Version 6.** This is the latest version of the popular IP or Internet Protocol. With Version 6, the IP address structure, routing and class of service changes. It is part of the TCP/IP suite of protocols sponsored by the IETF.
- **PAN ID – Personal Area Network Identifier.** This is the term for the network name assigned to particular personal area network.
- **CSMA – Carrier Sense Multiple Access.** This protocol defines the channel access technique deployed by Ethernet, Wi-Fi and bus oriented networks. It provides a method for detecting collisions and retransmitting as a method to acquire a communications channel.
- **TDMA – Time Division Multiple Access.** The protocol defines the channel access technique used by TSMP and GSM networks in which a communications channel is divided into time slots. Each node is allocated a specific time slot for communication.

Network Types. Now that we have a feel for the terminology, we first look at different network topologies commonly used. The figure below shows three such topologies: Star, Bus and Ring.

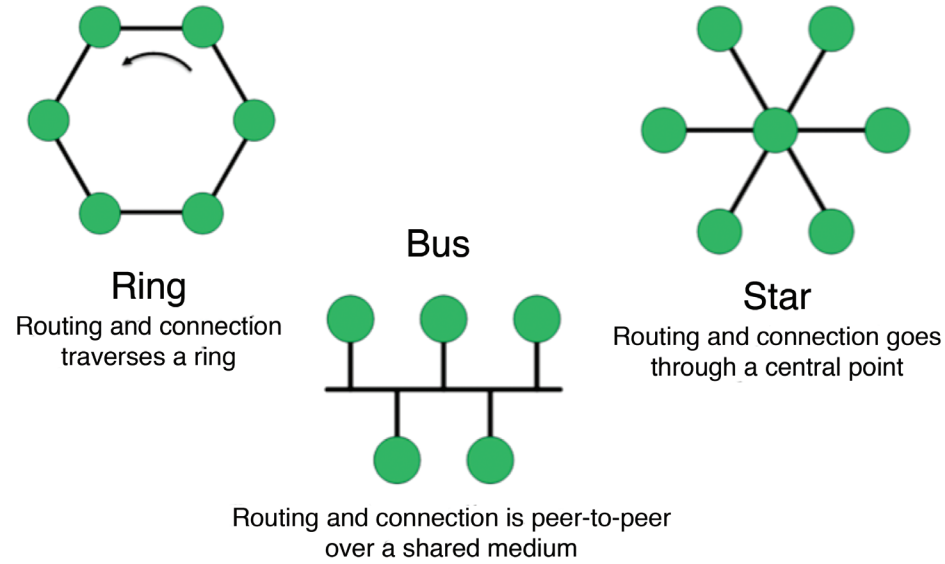

In the Ring, nodes are connected from one to the next. Communications messages are forwarded around the ring in either a clockwise and/or counter-clockwise fashion. As a message is forwarded, the node checks to see if the message is meant for itself, if so, it keeps the message, if not, it forwards it on. It is most common in cabled networks (wire or fiber), but could conceivably be used in a wireless fashion as well – but is not practical unless being used over long distances.

In the Bus, all nodes share a common communications medium and contend for using it. Typically this means some kind of CSMA type approach. Since a common medium is used, collisions and retransmissions increase with traffic loading. In the wired case, these types of systems are referred to as hubs – which are generally no longer used. In the wireless case, it is more complicated because open space is often a shared medium, so even if routing is handled in a star, ring or other topology, open space often appears as a bus. More on this later as it is one of the fundamental hurdles in wireless networking.

In the Star, nodes are connected through a master, central node. This central node is responsible for looking at each message and forwarding it out on the proper communications link. While various star architectures have been used over time, the most commonly known in the wired space is the Ethernet Switch. In the wireless space, the Wi-Fi access point is also a familiar example, since all messages are routed through the access point; however, even though messages are routed through the access point, open space is accessed via CSMA – a bus type protocol.

A mesh network employs some level of more complete interconnection among nodes. This means that paths are not defined by a specific architectural pattern, but rather by the connections themselves. In the full mesh topology, each node (workstation or other device) is connected directly to each of the others. In the partial mesh topology, some nodes are connected to all the others, but some of the nodes are connected only to those other nodes with which they exchange the most data. The figure below illustrates a full mesh, where each of the five nodes is connected to all the others.

