NXP Semiconductors
Application Note

Document Number: AN12099
Rev. 0 12/2017

Thread Large Network

1. Introduction

Thread is an IPv6-based mesh networking protocol for connecting products around the home and in buildings to each other, to the Internet and the cloud. Thread networks support full mesh connectivity between all routers in the network, is scalable to hundreds of devices and developed to run on low-power IEEE 802.15.4 chipsets.

2. Thread mesh networking

A mesh network topology provides an effective way of ensuring radio communication between two nodes; it consists of multiple nodes communicating with each other while relaying packets that addressed to other nodes in the network.

Mesh networking improves the reliability and reachability on the network by supporting multiple paths to reach a destination node.

In the Thread mesh topology, every router communicates with every other router within the network, distributing the wireless packets over the network.

Thread mesh networking involves identifying, configuring, and securing links to neighboring devices as the network's membership and physical environment change.

Refer to the Thread Overview whitepaper from threadgroup.org for additional details on Thread Network Topology and Mesh Networking.

Contents

3. Data transmission

3.1. IEEE 802.15.4 data frame

The IEEE 802.15.4 frame is composed of a header, addressing fields, security header, payload and a Frame Check Sequence (FCS).

The header contains the frame control field which indicates the type of frame; it specifies information about the frame and what it contains; there are four types of frames: beacon, data, acknowledgment and command frame.

Any data frame may contain both source and destination information with the size of the address field between 4 and 20 bytes. The security header indicates the security details contained in the frame. The payload field is variable in length, considering the full data frame must not exceed the 127 bytes for the Media access control Protocol Data Unit (MPDU). The FCS is added at the end of the frame to determine if the packet was received correctly.

The PHY handles the preamble sequence, Start Of Frame (SOF) and Frame length bytes. These bytes are not part of the MPDU.

Figure 1 shows the structure of the 802.15.4 data frame.

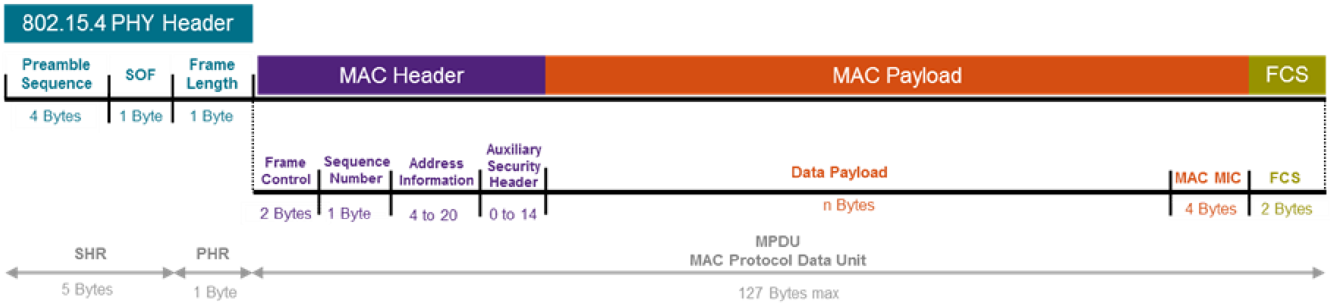

Figure 1. IEEE 802.15.4 data frame

3.2. IPv6 packet and 6LoWPAN

A system of interconnected networks with host-to-host datagram service uses the Internet Protocol (IP). The Internet Protocol Version 6 (IPv6) is the successor of IPv4 addressing standard developed by IETF. IPv6 is being introduced to solve IPv4 address exhaustion problem since IPv6 uses a 128-bit numbering scheme.

Conventional networks use the IP protocol as a base of many application protocols. The Internet of Things has raised the need to incorporate the sensor networks to the same ecosystem.

Most wireless sensor networks use the IEEE 802.15.4 standard for communication. This specification defines a maximum frame length of 127 bytes while IPv6 packets can be up to 1,280 bytes long. This limitation required a new standard to define the adoption and compression of IPv6 packets transmitting over IEEE 802.15.4 data frames. This adaptation is referred to as IPv6 over Low-power WPAN

(6LoWPAN). Figure 2 and Figure 3 illustrate scenarios for a Thread 6LoWPAN frame, some header sizes may vary from packet to packet.

