Application Report

AN123 – Breaking the 400-Node ZigBee® Network Barrier With TI's ZigBee SoC and Z-Stack Software

Suyash Jain

ABSTRACT

While traditional ZigBee networks encompass 5-20 nodes, there are specific use cases that require extending the ZigBee network to several hundred nodes. An example of such a use case is the deployment of a ZigBee wireless sensor network. This application report provides developers with all they need to build, configure and deploy a large network (>400 nodes) with TI's ZigBee SoC and Z-Stack™ software package using an MTO routing scheme.

Table of Contents

1 Introduction 2
2 Understanding the Large Scale ZigBee Network Deployment 3
2.1 Joining a ZigBee Network 3
2.2 Message Reporting 5
2.3 Link Status Messages 6
3 Test Application 7
3.1 Test With Only MAC ACK Required for Transmitted Messages 8
3.2 Test With APS ACK Required for Transmitted Messages 9
4 Z-Stack Knobs 11
5 Conclusion 13
6 References 14
7 Revision History 15

List of Figures

Figure 2-1. ZigBee Network MAC Join Procedure 3
Figure 2-2. (a) Network Join Sequence When the Parent is the Trust Center (b) Ubiqua Log Snapshot Showing Over the Air Messages During Network Join Procedure 4
Figure 2-3. (a) Network Join Sequence When the Parent is not the Trust Center (b) Ubiqua Log Snapshot Showing Over the Air Messages During Network Join Procedure 5
Figure 3-1. Screen Shot of the Test Data Collector Application Running on the PC 8
Figure 3-2. Screen Shot of the Test Data Collector Application Running on the PC 10

List of Tables

Table 3-1. Format of the Message Sent to the Concentrator 7
Table 3-2. Total Network Statistics of the Network 9
Table 3-3. Per Hop Latency Statistics of the Test Network 9
Table 3-4. Percentage of Packets That did not Receive APS ACK Per Hop 10
Table 4-1. Network Parameters in the Z-Stack 11

Trademarks

Z-Stack™ are trademarks of Texas Instruments.
All other trademarks are the property of their respective owners.

1 Introduction

One of the common applications for the low-power wireless ZigBee sensor networks is to collect data from the sensor nodes at a data concentrator. In a network deployment where all nodes are reporting data to a central node, the many-to-one routing scheme provide an efficient routing solution. Since all the nodes in the network maintain a valid route to the central node to send the collected data, one route request broadcast (in the case of the many to one routing scheme) from the concentrator creates the path towards the concentrator. If the AODV routing scheme is used, all the nodes have to initiate the route discovery for the concentrator. The route request broadcasts from each network node add up owing to the large number of nodes and produce huge network traffic overhead. To optimize the routing solution, many-to-one routing is recommended in such networks introduced by ZigBee-Pro 2007 to allow a data concentrator to establish routes from all nodes in the network with one single route discovery and minimize the route discovery broadcast storm.

The *Many-to-One Routing Protocol* section in the *Z-Stack Developer's Guide* [2] provides an introduction of how the many-to-one protocol works and helps to understand its advantage when used in the network where many nodes in the network are reporting the data to a central node.

As the number of nodes reporting data to a concentrator in a ZigBee network are deployed (using Texas Instruments ZigBee Compliant software offering Z-Stack), it becomes important to understand various ZigBee network deployment considerations. Additionally, it is also important to understand various Z-Stack parameters that allow achieving desired network performance, as shown in Section 2. Texas Instruments has setup a 400-plus node network at its San Diego office to demonstrate the robustness of its ZigBee-Pro stack. Section 3 presents information about the test network and presents the test data from the deployed test network using the MTO routing scheme. Section 4 provides parameter values used in the test for the large scale ZigBee network deployment with many-to-one routing scheme.

Note

Results produced in this document involve a CC2538 device as the Zigbee Coordinator along with CC2530s for the Zigbee Routers/End Devices. It uses is a customized version of Z-Stack 2.5.1 that is no longer available or encouraged for project development. The solution accomplishes basic over-the-air communication and assumes a clean radio environment with minimal interference. For newer Zigbee 3.0 designs, it is recommended that developers consider the SimpleLink CC13X2/CC26X2 family of devices. Several LaunchPad EVMs such as the LAUNCHXL-CC26X2R1 or LAUNCHXL-CC1352P-2, which includes an internal PA for up to +20 dBm output TX power, are available for evaluation. The SIMPLELINK-CC13X2-26X2-SDK contains the latest application examples and is regularly updated to include newest features and bug fixes. The other resources available on TIREx such as SimpleLink Academy provide straightforward instructions to get started with the platform. An update to the findings of this application report is provided in *Z-Stack Large Mesh Network Performance Using the SimpleLink™ Wireless MCU Family*.