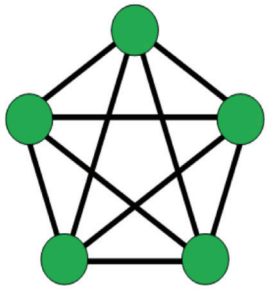

Another important thing to note about a mesh is that some or all nodes may be routers and some or all nodes may be end points. Typically, full interconnection is not achieved, unless the network is very small. Full interconnection gets very complex very quickly. In addition, wired mesh networks tend not be practical due to the complexities of connecting all the wires. The following figures illustrate three different instantiations of mesh networks. The green nodes are end devices, the yellow are routers (which may also be end devices) and the purple is the network coordinator – responsible for allowing joining and departing from the mesh (more on this later). Note that one instantiation of a mesh can be a star – a mesh with one router and the rest end points. The Cluster-Tree network is a combination of near full connectivity among routers and end points hanging off individual routers. The Peer-to-Peer mesh generally gives equal rights to all nodes, including routing and end point functionality.

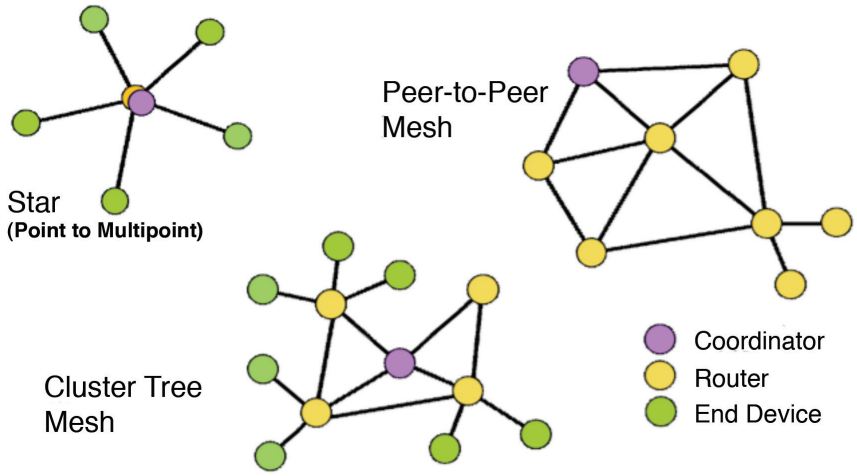

While we have discussed the fact that mesh networks aren't really practical for wired networks, it is also important to look at the other differences between wired and wireless domains. So what else makes wireless different? On the positive side, wireless makes it possible to have more connections since it is not practical or cost effective to create a full mesh with wires. However, on the negative side, wires are predictable, reliable and well understood. Wireless forces the sharing of an already noisy, uncontrollable medium called open space. Hence, while wireless gives us more flexibility, the uncertainty of wireless drives the need for more connection paths…and more complexity.

Then we look at wireless, particularly mesh networks, there are a number of hard problems that need to be solved.

Accessing the medium. Since we all share open space, listening is more important than talking. If everyone talks at once, listening is difficult. So radios must be good listeners if they are going to have a chance to get a word in edge wise, so to speak.

Discovering routes. Determining paths in a wireless mesh network is difficult because the environment is dynamic. In this case, there are two choices: Planning the trip in advance, or taking it one step at a time. Oftentimes doing both is best – this usually involves retracing ones steps and repeating well traveled routes.

Adapting to a changing environment. In a wireless world, paths to nodes can disappear and re-appear. This is due to changing signal conditions or traffic conditions.

Sleeping and Waking. Once we go wireless, the next step is often to eliminate the traditional power cable. This means batteries and the need for effective power management. The most common way of handling power management is putting the nodes to sleep when they are not being used. This sounds well and good until it is time to wake up.

How to Compare

Given that we now understand the basics of networks, in particular mesh networks, the next question that should be addressed is how to compare them. For this, we should consider security, reliability, power management, scalability, data movement, and cost.

Security. This is as much about the perception of threat as actual threat. Nonetheless, security is easily evaluated by the traditional factors that are well understood in the industry. The first is encryption – protecting the information itself. Modern encryption wants

at least AES128 as an algorithm (128 bit key). The next is authentication, which is validating that the users (or nodes) are who they say they are. This is typically handled by key exchange or authenticated certificate. Last is authorization, which should be thought of as granting permission associated with having the right key or certificate. Beyond these, there are other factors which are associated with the ease of distributing and configuring the authorization and authentication mechanisms.