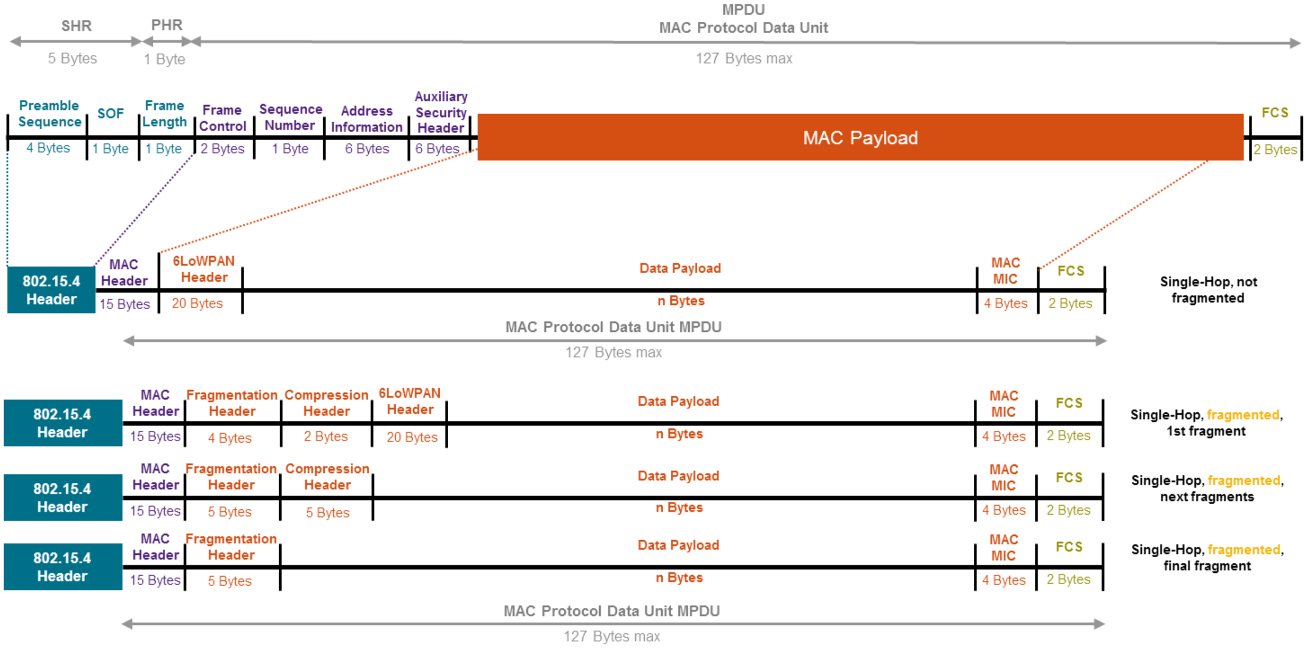

Figure 2. Single-Hop 6LoWPAN frame formats

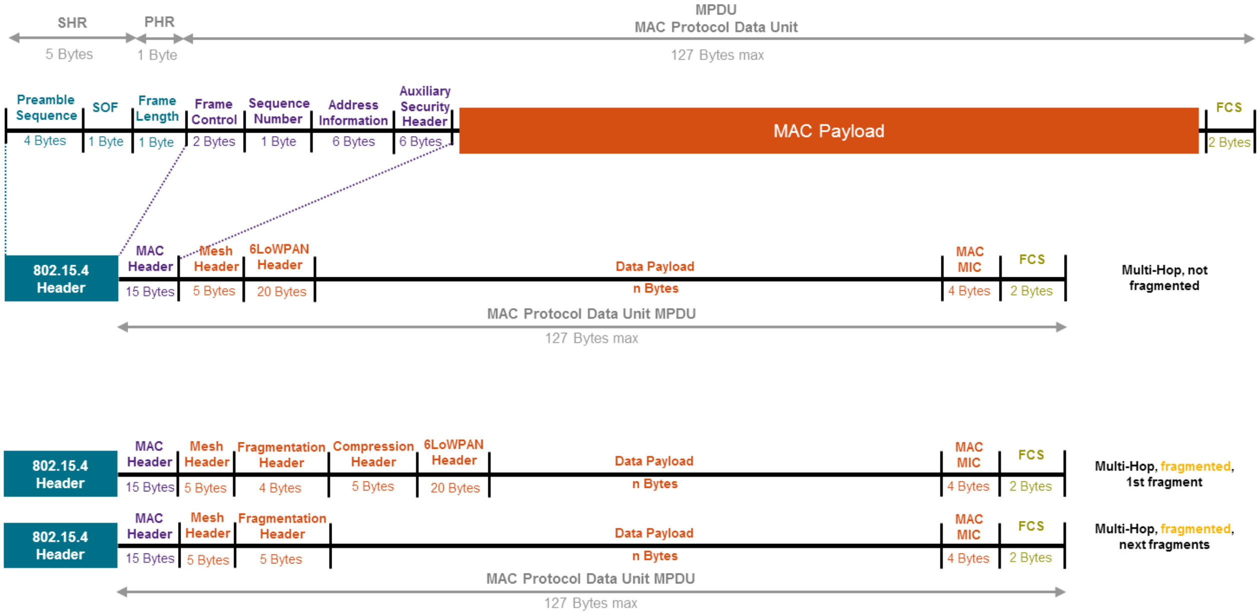

Figure 3. Multi-Hop 6LoWPAN frame formats

3.3. ICMP

The Internet Control Message Protocol (ICMP) is defined by the RFC 792.

The ICMP-Echo Request/Reply is typically used by network devices to verify device reachability over IP-based links. The source address in an echo message is the destination of the echo reply message. To form an echo reply message: first, the source and destination addresses, then type code gets changed, and the checksum recomputed.

