Enabling Industry 4.0 Business Models For MSOs Using Wireless Mesh Networks At 60 GHz

A Technical Paper prepared for SCTE•ISBE by

Elliott Hoole
Director, Wireless R&D
Charter Communications
6360 S. Fiddlers Green Circle, Greenwood Village, CO 80111
(720) 536-9424
Elliott.Hoole@charter.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

The term "Industry 4.0" refers to the 4th industrial revolution brought about by digital technologies to dramatically increase operational efficiencies primarily through the use of analytics and automation. Wireless systems are seen as playing a primary role in this revolution to more easily allow for on-demand manufacturing process reconfigurations along with high volume data collection via industrial internet of things (IIoT) networks. The motivation of the project described in this paper is to investigate the performance and feasibility of using multi-hop pseudo-mesh indoor networks in the 60 GHz V-Band for industrial automation applications to support Industry 4.0 use cases. These use cases include the real-time control of multiple robotic apparatuses using a dedicated multi-access edge compute (MEC) system. Another goal of the project is to characterize the deployed wireless connection robustness and their practical limitations in both static and dynamic links. In addition to performace evaluation it is also important to understand the current state of supporting component technologies and the viability of a robust technology ecosystem for deploying and supporting these networks.

The main challenge for the initial phases of this project was the fact that there are no commercially available 60 GHz true mesh products currently available on the market. Mesh connectivity for this project was emulated using multiple point-to-multipoint links. Another challenging aspect was that the 60 GHz Consumer Premises Equipment (CPE) endpoints in our trials are the same physical devices as the access nodes with a different operational configuration which would not be the case in an actual commercial deployment. Lastly there were some use cases on our trials which required connectivity translation to interwork with 60 GHz network, so there may be issues addressing some of the needs of some potential customers depending on their specific objectives and equipment. Many of these issues are being addressed in later phases of this project and will be detailed in a future publication.

However despite the challenges of these initial project phases, it has been found through measurements and practice that networks with 60 GHz technology can meet needs of (Industrial Automation) (IA) use cases today if deployed and managed properly.

2. Industrial Automation System Requirements

There are many specific situations within the Industrial Automation use case umbrella. One expert who is active in the field has stated that a good place to start would be to "duct-tape an iPhone to a milling machine"[1] which would enable the collection of a production asset's operational parameters. A very high value use case that does not require huge bandwidth and low latency is to perform what is known as "finding the hidden factory". This refers to operational inefficiencies which are found through process monitoring, data collection, and targeted analytics to determine the efficiency of the end-to-end production process and identify sections of it that may not be performing as intended. In so doing, surprisingly large amounts of money are saved simply by reducing otherwise unknown waste [2].

The 3rd Generation Partnership Project (3GPP) have identified target requirements for specific IA use cases, and these can be found in TS 22.104 [3]. **Error! Reference source not found.** below is a table showing the 3GPP use cases (columns) with their associated application areas (rows). As seen, not every application area employs every use case. So different industrial customer segments can potentially be addressed with different target service and application packages, although this concept is not explored further in this paper.

Table 1 - Industrial automation application areas and associated use cases

3GPP TR 22.104 Annex A	Motion Control	Control-to-Control	Mobile Control Panels With Safety	Mobile Robots	Remote Access and Maintenance	Augmented Reality	Closed Loop Process Control	Process Monitoring	Plant Asset Management
Factory Automation	X	X		X					
Process Automation				X			X	X	X
HMIs and Production IT			X			X			
Logistics and Warehousing		X		X					X
Monitoring and Maintenance					X				

Each of the use cases listed in Table 1 has a set of system level parameters which should be met in order to provide the expected level of performance. A representative sampling of these requirements is given below in Table 2.