2 Understanding the Large Scale ZigBee Network Deployment

This section provides an overview of the Zigbee network join process of a new node at the MAC layer: obtaining the current network layer key, announcing the device on the network, and data communication on the network. It is important to understand various considerations for a large network deployment. Starting with an overview of single node network join, this section extends the understanding to be relevant to large network deployment to understand challenges in such a network deployment.

2.1 Joining a ZigBee Network

This section details the consideration for network join in a large network. It divides the network join considerations in four parts: MAC association, obtaining current network key, the device announce and message communication in the network.

2.1.1 MAC Association

Packet Information	Source PANID	Destination PANID	MAC Source Address	MAC Destination Address
Beacon Request		0xFFFF		0xFFFF
Beacon	0x7C47		0xE1F3	
Beacon	0x1235		0x0000	
Beacon	0xE44E		0x0B75	
Beacon	0x0C28		0xC5DE	
Association Request	0xFFFF	0x1235	00:12:4B:00:00:00:00:02	0x0000
Acknowledgement				
Data Request		0x1235	00:12:4B:00:00:00:00:02	0x0000
Acknowledgement				
Association Response		0x1235	00:12:4B:00:00:00:00:01	00:12:4B:00:00:00:00:02
Acknowledgement				
Transport Key		0x1235	0x0000	0x41CF
Acknowledgement				
Device Announce		0x1235	0x41CF	0xFFFF
Device Announce		0x1235	0x0000	0xFFFF

Figure 2-1. ZigBee Network MAC Join Procedure

The joining device issues a beacon request in response to which all routers and the coordinators within RF range respond with 802.15.4 beacon. The device selects a parent and then issues an association request and receives association response with short-address on the network. The concentrator provides the current network key via the Transport Key command and the device then announces itself on the network.

A ZigBee network is started by a coordinator and then the network devices (routers and end-devices) join the network. The first step of the network join process involves active scan of the channels (IEEE 802.15.4 channels 11 to 16 or specific channels configured by the application) to identify the existing ZigBee networks in the vicinity. Active scan involves sending out beacon request(s) to which coordinator and router(s) within RF range of the device trying to join the network respond with a IEEE 802.15.4 beacon (see Figure 2-1). The new device selects one parent from various beacons it may receive and issues an “Association request”. The selected parent, if allowing network join, then issues an Association response providing the new device with a short address on the network (see Figure 2-1).

For large network deployments, it is important to understand that if several routers are present in the vicinity of the joining device, there can be a lot of devices sending out IEEE 802.15.4 beacons in response to an active scan beacon request, and can cause increased overhead in the network at the time a new device tries to join the network. Additionally, if several devices are trying to join the network at once, in a large network deployment with hundreds of nodes, the increased traffic can cause packet transmission failure for brief periods if any data

communication is on-going at the time the new devices are trying to join the network. It is recommended to deploy the network where new device network join times are randomized to avoid this increased traffic.

2.1.2 Obtaining Current Network Key

If network security is enabled and the key is not pre-configured, then network key is delivered over the air.

If secured transport of the network key is used, then it is required that all devices have a pre-configured Trust Center Link Key (TCLK) and that the network key is delivered to joining devices secured (encrypted) with this key. There are basically two joining scenarios for the device, which are important to understand for a large network deployment.

Figure 2-2 shows the network join when the parent is the trust center of the network. The device issues an association request to the parent and obtains the current network key via a transport key command from the trust center, which is APS encrypted with the trust center link key between the trust center and the joining device. After that, the new device transmits a Device Announce message over the air.