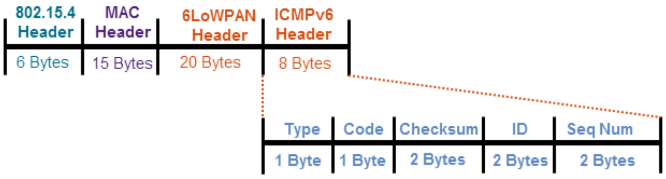

Figure 4. ICMPv6 Header Format

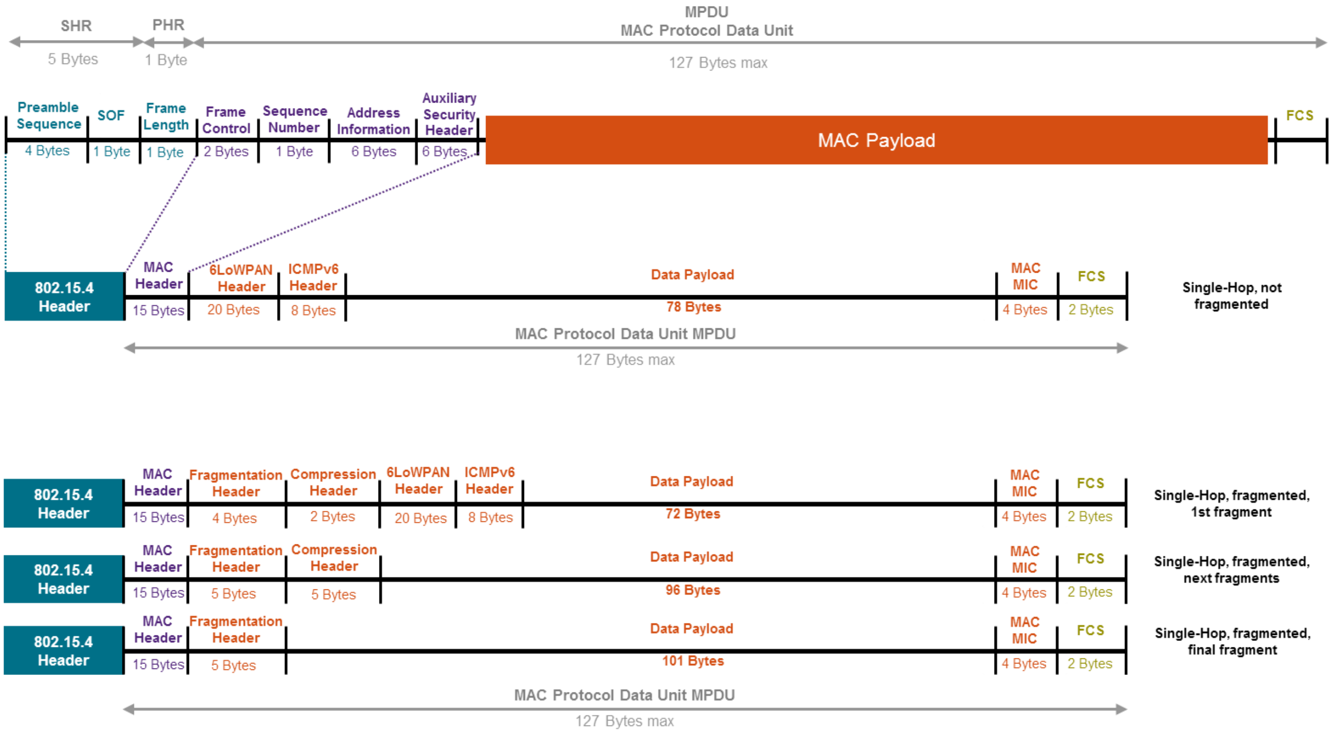

Figure 5. ICMPv6 ping ML64 Single Hop

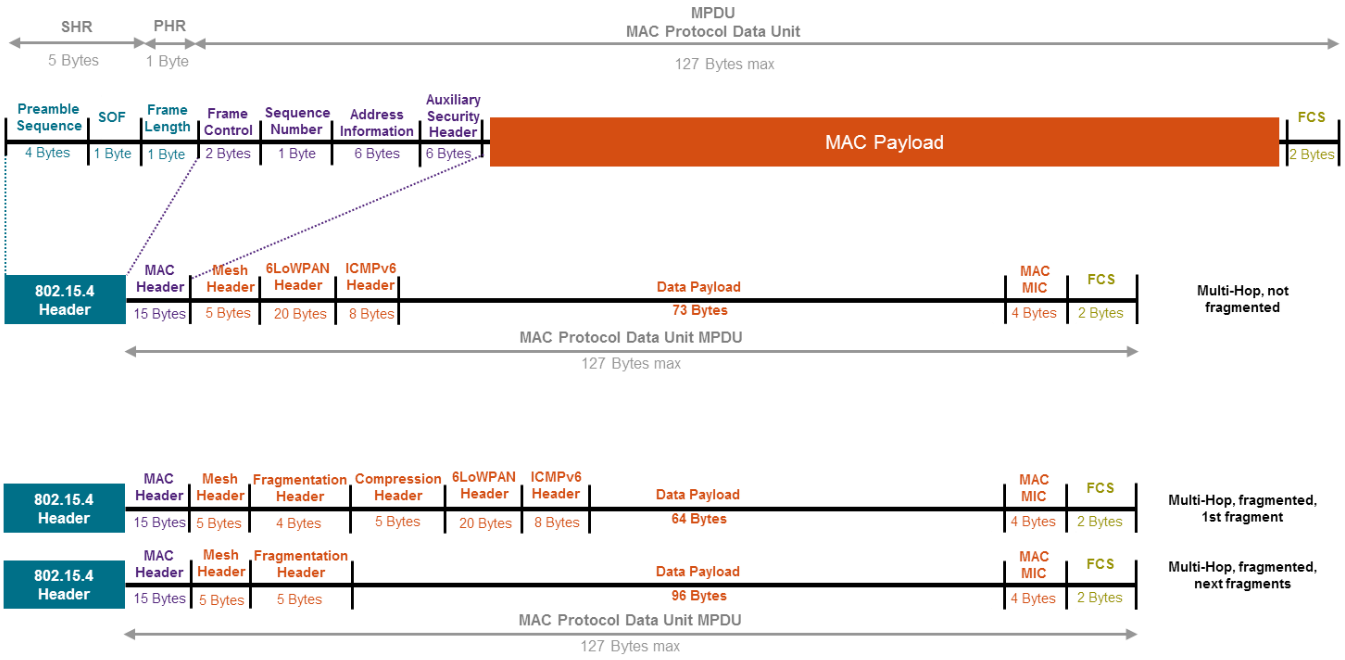

Figure 6. ICMPv6 ping ML64 Multi Hop

3.4. Latency and RTT

In the networking field, the latency is the time it takes a data packet to travel from one node to another. On the other hand, the Round Trip Time (RTT) is the total time a packet takes to reach its destination and return with a reply from the destination node.

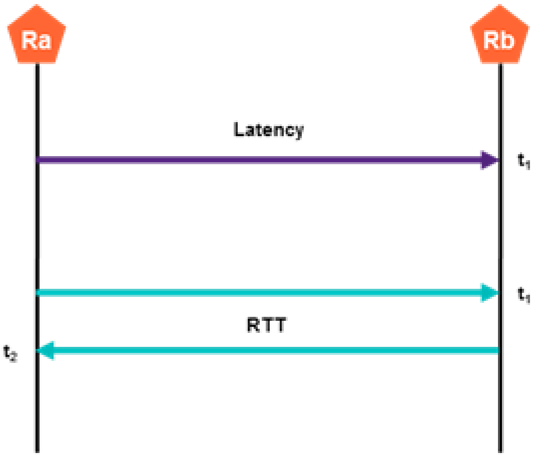

Figure 7. Latency (up) and RTT (bottom)

3.5. Thread: Calculated vs. Measured unicast RTT

Every IPv6 Thread packet translates to an IEEE 802.15.4 data frame. The size of the packet gets defined by the size of headers and payload content. See Application Note JN-AN-1035 for a detailed description on how to calculate data rates of an IEEE 802.15.4 wireless network containing nodes that employ NXP wireless microcontrollers. Depending on the size, a single IPv6 packet may be split into multiple IEEE 802.15.4 frames by the 6LoWPAN module.