Table 2 - - Selected 3GPP technical requirements for industrial automation use cases

Use Case	Service Avail	Message Size	Transfer Interval	Survival Time	UE Speed	# of UEs	Service Area
Motion Control	5-7 9's	50 bytes	0.5 ms	0.5 ms	≤ 75 kph	≤ 20	50x10x10 m
Mobile Robots	6 9's	40 - 250 bytes	1-50 ms	1-50 ms	≤ 50 kph	≤ 100	< 1 sq km
Process Monitoring	4 9's	20 – 255 bytes	100 - 60k ms	3 x Trans Intv	0	10k – 100k	10x10x0.05 km
Mobile Control Panel	6-8 9's	40 - 250 bytes	4-8 ms	4-8 ms	≤ 8 kph	TBD	50x10x4 m
Process Control	6-8 9's	20	10 ms	0	0	10-20	100x100x50 m
Control-to-Control	6-8 9's	1k bytes	10 ms	10 ms	0	5-10	100x30x10 m
Augmented Reality	3 9's	unspec	< 10 ms	unspec	< 8 kph	unspec	20x20x4 m
Asset Management	4 9's	20 – 255 bytes	'several seconds'	3 x Trans Intv	0	Up to 100k	10x10x0.05 km

In addition to the requirements listed in Table 2, 3GPP TR 22.104 [3] further segments each use case into sub-categories so that there are multiple sets of requirements for each of the use cases listed in the columns of Table 1. For example, the Motion Control use case (UC) has 3 sub-categories which are listed below in Table 3.

Table 3 - Motion control use case sub-categories

Use Case	Service Avail	Message Size	Transfer Interval	Survival Time	UE Speed	# of UEs	Service Area
Motion Control UC #1	5-7 9's	50 bytes	0.5 ms	0.5 ms	≤ 72 kph	≤ 20	50x10x10 m
Motion Control UC #2	6-8 9's	40 bytes	1 ms	1 ms	≤ 72 kph	≤ 50	50x10x10 m
Motion Control UC #3	6-8 9's	20 bytes	2 ms	2 ms	≤ 72 kph	≤ 50	50x10x10 m

The most stringent use case sub-category of all those listed in [3] is Motion Control #1. This is the use case that is typically focused on in marketing literature and used as ***the*** requirements for Industrial Automation, but as can be seen in Table 2 there are other sub-categories of Motion Control that are not as demanding as Case #1 and many other use cases and their sub-categories which are not as demanding as Motion Control #1 but are very useful for Industrial Automation networks and services.

Consider a 100 byte packet sent every 1 msec. This results in a required throughput of 8*100*1000=800 kbits/sec. As can be seen from the previous figures this represents a sufficient amount of bandwidth and transmission interval for every listed use case above. So a system throughput budget of 1 Mbps per Industrial Automation endpoint can be viewed as being representative of what IA systems require.

Another important requirement for IA networks is service availability which is an end-to-end network requirement for delivering specified services. System redundancies, both at the node level and in available paths, can be employed to increase the overall end-to-end availability beyond what a single node/path can deliver.

An illustration of service availability with respect to redundancy is shown below in Figure 1. In this case reliability concerns only the wireless network portion of the end-to-end network. The table embedded in Figure 1 gives the relationship between service availability and reliability when the survival time is equal to the transfer interval. These will be discussed shortly. More information concerning service availability and its relationship with reliability can be found in 3GPP TS 22.241[4].

In order to increase the reliability of the wireless network portion of Figure 1 and the overall availability of the end-to-end service, a wireless mesh network can be employed. This type of network provides an endpoint with multiple paths through which to send data so as to eliminate the sole dependency on any particular node for communication.