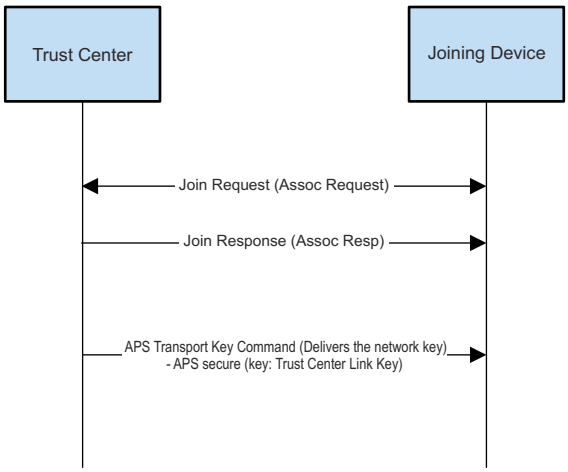

Packet Information	Source PANID	Destination PANID	MAC Source Address	MAC Destination Address
Beacon Request		0xFFFF		0xFFFF
Beacon	0x7C47		0xE1F3	
Beacon	0x1235		0x0000	
Beacon	0xE44E		0x0B75	
Beacon	0x0C28		0xC5DE	
Association Request	0xFFFF	0x1235	00:12:4B:00:00:00:00:02	0x0000
Acknowledgement				
Data Request		0x1235	00:12:4B:00:00:00:00:02	0x0000
Acknowledgement				
Association Response		0x1235	00:12:4B:00:00:00:00:01	00:12:4B:00:00:00:00:02
Acknowledgement				
Transport Key		0x1235	0x0000	0x41CF
Acknowledgement				
Device Announce		0x1235	0x41CF	0xFFFF
Device Announce		0x1235	0x0000	0xFFFF

Figure 2-2. (a) Network Join Sequence When the Parent is the Trust Center (b) Ubiqua Log Snapshot Showing Over the Air Messages During Network Join Procedure

The second scenario is when a device joins the network, but its parent are not the Trust Center. The transport key command is tunneled from the Trust Center, through the parent of the joining device, to the joining device. The joining procedure is illustrated in Figure 2-3. Notice that the APS Update Device command, sent from the parent to the trust center, is network layer encrypted. The APS Tunnel Command with theAPS Transport Key command as the payload is also network layer encrypted, but the payload is APS layer encrypted with the trust center link key between the trust center and the joining device. Finally, the APS Transport Key command forwarded from the parent to the joining device is APS encrypted with the trust center link key between the trust center and the joining device.

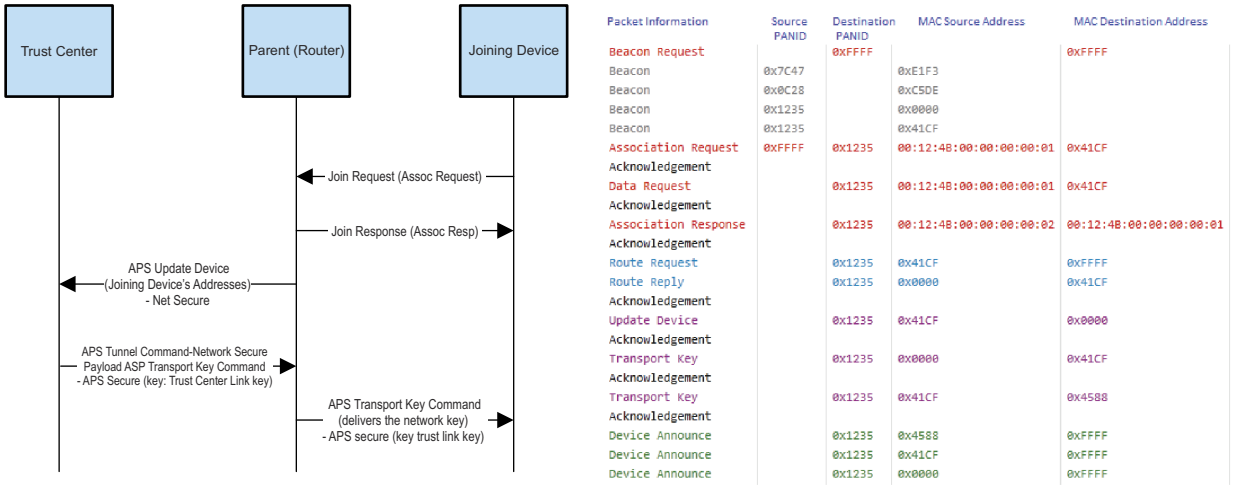

Packet Information	Source PANID	Destination PANID	MAC Source Address	MAC Destination Address
Beacon Request		0xFFFF		0xFFFF
Beacon	0x7C47		0xE1F3	
Beacon	0x0C28		0xC5DE	
Beacon	0x1235		0x0000	
Beacon	0x1235		0x41CF	
Association Request	0xFFFF	0x1235	00:12:4B:00:00:00:00:01	0x41CF
Acknowledgement				
Data Request		0x1235	00:12:4B:00:00:00:00:01	0x41CF
Acknowledgement				
Association Response		0x1235	00:12:4B:00:00:00:00:02	00:12:4B:00:00:00:00:01
Acknowledgement				
Route Request		0x1235	0x41CF	0xFFFF
Route Reply		0x1235	0x0000	0x41CF
Acknowledgement				
Update Device		0x1235	0x41CF	0x0000
Acknowledgement				
Transport Key		0x1235	0x0000	0x41CF
Acknowledgement				
Transport Key		0x1235	0x41CF	0x4588
Acknowledgement				
Device Announce		0x1235	0x4588	0xFFFF
Device Announce		0x1235	0x41CF	0xFFFF
Device Announce		0x1235	0x0000	0xFFFF