A Thread network usually consists of multiple routers that forward packets until they reach their destination or until the max-hop limit is reached. Some packets request a response; this response may or may not travel through the same number of hops until the reply packet reaches its destination. As Thread links are asymmetric, the reply may have a different number of hops to reach its destination.

Based on multiple parameters like packet size, packet processing time, CCA back-off time and MAC acknowledge delay; we can estimate the expected RTT of a packet under ideal conditions. The following estimations and tests consider multiple packet lengths and distances (hop-count).

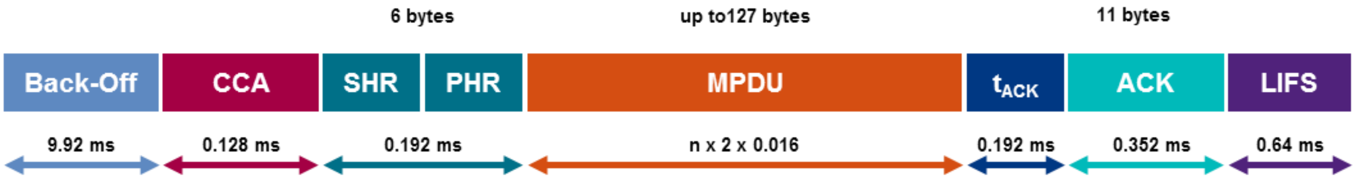

Figure 8. Frame Time

Table 1. MAC max Frame time parameters

Backoff time	9.92 ms	$t_{backoff}$ = random number up to $2^{(minBE-1)}$ × BP.
		BP = Backoff Period = 320 us
		BE = Backoff Exponent = 5 for Thread

CCA	0.128 ms	Clear Channel Assessment time.
802.15.4 Header	0.192 ms	PHY header = SHR + PHR (6 bytes)
MPDU	4.064 ms	MAC Protocol Data Unit (127 bytes)
Tack	0.192 ms	Timing of the transmission of an acknowledgement frame.
ACK	0.352 ms	Max wait time for an Acknowledge frame.
LIFS	0.064 ms	Long Interframe Spacing

The Table 2 calculates RTT for a ML64 ping with payloads 32, 78, 256, or 450 bytes. The RTT considers the previously explained frame time plus the ping software processing time.

Table 2. Calculated Single Hop RTT

ping Payload	Packet Size	# Frag	Latency	RTT
32	81	1	13.4 ms	28.3 ms
78	127	1	14.8 ms	31.2 ms
256	368	3	45.4 ms	92.4 ms
450	624	5	76.4 ms	154.4 ms

Table 3. Calculated Multi Hop RTT

ping Payload	Packet Size	# Frag	Frame Time	RTT Multi Hop				
				2	3	4	5	6
32	86	1	13.5 ms	55.7 ms	82.8 ms	109.8 ms	136.9 ms	164.0 ms
78	172	2	27.7 ms	112.4 ms	167.8 ms	223.2 ms	278.7 ms	334.1 ms
256	381	3	45.8 ms	184.8 ms	276.5 ms	368.1 ms	459.8 ms	551.4 ms
450	668	6	89.3 ms	358.7 ms	537.2 ms	715.8 ms	894.3 ms	1072.9 ms