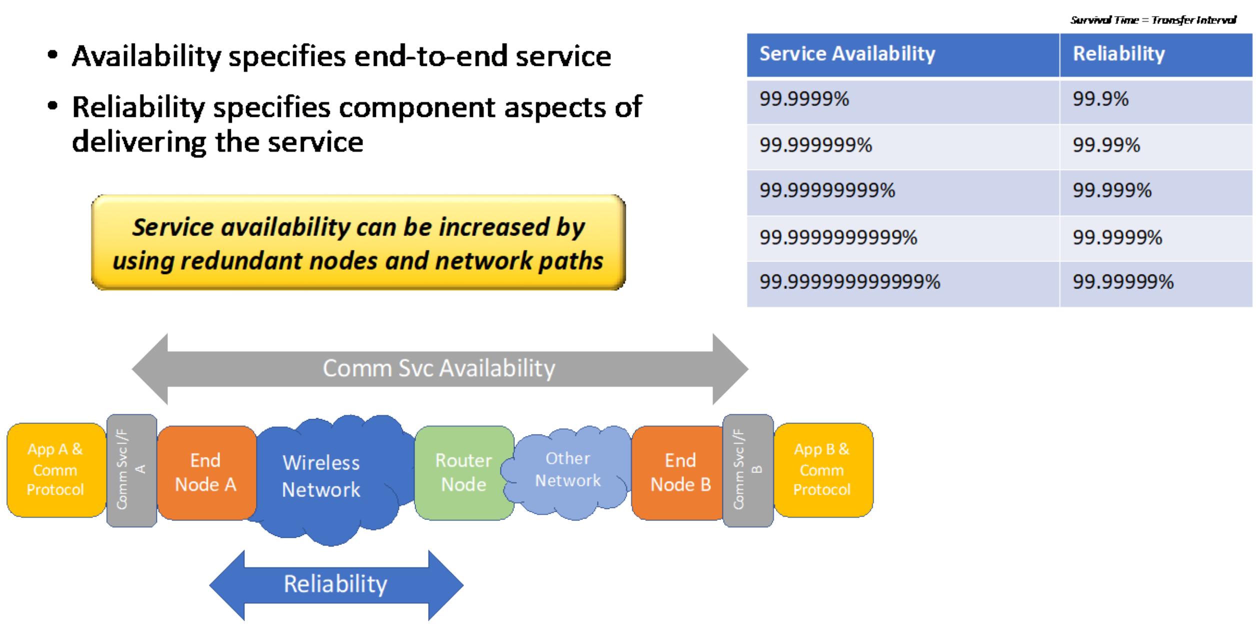

Service Availability	Reliability
99.9999%	99.9%
99.999999%	99.99%
99.99999999%	99.999%
99.9999999999%	99.9999%
99.999999999999%	99.99999%

Figure 1 – The relationship between service availability and reliability

As seen in Figure 1, service availability and reliability are related by parameters called transfer interval and survival time. These are system delays that affect an endpoint's response time to a command from a control system. These delays are illustrated below in Figure 2. In this illustration the transfer interval from the MEC to the endpoint is the same as the one from the endpoint to the MEC since they traverse the same path. This is not always the case and must be taken into account for round trip response time calculations. The survival time is the amount of time it takes an endpoint to respond to a command and issue a status response back to the control system.

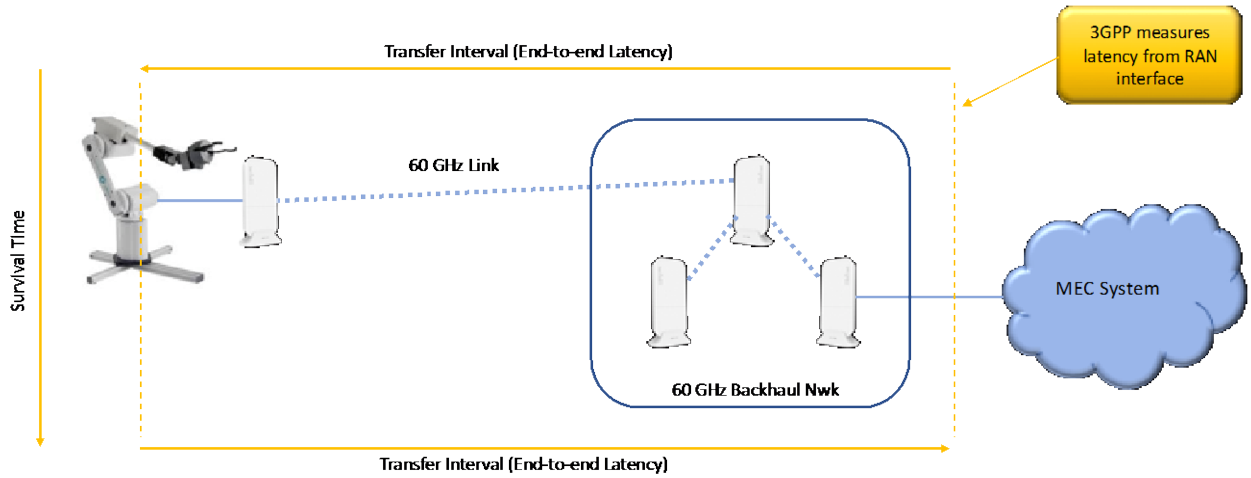

Figure 2 - Industrial Automation network delays

The IA network requirements discussed in this section illustrate the level of performance expected to be able to perform the target IA use cases. As seen in Table 2 not every use case requires a very high degree of performance and many of them can be enabled with equipment that is widely available.