Reliability. The best way to think of reliability is the likelihood a message will arrive at its intended destination on time. If the message always are arrives at the destination when expected, the network is very reliable. Secondarily, we want the message to arrive at the destination, even if it is a bit late. The components for evaluating reliability for wireless mesh networks have to do with the following:

- **Frequency agility –** This is detecting and adapting the network around potential interference.
- **Message loss potential –** This pertains to whether messages get lost in the shuffle. With all the re-routing and different paths, the network must be very careful to ensure messages don't get lost and that duplicate messages following different routes get discarded.
- **Adaptability –** This is best described as the network's ability for changing the routing to accommodate for nodes disappearing – while still preventing lost messages. This is most effective if done quickly.
- **Single points of failure –** Simply put, are there any single points of failure, what is the risk of them failing, and how is recovery handled?

Power Management. The most frequent question asked by customers when discussing wireless sensor networks is how long will my batteries last? As soon as the cord is cut, everyone wants to still keep maintenance low. Viewed in the context of the network architecture, power management is analyzed in terms of end nodes, router nodes and network coordinators. It is most important to have low powered, power efficient end nodes because they are most likely to be far from traditional wired power sources. The routers are second. Battery powered routers, or routers that sleep extend the flexibility of the architecture. Finally, the coordinator is usually always powered. Now, in the context of nodes that can sleep, we then look at their average power consumption. This is best assessed by looking at the combination of how they wake up, how frequently they wake up, total transmitting time and total listening time. Since the most power is consumed when radios transmit, it is important to keep transmit time to a minimum.

Scalability. How big can the network get before it fails, at least on a practical level? All the networks have large physical limits in the 10s of thousands, but the practical design of the network is always much smaller. This is because scalability is related both to reliability mechanisms and nature of the application. If a network never has any problems which cause rerouting, then network routing tables will never change, meaning cached routes will always work and there will be few retransmissions or reroutes because of failures. The end result is a very stable network that can be very large. Scalability is also dependent on the type and volume of data. Data flow can be placed in three categories: Dribble Data, Bursty Data and Streaming Data. Dribble data is periodic, infrequent and slow, while streaming data is constant, etc. A network can be very large if the traffic is dribble data because the flow follows consistent patterns, with plenty of bandwidth. Sleeping networks do well with dribble data, but scale poorly with streaming data.

Data Movement. Now we look in more detail of data flow – not for network scalability purposes, but for raw capacity. Here there is a classic trade-off in needs: Does the application require lots of data with low latency or does it require dribble data with long, non-deterministic latency? As such, in evaluating networks for data movement, a combination of the following five variables needs to be considered, namely data rate, latency, packet size, fragmentation, and range.

Cost. Cost is measured by the individual unit cost per node as well as the cost to maintain the network. Maintenance is often difficult to quantify and deployment cost is often forgotten. It is easiest to quantify those variables that are most perceptible, namely the actual purchase cost of a transceiver system per node. Things become a bit more complicated in the case of battery powered sleeping nodes. For example, assume all end points need to be sleeping and a point to multipoint system is not practical due to range. Then a network that does not have sleeping routers will need to deploy powered routers in addition to the end points as compared to a network that has sleeping routers. Hence, even if all radios share the same unit cost, more radios are needed in the powered router system. However, if power is available, then it becomes a non issue. So, the cheapest radio may not be the best for the application. The other point is that the cost of the radio tends to be evaluated relative to the cost of the device to which it is connected.

Network Architectures

Point-to-Multipoint

Key Characteristics. This is also known as a simple star and it is not really a mesh network, but it is often confused with one. These networks tend to use the modern air Interfaces of either Frequency Hopping Spread Spectrum or Direct Sequence Spread Spectrum (802.15.4). They need to be statically configured for PAN ID, routes, and security. The keys to note are that all nodes can see all other nodes and they need to be told who to talk to. Security tends to be pair-wise for both the encryption and key. End points may go to sleep or stay awake, but the central router is always awake.

Network Architecture. The figure below illustrates a typical topology. All nodes are on the same channel (or hop to the same channel). Bandwidth / throughput limited by simultaneous data at concentration point. Collisions happen with lots of traffic or lots of nodes.