Figure 2-3. (a) Network Join Sequence When the Parent is not the Trust Center (b) Ubiqua Log Snapshot Showing Over the Air Messages During Network Join Procedure

For large network deployments, it is important to note in Figure 2-3 that if the parent is not the trust center, there may be a condition that the parent of a new device may not have a path established towards the trust center. So, this parent will issue a route request message broadcast (AODV) to find the route towards the Trust center. This enables the new parent to send the update device command to the trust center. The broadcast increases network traffic. If many devices try to join the network simultaneously, various route requests (broadcasts) from multiple parents will occur. These being broadcast messages will be re-broadcasted by coordinator and all routers in the network, increasing the network traffic. Also, these route requests, due to increased traffic, can take considerable time to be resolved and the network throughput may go down during the time various devices are trying to join the network. Additionally, if following the route request and the corresponding route reply message does not come in time, due to increased network traffic, the joining device may time out waiting for transport key command. The new device will then repeat the whole process of network join again. This can cause significant network overhead especially as the device count increases.

Additionally, if the network communication is on-going when several new devices are joining the network, then such route request broadcast can add up affecting network performance (throughput) during network join of the devices. It is recommended to have devices join the network not all at once but with some delay between the network joins of these devices.

2.1.3 Device Announce

Device after obtaining the current network key will issue a device announce as a broadcast. This will be re-broadcasted by all the routers and coordinator. If simultaneously, many devices try to join the network. These broadcasts may cause network throughput to go down during that time. And may delay network join of devices that are trying to attempt network join at this time.

When several devices join the network, or are attempting to join the network, these devices announce broadcasts in addition to other broadcasts (IEEE 802.15.4 Beacons, possible route requests from parents as explained in Section 2.1 to find the path to the concentrator) can add up and cause the network performance to suffer. To avoid this, the devices should be brought in a large ZigBee network one by one, (with some delay between network joins of various nodes).

2.2 Message Reporting

After successfully joining the network, network devices would transmit data to the concentrator which can be periodic or based on some event on the nodes. It is important that the network traffic in large network be kept low to not burden the concentrator on processing the received messages. To understand the network traffic

considerations, we will consider two scenarios for packet transfers in the network based on the type of acknowledgment required by the sender node of the messages.

Case 1: MAC ACK Required: If the concentrator in the network is not required to send the APS ACK to the messages from the nodes in the network then the network only operates in one direction. In this case concentrator does not need to know or create the routing information to send packets to the nodes in the network. If all nodes in the large network require only the MAC ACK from the next hop then care must be taken that all devices do not transmit the messages at the same time and that the concentrator has sufficient time to process the messages from the nodes.

Case 2: APS ACK required: If APS ACK or other application layer reply message is required to be sent to the messages sent from the nodes then it is required that concentrator knows the path back to the devices. To avoid the network overhead of concentrator discovering paths to all the nodes in the network using the AODV routing, ZigBee-Pro-2007 also defines Source Routing. When using the source routing, the reporting devices in the network using the 'many to one routing' scheme sends out a route record command before sending the packets (as explained in the *Many-to-One Routing Protocol* section of the *Z-Stack Developer's Guide* [2], which has the path that the message takes to reach the concentrator. This path is stored in the source routing table on the concentrator. Depending on the implementation, if there is space in the source routing table to store paths to all the nodes in the network or if there is an expired entry that can be replaced. The concentrator will store the path back to the sending device. Then to send the APS ACK or the application layer packet in to the node the concentrator will use this path without having to discover the route to the sending node.

It is important to understand that the size of source routing table on the concentrator is very critical in Large ZigBee network operation using MTO routing scheme. Ideally, the concentrator should have sufficient memory to store path back to all the devices so that it does not do an AODV route request to find the path to the network devices for sending acknowledgment or message packets.

If the concentrator is memory constraint, it is very important to note the source route expiry timeout (SRC_RTG_EXPIRY_TIME) value should be such that no more than (MAX_RTC_SRC_ENTRIES) devices reporting within this interval. In this case concentrator should be able to store the paths back to all the nodes as they report data (even if the concentrator does not have memory to store path to all devices) as expired routes will be replaced with new routes and the network should not see broadcast route request storms originating from the concentrator which can cause the network throughput and reliability to go down significantly.