3. 60 GHz Indoor Deployment Aspects

3.1. Indoor Scattering

Historically 60 GHz wireless transmission has been regarded as being problematic due to the oxygen absorption characteristic in the radio frequency band. This phenomenon really only affects very long links as the effect is described in dB/km. However for shorter links (< 100m) this effect is negligible.

60 GHz radio signals do not pass through materials as well as radio signals do in other bands. As such there is a tendency for a large amount of environmental scattering which makes antenna directionality indoors much less critical or even important. For indoor 60 GHz applications this means that line-of-sight and precise beam pointing deployment aspects are not as critical as in other bands where the radio signals scatter less.

3.2. Antennas

The typical antenna devices in 60 GHz products is found to be a patch array that is used in conjuction with beamforming/beam-switching to provide higher gain for longer links and alleviate the need for antenna alignment during installation. The standard antenna array module is similar to the one shown below in Figure 3.

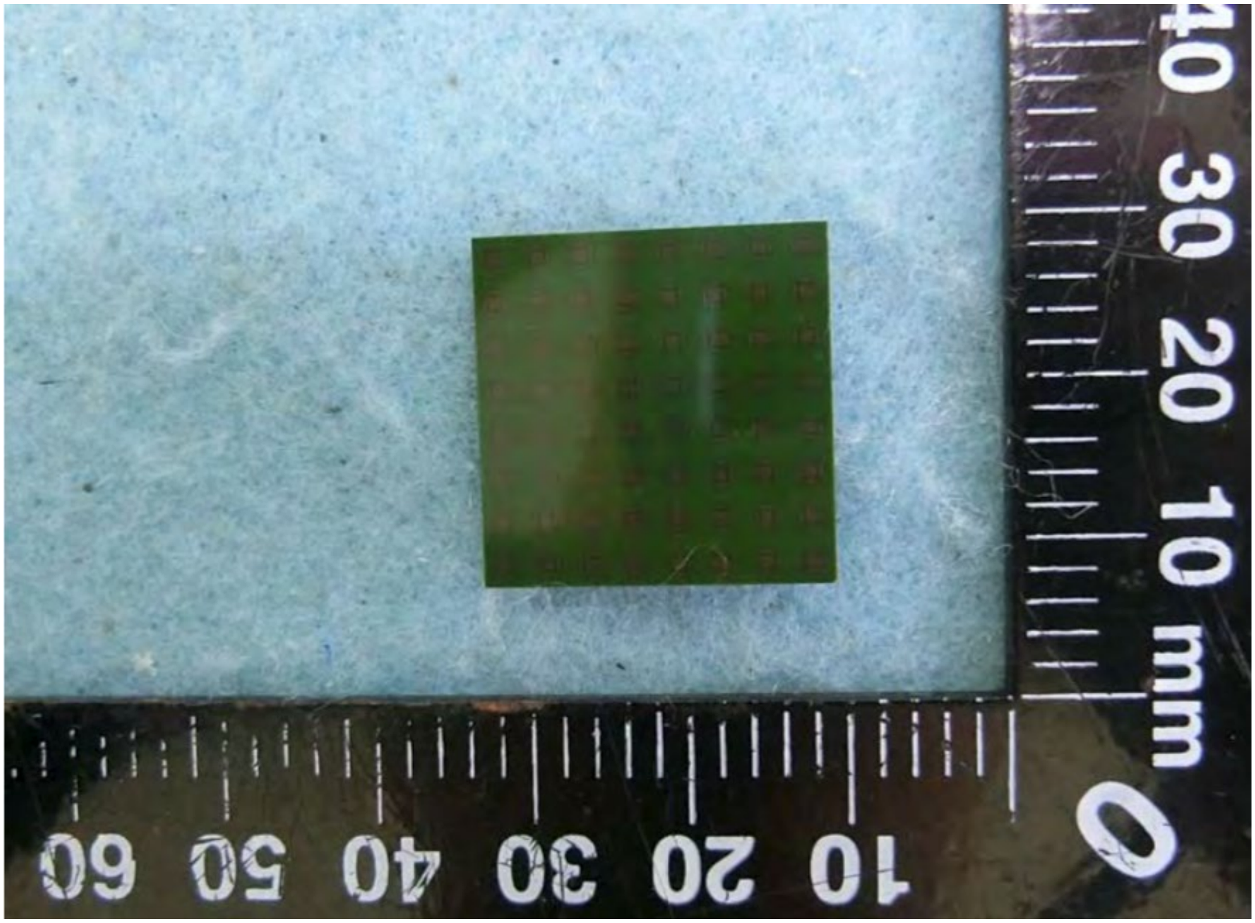

Figure 3 - Typical 60 GHz MIMO antenna array module

These standard multiple-input and multiple-output (MIMO) patch array modules have the typical specifications given below in Table 4.

Table 4 - MIMO antenna module parameters

MIMO Module Parameter	Value
Number of elements	64
Coverage angle	90° horizontal / 40° vertical
Max Effective Isostropic Radiated Power (EIRP)	36 dBm

In a representative warehouse indoor environment we have seen through field testing, which will be further described in Section 4.2, that a hemi-spherical pattern from the same AP is better overall in terms of coverage and robustness. This was done by substituting a hemispherical antenna module for the standard MIMO array module and then testing both configurations in the same deployment locations. The hemispherical antenna has the specifications given below in Table 5.

Table 5 - Hemi-spherical antenna module parameters

Hemi-Spherical Module Parameter	Value
Number of elements	32
Coverage angle	180° horizontal / 180° vertical
Max EIRP	30 dBm

For product commercialization, the radio units with hemi-spherical antenna modules can simply be variants of existing products with an antenna module substitution.

3.3. Diversity and Coverage Robustness

In order to meet the more stringent requirements for IA use cases (e.g. motion control), deployments should make use of overlapping coverages from more than one access point (AP). During the network planning phase, the coverage areas from different APs should be planned with substantial overlap to allow for secondary or alternate connectivity options for each endpoint.