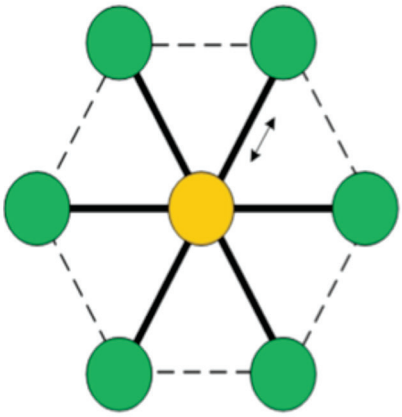

Green End points (green) see other end points but are told only to talk to central point (orange)

Strengths. The beauty of the basic non-mesh point-to-multipoint network is simplicity. Communication, unless traffic is very heavy, is relatively deterministic since there are no hops and minimal or managed collisions. It also allows for maximum throughput because there is no added routing and no added route discovery. Finally it is easy to understand and easy to manage. Because of the simplicity, it also tends to drive the lowest cost for its specific size and function.

Limitations. Unfortunately, the simplicity described in the strengths also drives a number of limitations. The networks will tend to be small. Large networks only work if polled from central point. This requires very specific message management. There are also single points of failure and no ways to route around changing conditions. The network follows the belief that if it worked the first time, it will work forever so you must be sure of good RF conditions.

ZigBee® 2007

Key Characteristics. ZigBee is built on top of 802.15.4 using DSSS in 2.4 GHz. End points sleep, Routers don't sleep and a coordinator is needed to start the network and to allow points to join the network. The ZigBee standard has had three different versions: 2004, 2006 and 2007. ZigBee 2004 is no longer used and ZigBee 2006 had significant limitations. ZigBee 2007 includes key features for frequency agility, message fragmentation, and enhanced security associated with key management. The routing of messages follows the previously described Cluster-Tree methodology where routes to all points are maintained at each cluster. This allows a very short routing time, but requires lots of routes. Discovery of routes uses the AODV algorithm where paths are explored between clusters.

Network Architecture. The network consists of three specific types of points. A ZigBee Coordinator (ZC) is required for each network to initiate network formation. The coordinator may act as a router once the network is formed. The ZigBee Router (ZR) is actually an optional network component, although a network without routers becomes a point-to-multipoint network described earlier. The router participates in multi-hop routing of messages. Finally, the ZigBee End Device (ZED) does not allow association and does not participate in routing. As such it is often referred to as a child because it doesn't really have any responsibilities. The following figure illustrates a network.

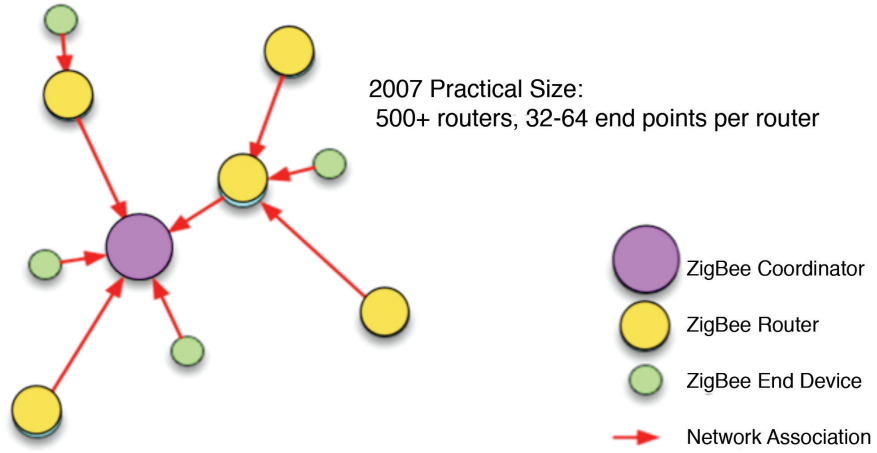

Strengths. End devices are very low power because they are subservient to parental routers. Cluster-Tree routing provides quick knowledge of routes and thereby efficient routing. With ZigBee 2007, frequency agility switches from problem channels automatically in a sort of on demand frequency hopping. Long messages are allowed with message fragmentation support and security is flexible with support of separated keys. Finally, the network can scale to be very large.

Limitations. The biggest limitation tends to be that routers must be powered - they can never go to sleep. In addition, the benefit of Cluster-Tree routing also means that network changes require a lot of route discovery traffic. Heavy traffic volume means lots of collisions and potential message loss. Finally, a coordinator is needed to start and manage the network - so if the coordinator goes down, then no one can join and the network can't start.