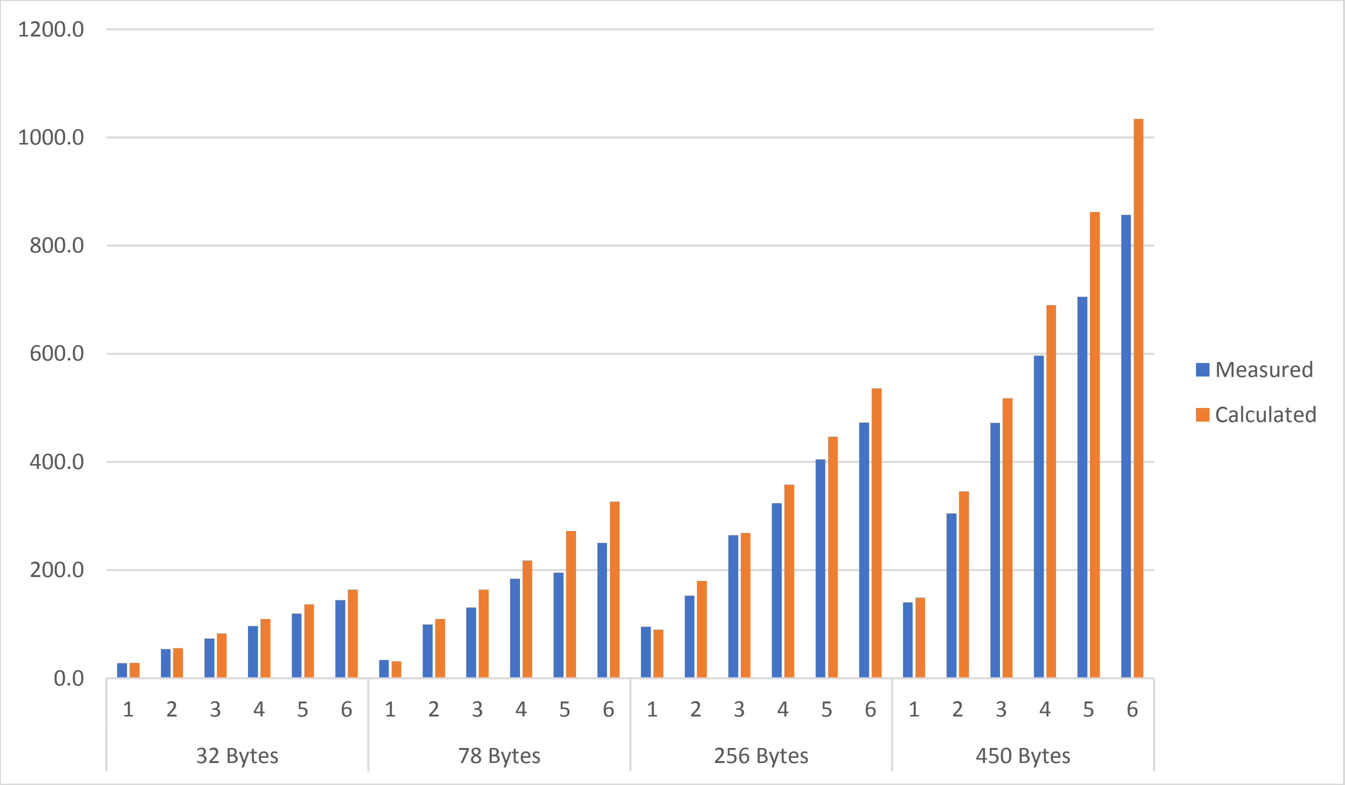

Figure 9. Calculated vs Measured RTT

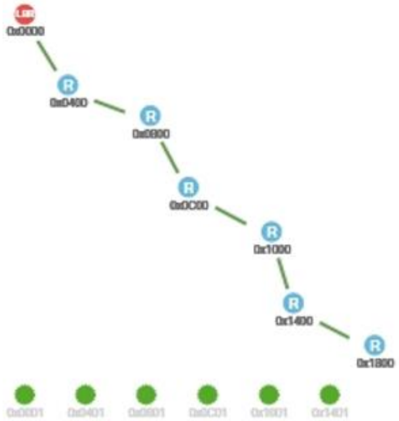

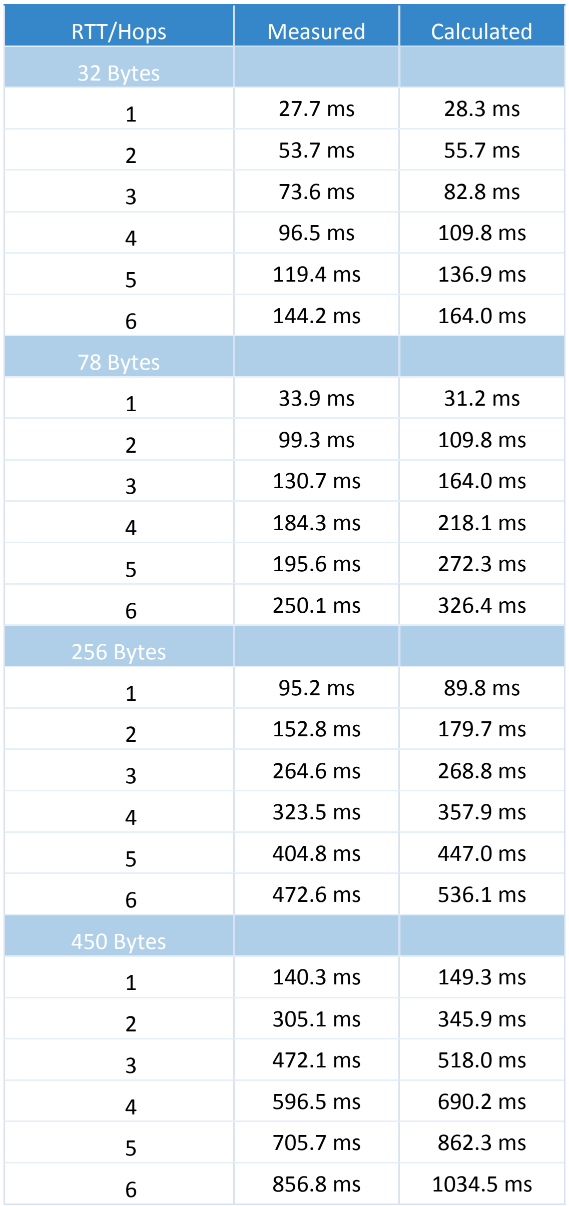

RTT/Hops	Measured	Calculated
32 Bytes		
1	27.7 ms	28.3 ms
2	53.7 ms	55.7 ms
3	73.6 ms	82.8 ms
4	96.5 ms	109.8 ms
5	119.4 ms	136.9 ms
6	144.2 ms	164.0 ms
78 Bytes		
1	33.9 ms	31.2 ms
2	99.3 ms	109.8 ms
3	130.7 ms	164.0 ms
4	184.3 ms	218.1 ms
5	195.6 ms	272.3 ms
6	250.1 ms	326.4 ms
256 Bytes		
1	95.2 ms	89.8 ms
2	152.8 ms	179.7 ms
3	264.6 ms	268.8 ms
4	323.5 ms	357.9 ms
5	404.8 ms	447.0 ms
6	472.6 ms	536.1 ms
450 Bytes		
1	140.3 ms	149.3 ms
2	305.1 ms	345.9 ms
3	472.1 ms	518.0 ms
4	596.5 ms	690.2 ms
5	705.7 ms	862.3 ms
6	856.8 ms	1034.5 ms

Figure 10. Calculated vs Measured RTT (multi-hop, multi-fragment)

4. Multicast

Multicast packets allow a device to send data to a group of nodes within the network. Thread makes use of the Multicast Protocol for Low-Power and Lossy Networks (MPL). With MPL implementations networks can communicate using low-power and lossy links with widely varying topologies.

Thread defines a set of Mesh-Local scopes boundaries for interfaces participating in the same Thread Network. Multiple Mesh Link Establishment (MLE) packets require the use of multicast packets to establish the link costs or for device synchronization.

All Thread interfaces must subscribe to a Link-Local All-Nodes multicast address (FF02::1) and Realm-Local all-nodes multicast address (FF03::1). Interfaces operating like Router, REED, or Border Router must subscribe to the Link-Local All-Routers multicast address (FF02::2) and a Realm-Local all-routers multicast address (FF03::2).

Routers transmit MLE Advertisements to the Link-Local All Nodes multicast address (FF02::1) with a variable interval (see Trickle algorithm) from 0 to 32 seconds. The Realm-Local scope addresses are mostly used for application purposes.