2.3 Link Status Messages

Wireless links may be asymmetric, that is, they may work well in one direction but not the other. For many-to-one routing, it is a requirement to discover routes that are reliable in both directions. To accomplish this, routers exchange link cost measurements with their neighbors by periodically transmitting link status frames as a one-hop broadcast. The reverse link cost information is then used during route discovery to ensure that discovered routes use high-quality links in both directions. [1]

In large network deployments where many nodes are placed close to each other, the neighbor table entries (MAX_NEIGHBOR_ENTRIES) should be increased to allow the possibility of choosing the best neighbor for network communication.

Also to reduce the network traffic due to too many nodes transmitting the link status message (Z-Stack Default timeout is 15 seconds), the timeout value can be increased to say 30 seconds to reduce OTA messages. This may help to reduce the CSMA CA back-offs on the packet transmissions in the network, especially in case of large network deployments where many nodes are within RF range of each other.

3 Test Application

Texas Instruments has setup a 400-plus node network at its San Diego office to demonstrate the robustness its ZigBee-Pro stack. This section presents test data from the deployed test network using Many To One Routing scheme and also provides parameter values used in the test for large scale ZigBee network deployment with many-to-one routing scheme.

The test application implements a ZigBee private profile-based network where each node transmits a 46 byte payload to the data concentrator. Table 3-1 shows the fields of the data being transmitted and index in the transmitted data.

Table 3-1. Format of the Message Sent to the Concentrator

Field	Byte Index
Command	0
IEEE Address of the Node	1
Parent Short Address	9
Application Build Code	11
Temperature	11
Voltage	12
Node Type	13
Source Routing Enabled or Disabled	14
Routing Index	15
Attempted TX Packets	16
Successful TX Packets	19
Received APS ACK to TX Packets	23
Minimum Latency	25
Maximum Latency	28
Average Latency	31
Invalid Network Packets Counter	34
MAC Packets RX Counter	38
MAC Packets With Bad FCS	42

The entire network uses security at the network layer and the key is distributed by the trust center (also the concentrator in the test deployment) to the joining device. When the network is first started and devices have joined the network, the concentrator sends a many-to-one route request (MTOR) so that all 400 nodes can discover a route to the concentrator. The concentrator is configured as one with memory. The source routing table size is set to 430. The concentrator uses the source routing table to be send the APS_ACK back to the nodes. For the test, network is setup in both configurations, one where all message require APS ACK and second where APS ACK is not required, messages only receive MAC ACK.

On start-up, each node does not transmit the application data automatically. Once the nodes have joined the network. The concentrator application sends a broadcast “announce” message that tells each node to start transmitting and how often to transmit. Each node is by default configured to transmit within an interval window of 3 minutes (180 seconds) where 60 seconds is a jitter period. This is done in order to ensure that the number of nodes trying to transmit around the same time is minimized. This is critical for achieving the best network performance as the number of collisions and thus unsuccessful transmissions would increase if every node attempted to transmit at roughly the same time. The data concentrator is a CC2530 device with a universal asynchronous receiver/transmitter (UART) connection at 115200 bps (8-N-1). The data is collected and provided to a PC application that displays the performance characteristics of the network as shown in Figure 3-1 and Figure 3-2.

Each node keeps track of performance statistics and sends this information to the concentrator as part of the 46 byte payload as listed in Table 3-1. The information in Section 3.1 and Section 3.2 represents performance data gathered from the network.

Several Z-Stack parameters were tuned and optimized for deployment in the large network. Section 4 presents a list of optimized parameters and their descriptions, which can be used by the developers to implement large networks based on the many-to-one routing scheme.

3.1 Test With Only MAC ACK Required for Transmitted Messages

A test network with packets only requiring MAC ACK from next hop in the path towards concentrator was setup with 405 nodes. The network statistics are presented below. Also, Section 4 presents the important stack parameters and their values that were tuned for the network operation.

3.1.1 Node Performance Data

This section presents the averaged data statistics for each router node in the network.

- Average Latency (data confirm) = 23.4 ms
  - The application at each node calculates latency as time difference when the packet was transmitted and the MAC ACK is received from the next hop for the message.
  - This number is average of the minimum latency observed per node in the network.
- Average TX efficiency = 99.98 %

$$\text{Defined as } \frac{\text{Number of packets transmited which recieved MAC ACK}}{\text{Number of packets attempted to be transmitted}} * 100 \quad (1)$$

- Max hop count = 4

3.1.2 Throughput Statistics (At Concentrator)

Network throughput statistics as collected at the concentrator node are presented below:

- Average reports received per interval (180 seconds) = 495
- Interval window = 180 sec
- Average reports per second = 3
- Average bps (UART data transfer to PC application) = 1128 bps