Wireless HART

Key Characteristics. Wireless HART uses the Time Synchronized Mesh Protocol (TSMP) created by Dust Networks. Unlike other networks, the time based system uses TDMA (Time Slots) for an access method. The network is optimized for low power and all nodes can be sleeping routers and every node is a router. A Gateway is required to keep the network synchronized due to the critical time synchronization of sleeping and waking functions. Like ZigBee, it is built on top of 802.15.4 DSSS, but it adds a more deliberate frequency hopping algorithm. Security includes encryption and authentication.

Network Architecture. The figure below illustrates a typical network topology. Note that all the nodes are routers. The illustrated routes change dynamically based on visibility within specific time slots as it hops through the different DSSS channels. The relationship between any two nodes is negotiated to be in a specific time slot, thereby minimizing the probability of any collisions. When sleeping, nodes awaken during their time slot and listen to see if there is any traffic. Clocks are kept synchronized by the gateway.

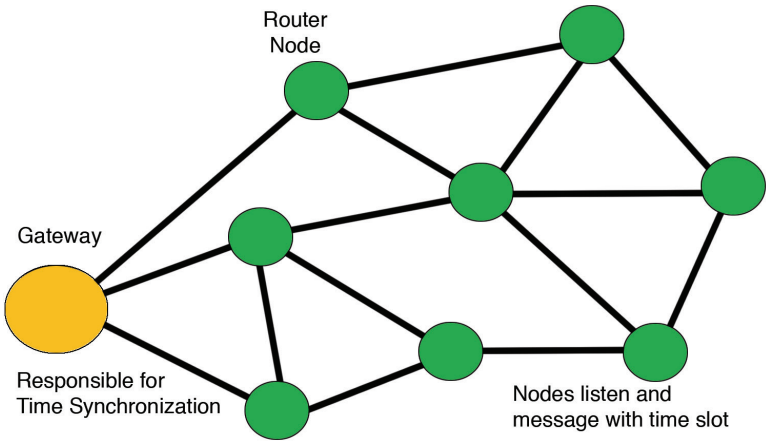

Strengths. Every node is a router at very low power consumption. Most of the time is spent listening. Since transmissions occur only within the allocated time slot, retransmissions are minimized. Communications are very reliable with every message acknowledged. Networks are able to scale to moderate level or around 1000 nodes. Frequency hopping minimizes the probability of interference. Security includes encryption and appropriate authentication.

Limitations. Because of the time slot approach, latency is long and non-deterministic. It takes a network a while to form and all of the nodes to negotiate their individual time slots. Because communications is slotted, the available 802.15.4 bandwidth is split up, meaning that throughput is minimized for bursty traffic. A powered gateway (coordinator) is required for network to stay functioning – opening up a single point of failure if the gateway is unavailable for an extended period of time. Finally, the radios are very expensive compared to the other available solutions.

6LoWPAN

Key Characteristics. 6LoWPAN is a distorted acronym for IPv6 over low power wireless personal area networks. Presently it is a proposed standard based on the IETF RCF 4944. It is designed to be used over 802.15.4 chips and radios. Unlike traditional IPv6, 6LoWPAN deals with Packet size incompatibilities in message transport (128 bytes vs. MTU of 1280 bytes in IPv6) and it is designed for a small memory footprint systems. Today it is a point-to-multipoint architecture and it proposed to be augmented with a mesh routing scheme.

Network Architecture. The figure below illustrates an example network topology. Note that for now it is only point-to-multipoint. Unlike the other networks discussed in this paper, the figure shows an end to end IP based link from a host computer to an end device. In this case it is illustrated by a meter. The end device is directly addressable by the host computer on the far end of the network. The interworking function provided in the pictured box provides a transport change and re-packetization at the data-link level.