5. Large network implementation

5.1. Hardware infrastructure

The setup consists of an array of clusters distributed through a building floor. Each cluster contains multiple development boards (FRDM-KW41Z and FRDM-KW24D512) connected via serial to an i.MX6 with Wi-Fi enabled to report processing and network activity to a Central Server. The Wi-Fi network was configured on Channel 11 while the Thread network was settled on Channel 23. The wireless interference was minimal as the Wi-Fi communication was mostly used to report results from each cluster to the Central Server after each test is finished.

Software

Version

- Kinetis Thread Stack v1.1 release. Download

Application

- Host Controlled Device

Configuration

- Devices were configured for Out-of-Band commissioning

Hardware

- 250 development boards (109 FRDM-KW41Z and 141 FRDM-KW24D15)
- Boards distributed in 12 clusters. Each cluster with:
  - 18 ~ 22 devices (KW41Z, KW24D)
  - 2 USB Hubs for powering the boards
  - 1 i.MX6 gateway with Wi-Fi module enabled
  - 1 sniffer (USB-KW41Z)
- 1 x Central Test Server

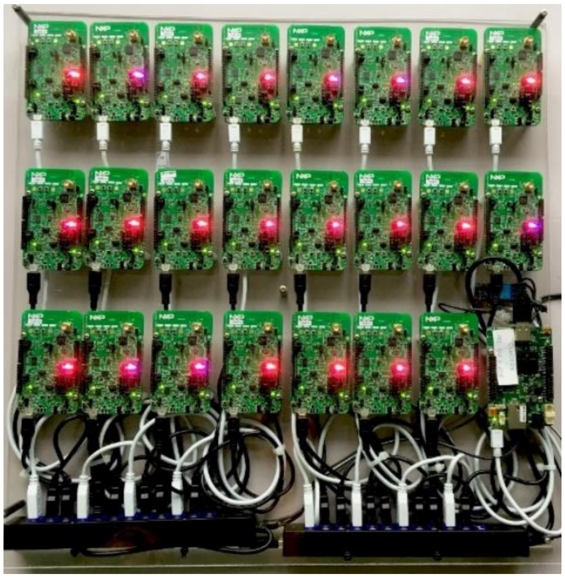

Figure 11. Cluster components

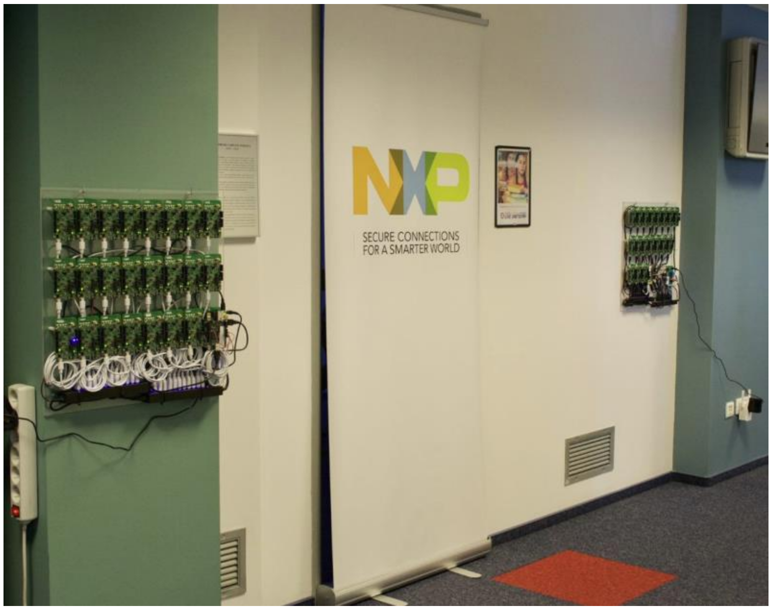

Figure 12. Mounted clusters

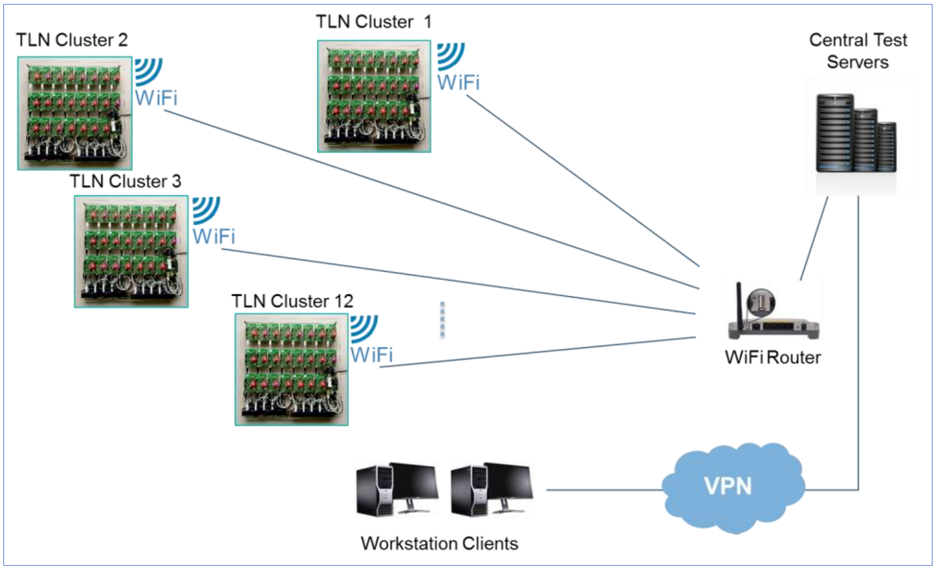

Figure 13. System architecture

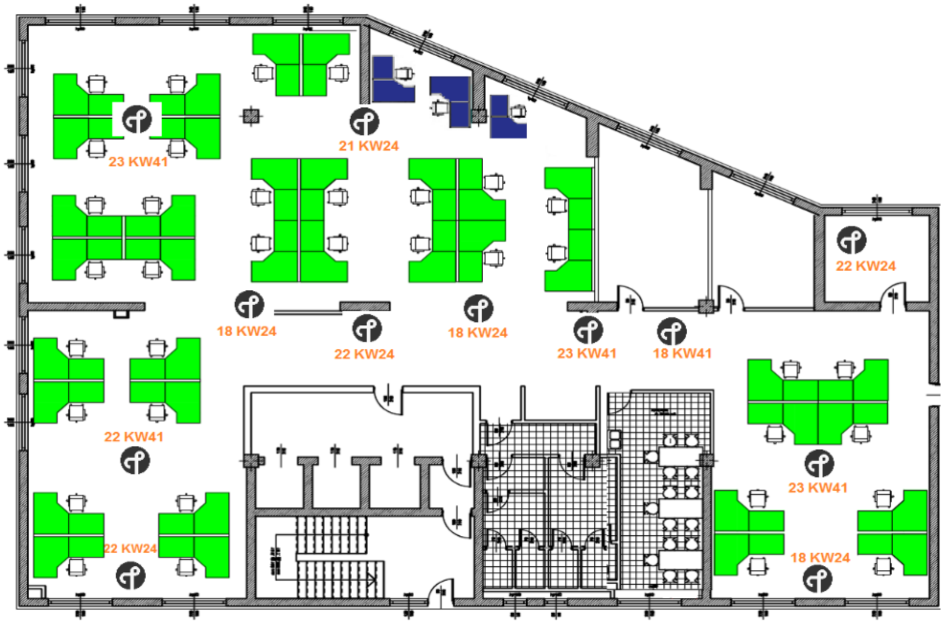

Figure 14. Clusters distribution (NXP Bucharest Office). Area: 450m^2

5.2. Topology

A Thread network topology is allocated and self-configured dynamically depending on multiple factors. Each device joins the network as a Router Eligible End Device. If the number of active Routers on the Network is less than ROUTER_UPGRADE_THRESHOLD, a REED must wait a random period between 0 and ROUTER_SELECTION_JITTER from the time it detected the condition and then, if the condition still holds, must attempt to become an active Router by requesting a Router ID from the Leader (Address Solicit message).