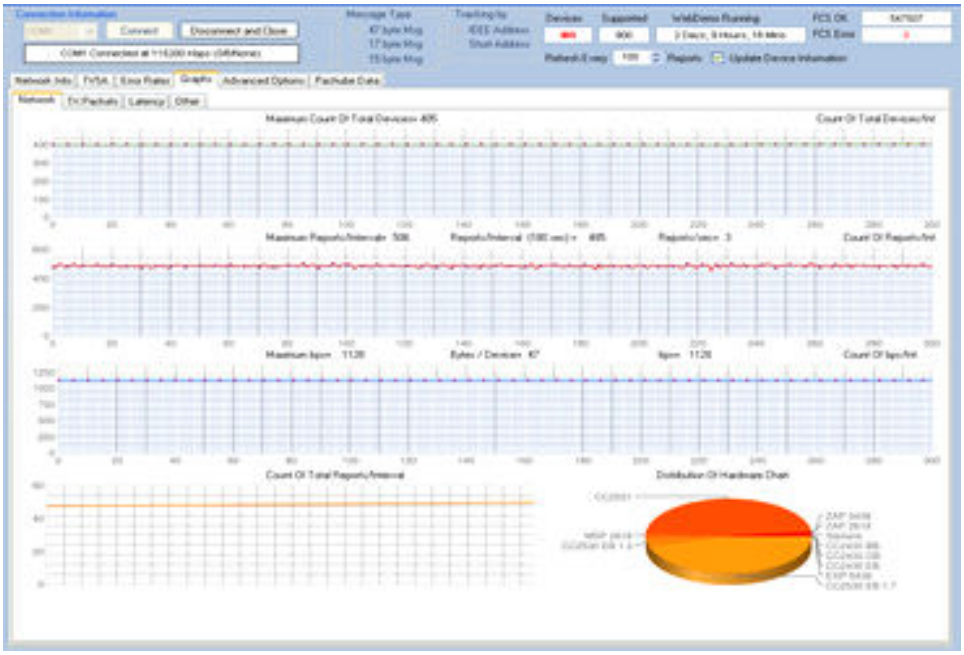

Figure 3-1. Screen Shot of the Test Data Collector Application Running on the PC

The application is interfaced with the concentrator via UART. Figure 3-1 shows the Total Number of devices in the network, Number of Report/Interval (Interval = 420 seconds), Maximum bits per second, total number of reports that are received per interval, and the pie chart showing distribution of various TI ZigBee platforms, Figure 3-1 shows half of the devices in the network are CC2531 System-on-Chip (SoC) and the half are the CC2530 SoC running the Z-Stack and a private profile application.

3.2 Test With APS ACK Required for Transmitted Messages

This section presents the averaged data statistics for each router node in the network.

- Average Latency (data confirm – APS ACK) = is provided in Table 3-3.
  - The application at each node calculates latency as time difference when the packet was transmitted and the APS ACK was received from the next hop for the message.
  - The numbers in Table 3-3 are an average of the latency observed at nodes at a given hop in the network.
- Average TX efficiency = 99.79 %

$$\text{Defined as } \frac{\text{Number of packets that received MAC ACK}}{\text{Number of packets transmitted}} * 100 \quad (2)$$

- Max hop count that was observed in the test network = 4

3.2.1 Total Network Statistics

Network throughput statistics as collected at the concentrator node are presented below:

- Average reports received per interval (420 seconds) = 440
- Interval window = 420 sec
- Average reports per second = 1

Table 3-2 presents the total network statistics while network data was collected for this application note. It shows how many packets were transmitted by all the nodes in the network combined, number of packets that received APS ACK and percentage of the packets that did not receive the APS ACK. About 0.2% packets did not receive APS ACK in the test network. This number can be reduced for example by reducing the network traffic and thus avoiding packet collisions, increasing application layer retries.

Table 3-2. Total Network Statistics of the Network

Total Packets TXed	Total Packets Ack'ed	% of Packets That did not Receive APS ACK
706017	704542	0.20

3.2.2 Per Hop Statistics

Table 3-3 shows the latency calculated at time between the message transmission and reception of APS ACK indication at the application layer for the devices classified based on how many hops the messages take to reach the concentrator. In a large network, the average network latency at hop 2 and above can be reduced by reducing the network traffic. The best achievable latency observed at each hop in the test network is listed under minimum latency.

Table 3-3. Per Hop Latency Statistics of the Test Network

Hop-No	Min Latency (msec)	Average Latency (msec)	Number of Devices
Hop-1	39.6	75	26
Hop-2	59.5	109.2	176
Hop-3	63.8	175.2	175
Hop-4	75.5	162.2	25

Table 3-4 shows the number of packets that were transmitted per hop, number of transmitted packets that received APS ACK and the percentage of packets that did not receive APS ACK per hop.