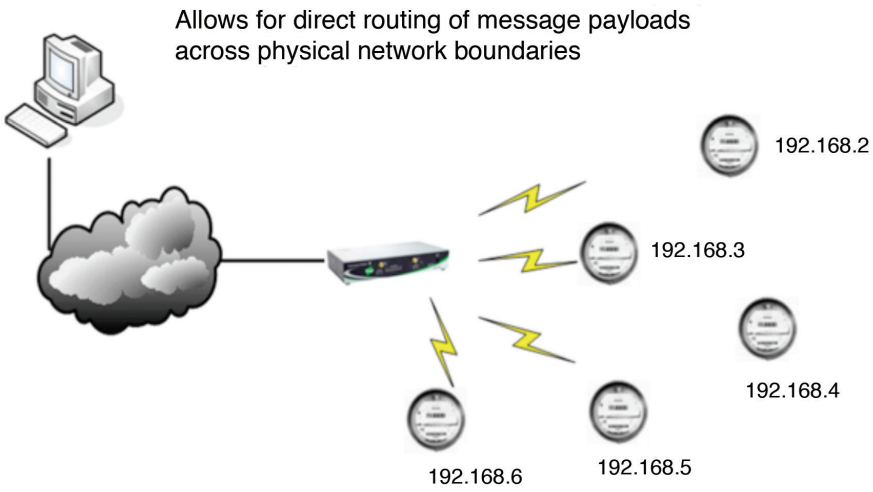

Strengths. The most powerful strength is that that 6LoWPAN is able to take advantage of the existing TCP/IP suite of internet protocols, all of which are well understood due to the proliferation of the internet. Hence it is able to capitalize on existing protocols, existing quality of service and security framework supported by the IETF. Hence, it enables seamless routing of message payloads.

Limitations. This system is still very new and is only a proposed standard. Because it is in officially in the public review stage, it will most likely undergo a number of changes. In fact, the mesh routing working groups are still being formed meaning that wide scale adoption is still a few years away. As such, interoperability is a nice concept that has not been proven yet. Finally, because it is still new, it has not yet been ported to a large group of chipsets.

DigiMesh™

Key Characteristics. Like its sibling Wireless HART, DigiMesh is designed to meet the very low power sensor networks where battery powered routers are required. It is available in multiple frequency bands, 2.4 GHz DSSS and 900 MHz FHSS. It does not rely on a full 802.15.4 implementation as it has some of these functions internal. For both message routing and discovery, it uses a variant of AODV. This means that routing tables built for only for needed destinations, leaving it to be referred to as a peer-to-peer mesh, instead of a cluster-tree. All nodes are viewed as equal participants meaning that they are all routers and they can all sleep. Channel access is a sort of time synchronized CSMA method, enabling bursty traffic, but the benefits of few collisions. It has a full security suite.

Network Architecture. The figure below illustrates a typical ad hoc network topology. Unlike the Cluster-Tree method described in ZigBee, routes are only determined on an as needed basis. This means that routes that are never used never get routing table entries and routes that are used frequently are continuously updated, optimizing there efficiency.

DigiMesh Nodes
No parent-child relationships
All nodes can route
Nodes are interchangeable
All can be low power

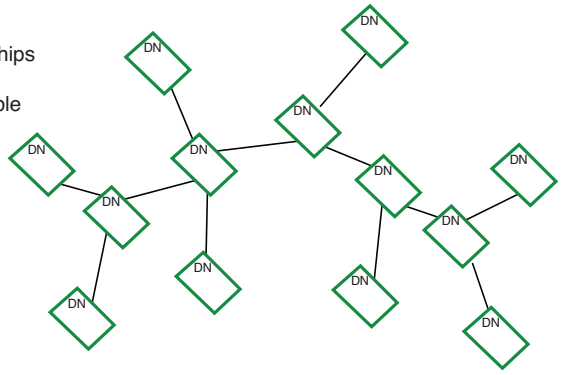

One of the other keys to note about the topology is that there is no coordinator or gateway function. Time synchronization is accomplished through a nomination and election process, enabling the network to operate autonomously.

Routing Methodology. The next set of figures illustrates the process of how routing failures are handled. The first shows the initial network configuration where a route has been established from one point to another. The second illustrates a failure where one of the nodes has been removed for an unknown reason, removing relationships in the center of the route. Finally, the last figure shows how this route is reconstituted using a path that didn't previously exist. The relationships were there, but they had never been used, but were newly discovered using AODV after the failure.