If the number of active Routers on the network exceeds ROUTER_DOWNGRADE_THRESHOLD, an active router that meets a specific criteria (refer to the Thread specification for details) must attempt to release its router ID and become a REED.

The tests executed for this application note resulted in a topology of ~20 routers and ~230 router eligible devices. This was settled automatically by the network based on the number of nodes and their distribution.

6. Wireless interference

Wireless communications are susceptible to interference; this may impact the reliability of a point to point communication link.

Thread is developed to run on IEEE 802.15.4 chipsets which operate in the 2.4 GHz ISM band. This frequency is shared between multiple technologies such as Wi-Fi, Bluetooth, cordless phones, microwave ovens, which may result in some interference, depending on the medium usage from each technology.

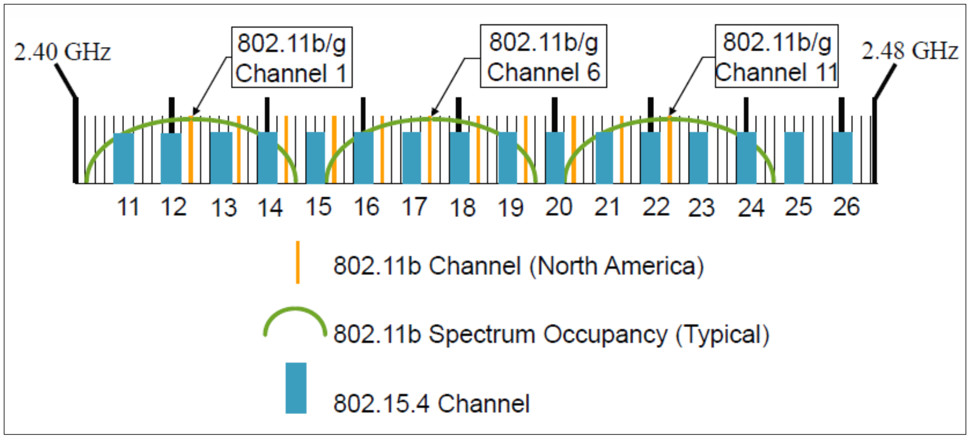

Figure 15. 2.4GHz Channel Occupancy

The tests performed for this application note involve minimal Wi-Fi interference. The Figure 16 is a Wi-Fi scan performed at the NXP Office. The tests were performed under Channel 11 which was mostly used to report results from each cluster to the Central Server after each Thread test (Channel 23) was finished.

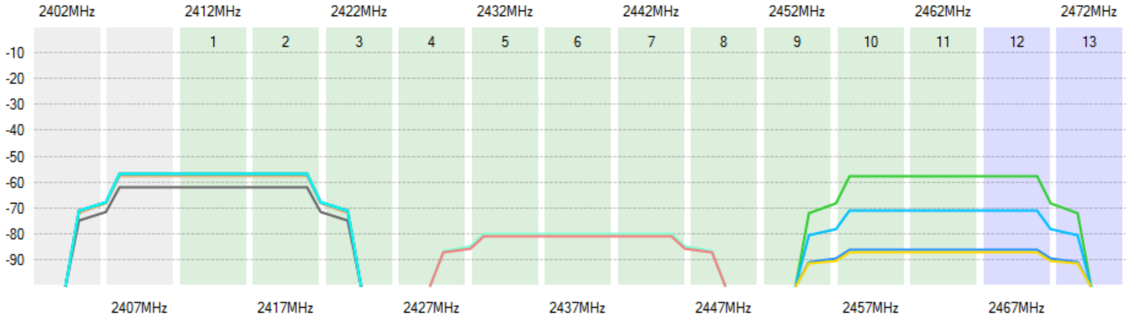

Figure 16. Wi-Fi Occupancy at NXP Office

7. Results

The network topology is described in Chapter 5. Large Network Implementation. The following are the RTT and latency tests results for the 250-node network. After each test, the data was collected and reported to the central hub for analysis via Wi-Fi.

7.1. RTT using ICMP Echo Request

Measure ICMP-Echo.Req (Ping) Round Trip Time (RTT).