Table 3-4. Percentage of Packets That did not Receive APS ACK Per Hop

Hops-No	Total Packets Attempted	Total Packets Ack'ed	Percent Error
Hop-1	45328	45281	0.10
Hop-2	309250	308704	0.17
Hop-3	307278	306510	0.24
Hop-4	44161	44047	0.25

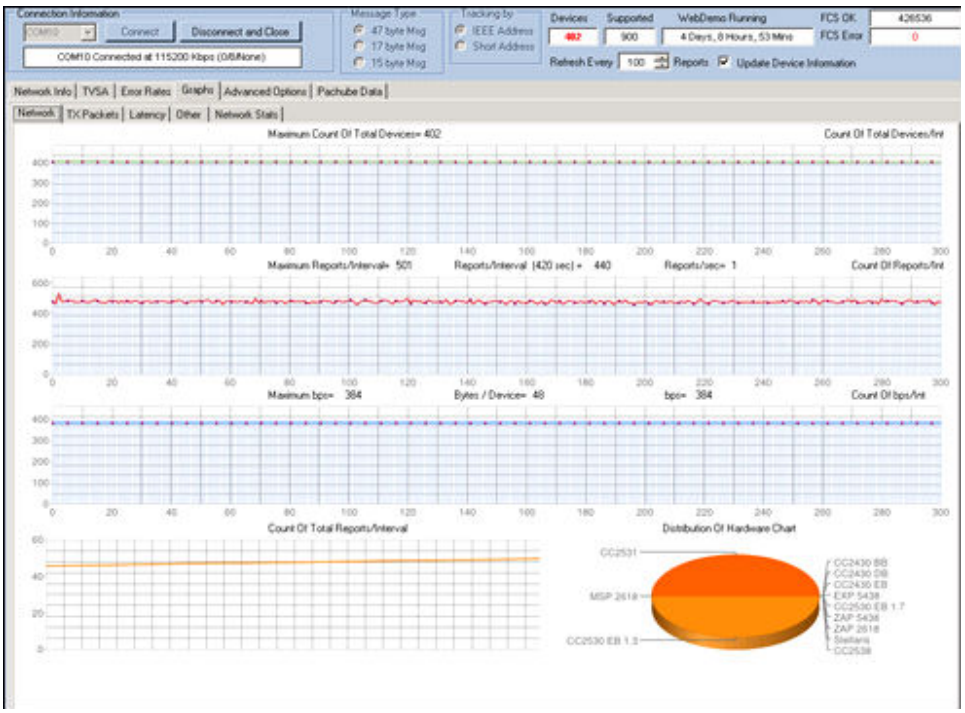

Figure 3-2. Screen Shot of the Test Data Collector Application Running on the PC

The application is interfaced with the Concentrator via UART. Figure 3-2 shows the Total Number of devices in the network, Number of Report/Interval (Interval = 420 seconds), Maximum bits per second, Total number of reports that are received per interval, and the pie chart showing distribution of various TI ZigBee Platforms. Figure 3-2 shows half of the devices in the network are CC2531 SoC and the half are the CC2530 SoC running the Z-Stack and a private profile application.

4 Z-Stack Knobs

Table 4-1 discusses network parameters in the Z-Stack that are optimized to achieve this 400+ node network. Table 4-1 also lists the advantage of optimizing these values so that developers can understand and tune the parameters as per their network deployments. The parameters are given in three parts: first parameters pertaining to the concentrator node only are provided, then for the concentrator and router nodes both and then a parameter for the router nodes only in the network.