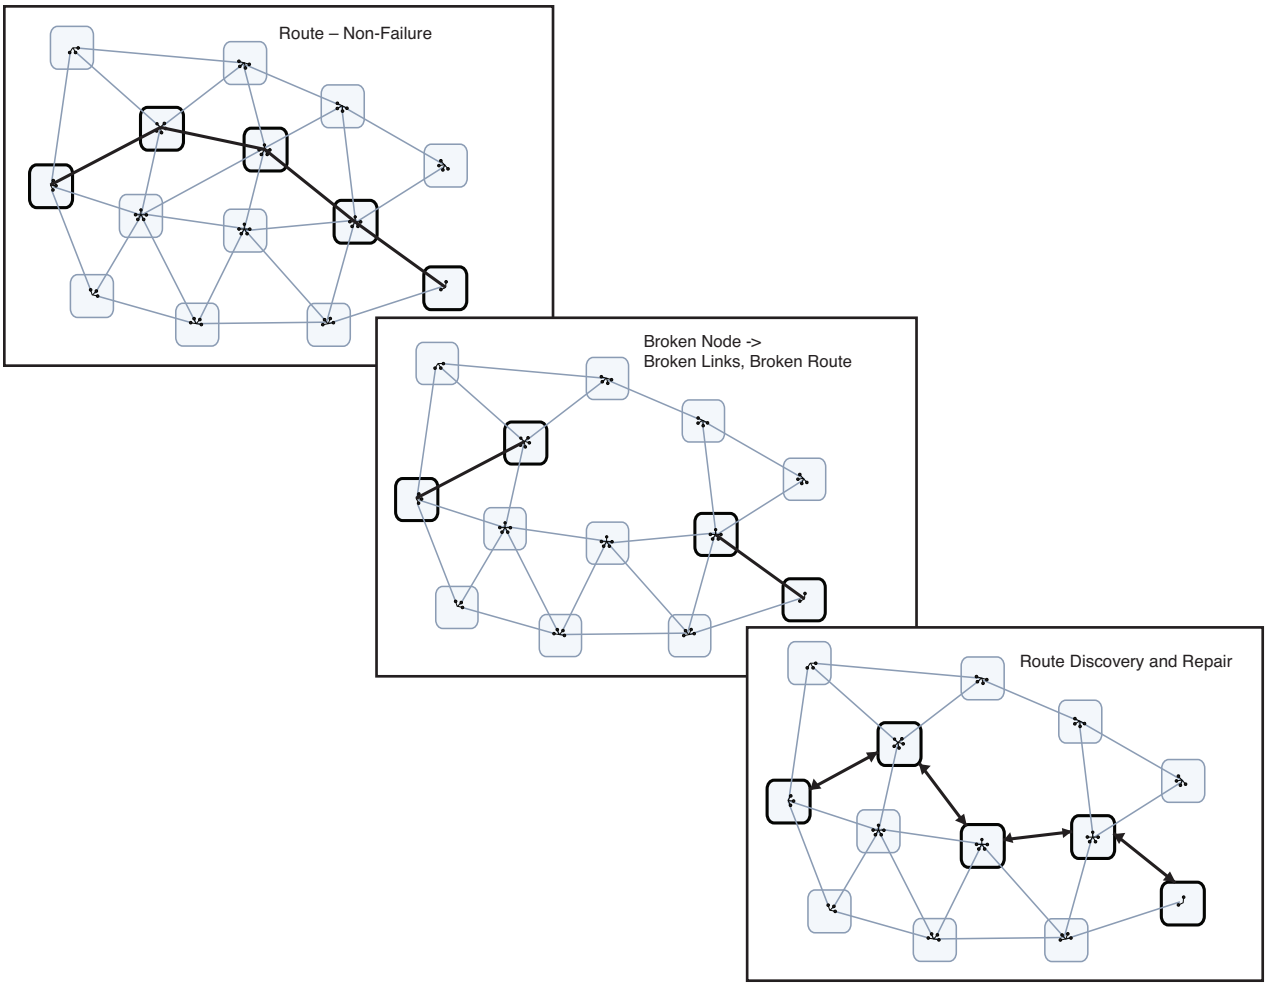

Strengths. Every node is a router at very low power consumption. Further, because every message is acknowledged and routes are determined on an as needed basis, the network is not overwhelmed with unnecessary discovery traffic – very important if the routers are battery powered and sleeping. Efficient route discovery and routing means that the network only learns routes that actually get used (AODV). Frequency agility is supported and security meets the requirements of both encryption and authentication. Reliability is projected at 99.99%. Finally, the system supports larger Payloads with support for message fragmentation.

Limitations. Unfortunately, efficient power management means latency is long and non-deterministic. Even though throughput is not limited by time slots, it is still limited depending on loading and discoveries. The network can scale to a moderate size of around 500+ nodes and can be very large if traffic is light and message flow doesn't change much.