1. Add a concentrator using a new router device in the network with THCI enabled.
2. From the new router, request the IPv6 addresses from each device in the network.
3. Send unicast ICMP-Echo.Request with **10 bytes** of payload to each device in a round-robin fashion with a 1-second interval.
4. Collect individual RTT information.

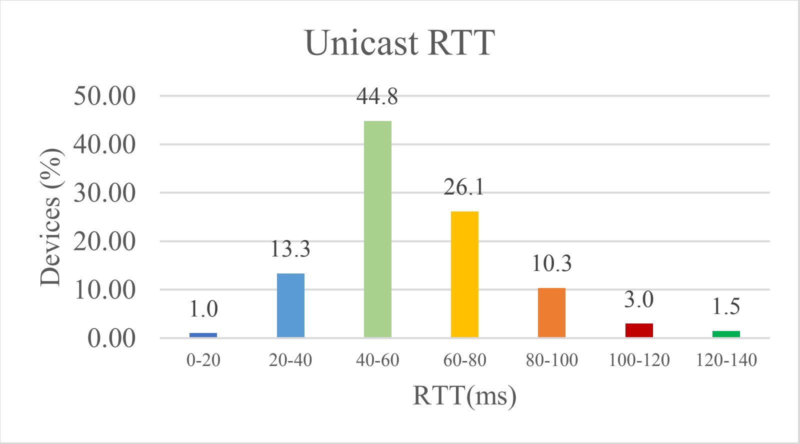

Figure 17. ICMP Echo request RTT

7.2. Unicast latency

Measure unicast latency using PTP (Precision Time Protocol).

1. Add a new router device in the network with THCI enabled that will be used as a concentrator.
2. From the new router, request the IPv6 addresses from each device in the network.
3. Send unicast packet to all devices in the network and wait for reply.
4. Calculate latency

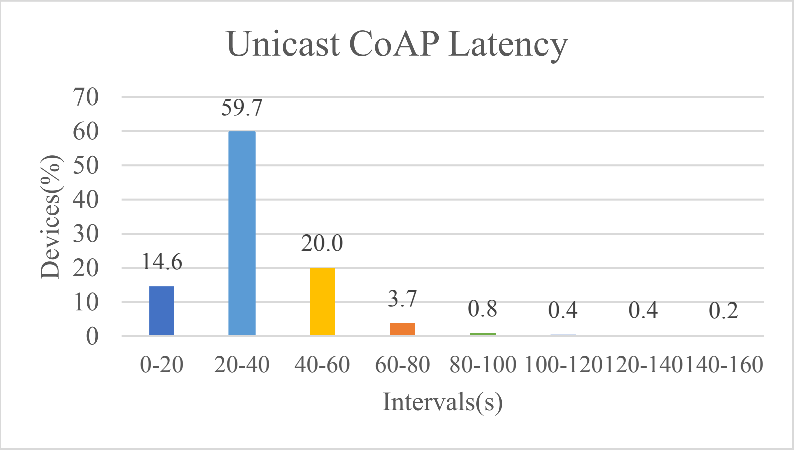

Figure 18. CoAP packet RTT using PTP

7.3. Multicast latency

Measure multicast latency using PTP (Precision Time Protocol).

1. Add a new router device in the network with THCI enabled that will be used as a concentrator.
2. From the new router, request the IPv6 addresses from each device in the network.
3. Send a multicast message.
4. Send unicast packet to get the latency for the multicast message.
5. Calculate the latency.

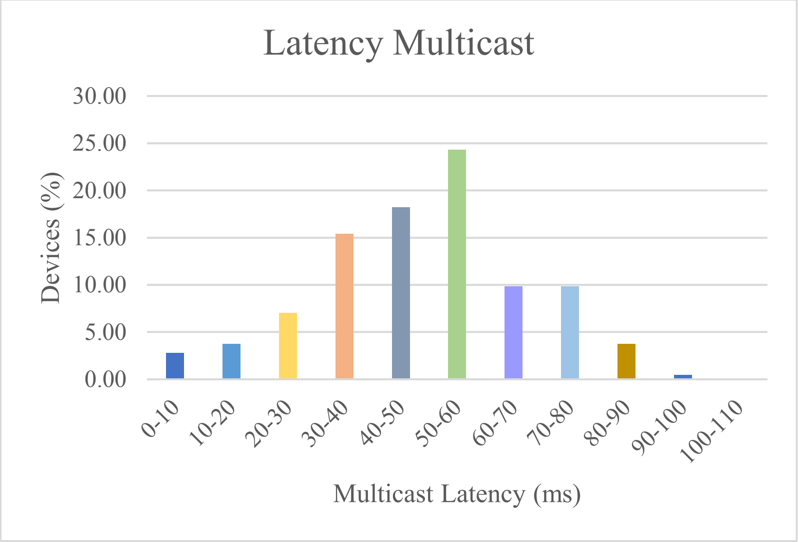

Figure 19. CoAP packet RTT using PTP

- Number of multicast packets sent: 200 packets
- Payload size: 20 bytes
- Success rate: 99.76%*

* From a total of 50,000 unicast packets sent to check if the multicast packet was received. For each multicast packet, 250 unicast packets are sent to read data (sequence number, latency) from all nodes about the multicast packet.

8. Conclusion

Thread networks provide secure and reliable wireless communication. A 250-node network distributed across a building floor provided useful data to understand the behavior of the network and verify the actual topology taken place. The tests executed for this application note involved latency and RTT measurements using ICMP packets, the data on each network node was collected using an i.MX6 present on each cluster and sent via Wi-Fi to an external Server for analysis and interpretation. Based on the different scenarios tested, it is clear that there are multiple packet collisions, resulting in multiple retransmissions during high data traffic in a crowded network but packet retransmissions at IEEE 802.15.4 level take care of many packets that did not receive an acknowledge, meaning it did not reach its destination.

9. Revision history

Revision number	Date	Substantive changes
0	12/2017	Initial release

How to Reach Us:

Home Page:
nxp.com

Web Support:
nxp.com/support

Document Number: AN12099
Rev. 0
12/2017