Table 4-1. Network Parameters in the Z-Stack

Z-Stack Compile Option	Description	Value	Advantage
Concentrator Only			
INT_HEAP_LEN	Defines the heap size for the coordinator	3280 (bytes)	
NWK_MAX_DEVICE_LIST	Defines the number of children the coordinator allows	10	Save RAM on the concentrator by reducing from default value of 20.
CONCENTRATOR_ENABLE	Enables this device to be a concentrator node and turns on the periodic MTOR event at the network layer	1	Needs to be defined equal to one only on the concentrator node
CONCENTRATOR_ DISCOVERY_TIME	Controls the MTOR period	120 (seconds)	Route Request to establish a path to the concentrator is sent every 120 seconds.
MAX_RTG_SRC_ENTRIES	Sets Source Routing table size	430	Default value is 12, increasing the size allows the concentrator to store the path back to various nodes in the network. This avoids concentrator to have to discover routes via AODV mechanism which would otherwise increase network traffic
SRC_RTG_EXPIRY_TIME	Time for which the route is valid in the Source Routing table	2	For details of this application note, see Section 2.2.
CONCENTRATOR_ENABLE	Needs to be defined for a node to become concentrator at compile time	1	
CONCENTRATOR_ ROUTE_CACHE	Sets the concentrator as with memory for Source routing table or one with limited memory	1	In test network concentrator was one with memory
MTO_RREQ_LIMIT_TIME	Controls the delay of the MTOR. In the case of a route error, the next MTOR by the concentrator is delayed by 5 seconds if an MTOR is already in progress.	5000	This limits the route request storm of the previous MTOR.
Concentrator and Routers			
LINK_DOWN_TRIGGER	Every router keeps track of successful and unsuccessful transmissions to neighboring routers. If the TX count for unsuccessful transmissions exceeds the value of LINK_DOWN_TRIGGER, the next hop is marked as bad and not used in subsequent routes	12	For cases when a neighbor is either removed or moved to a different location in the network neighbor table needs to be updated..But in a large network due to increased traffic the failures on transmission to the neighbor nodes may be higher so LINK_DOWN_TRIGGER is set to (increased) 12 from default value of 3.
NWK_ROUTE_AGE_LIMIT	Controls the number of link status messages missed before a neighbour is aged out from the neighbour table. Default value = 3	30	In a large network where there is a lot of traffic, more than the default value= 3 link status messages may be missed by a router, so the value is increased for the large network implementations to reduce false aging of the neighbours. This can cause route discovery message to be sent increasing the network traffic, even though a node is still a neighbor.

Table 4-1. Network Parameters in the Z-Stack (continued)

Z-Stack Compile Option	Description	Value	Advantage
BCAST_DELIVERY_TIME	Amount of time a broadcast message lives within the network	100	Increased from default of 70 as for large network it can take time for message to propagate in the network
DEF_NWK_RADIUS	Controls the number of hops that messages can travel in the network	15	
NWK_ROUTE_AGE_LIMIT	Controls the number of link status messages missed before a neighbour is aged out from the neighbour table. Default value = 3	30	In a large network where there is a lot of traffic, more than the default value= 3 link status messages may be missed by a router, so the value is increased for the large network implementations to reduce false aging of the neighbours. This can cause route discovery message to be sent increasing the network traffic, even though a node is still a neighbour.
DEFAULT_ROUTE_ REQUEST_RADIUS	Controls the radius of route request packets	8	Reduced from the default of 30 to reduce the route request storm for a large network.
ROUTE_DISCOVERY_TIME	Amount of time a route discovery lasts. When the network is formed, AODV routing is utilized to send the network key to each joining node. This value limits the broadcast route request storms.	13	For large network it can take time for the Route Response packets to come from the destination node. So, at each router node involved in the Route Formation timeout waiting for Route Response is increased from default of 5 seconds to 13 seconds.
ZDNWKMGR_MIN_ TRANSMISSIONS	Turns off frequency agility	0	Frequency agility is not implemented in this test.
NWK_LINK_ STATUS_PERIOD	Timeout after which a router/ coordinator node will send a link status message (Default 15)	30	Nodes send the link status messages every 30 seconds. Reducing the frequency of OTA messages reduces network traffic in large network especially in dense networks
Router Only			
MAX_NEIGHBOR_ ENTRIES	Controls number of neighbour entries for each router.	48	Increasing this number maximizes the number of “1-hop” routes, as the node can find many nodes as its immediate neighbour and can send data directly without performing a route discovery. In a large network this is useful as this reduces the Route discoveries that would otherwise be required.

5 Conclusion

This application report explains the large network (400+ nodes) deployed at the TI San Diego office and provides key parameters to be understood and optimized for a large ZigBee network deployment based on the many-to-one routing scheme. The network nodes are sending data to a central node (concentrator). Section 4 discussed the optimized parameters and explained the reasons for optimization.

The results presented are obtained in a network setting done at the TI San Diego office. The network performance can vary depending on the presence of interferer, jammers, and so forth. Range issues between the nodes and the application report do not ensure network performance if these guidelines are followed as there are several other factors such as interference that could influence the network operation.

6 References

1. ZigBee specification R20: https://zigbeealliance.org/
2. *Z-Stack Developer's Guide* (www.ti.com/z-stack)

7 Revision History

NOTE: Page numbers for previous revisions may differ from page numbers in the current version.

Changes from Revision C (April 2014) to Revision D (September 2020)	Page
• Updated the numbering format for tables, figures and cross-references throughout the document	2
• Updates were made to Section 1	2

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI’s products are provided subject to TI’s Terms of Sale (www.ti.com/legal/termsofsale.html) or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI’s provision of these resources does not expand or otherwise alter TI’s applicable warranties or warranty disclaimers for TI products.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265