Comparison

Using the criteria defined at the beginning of this document, the following table illustrates the author's best attempt at evaluating the different network approaches. It is important to note they all do very well in security in that they have well defined encryption, authentication and authorization schemes. ZigBee and 6LoWPAN get a slight nod here only in that their key systems should be easier to implement and a bit more flexible.

Category	Point-to-Multipoint	ZigBee	Wireless HART	6LoWPAN	DigiMesh
Security Encryption Authentication Authorization	★★★	★★★★ Need 2007 Pro	★★★	★★★★ Use existing IP security	★★★
Reliability Freq. Agility Message Loss Adaptability	★★	★★★ Need 2007 Pro	★★★★	★★★	★★★
Power Mgmt Sleeping Routers End Nodes Sleep Strategy	★★★	★★★ Routers: 1 star End node: 4 star	★★★★	? Not really defined	★★★★
Scalability Network Size Traffic Volume	★	★★★★ Need 2007 Pro	★★★	? Not really defined	★★★
Data Movement Data Rate Latency Range	★★★★	★★★	★	? Not really defined	★★
Cost	★★★★	★★★★	★	★★★ Assume current chipsets?	★★★★

With respect to reliability, point-to-multipoint takes the biggest hit because it inherently has a single point of failure. Some schemes may have frequency agility options while others do not. Prior to the 2007 standard, ZigBee has a weakness in the frequency agility area; this is fixed in the 2007 standard along with adding support for message fragmentation. The others are sim-

ilar – Wireless HART is designed to never lose a message so it gets the nod here while 6LoWPAN does well on the assump tion that the existing TCP/IP protocol suite has class of service built in. While DigiMesh has a similar approach to Wireless HART, it is still somewhat unproven in large deployments.

Power management will no doubt be hotly debated. The nod was given to Wireless HART and DigiMesh because they both define systems where all nodes in the network, including routers, can sleep. Even though sleeping ZigBee end devices are most efficient when it comes to power, the fact that routers can't sleep bumped the rating down. Until 6LoWPAN settles on a mesh and power management strategy, the rating will remain unknown.

The scalability rating follows directly from the question of how big can the network get and still function. This is where the ZigBee 2007 Pro stack shines. The Cluster-Tree architecture creates a hierarchy which enables scalability. DigiMesh and Wireless HART scale well; particularly if most communication is kept local – however, the networks can tend to get very slow when they get too big. Finally, point-to-multipoint has an obvious limitation in the number of nodes that can be attached to one central point.

The best data mover is no doubt the simplest system – namely point-to-multipoint. The simple network design means that focus can be made on short, deterministic latency and high data throughput. There is a direct trade-off here with power. Wireless HART and DigiMesh rate lower here because they are focused on minimizing power and maximizing reliability – this naturally leads to less deterministic latency and lower throughput. I recognize of course, that as a network gets bigger, these two networks will actually do better; however, this is represented in the high scalability ratings for these networks. ZigBee fits in the middle here because the back-bone of powered routers can move data very efficiently – but can get stuck if too many route discoveries are needed.

Cost may end up with the most debate. The ratings here were based primarily on the view of the cost of available chip set solutions under the assumption that the right architecture is chosen for the right job. If not, then the cost ratings go out the window. For example, trying to deploy a ZigBee solution where battery powered routers are desired means infrastructure costs will skyrocket. So given this caveat, point-to-multipoint, ZigBee and DigiMesh have common costs because they all use similar chipsets. 6LoWPAN is somewhat unknown – depending on resource requirements. The assumption is that similar to current chipsets can be used without substantial feature degradation. Wireless HART has a low rating predominantly because the limited number of suppliers has kept chipset prices 5X to 10X other solutions and customers have not demanded lower costs due to primary use on expensive assets in process control environments. This will most likely change as more competitors enter the market.

Conclusion

We have traced the architectures of wireless mesh networks and the respective architectural trade-offs. Each of the WirelessMesh Architectures has respective benefits as they optimize on different components. There is not a one size fits all approach as throughput is traded off against reliability and power consumption. Hence, it is important to match the needs of the application to the capabilities of the network. Further, it is important not to settle for the wrong network because of fad or hype in the market place. No doubt many of the conclusions here will be hotly contested by different network architectural advocates. This is always true where there are shades of gray in evaluation of different criteria. Finally, since this is a view of the current point in time. For example, had this paper been done a year ago, the results will have looked very different – as they will look different a year from now.