One possible system configuration could be to install multiple access point radio units feeding a single baseband unit which would then employ various diversity combining techniques. This concept is being investigated in an ongoing phase of this project.

3.4. Fast Reassociations

IEEE specification 802.11r-2008 [5] describes a fast transition technique for a client station (STA) to move its radio connection from one AP to another in a more expedited manner without having to undergo the full authentication procedure [7]. Doing so can reduce the reassociation time to perhaps 50 msec and is recommended to be done within enterprise networks using WPA2 Enterprise security. This technique should be utilized in IA systems to minimize the impact of reassociations on system performance.

It may be possible to make use of an edge compute system to help facilitate these transitions. This concept is being investigated in an ongoing phase of this project.

3.5. Edge Computing System

To truly enable the most meaningful IA use cases requires an edge compute system. A computing system is required to provide control and computing functionality for automated industrial endpoints and tasks. Having this computing system collocated with the endpoints minimizes latency and enables certain use cases that could not be enabled if the system was located further back in the network.

4. Performance Evaluations

4.1. First Indoor Trial

In order to test and evalutate the suitability of 60 GHz for use in Industrial Automation networks, we constructed an indoor network consisting of 10 nodes that were installed around the floor of a typical office building environment. Another 4 units were nomadic meaning their location was not permanent and could range throughout the coverage area. All of the radio nodes used in this trial were from the same vendor, identical, and based on a well known 802.11ad chipset. Three of the units were configured to be point-to-multipoint access points, and the rest were configured as client devices of those access points. Along with the radio nodes, a server cluster was installed in the telco room on the same floor to serve as a MEC system. The layout of the installed network is shown below in Figure 4.

Figure 4 - 60 GHz indoor test network

Since the deployed radio units do not provide a true mesh operation capability, we needed to emulate one by using multiple point-to-multipoint connections which could be configured and controlled through the switching infrastructure. To provide characteristic industrial network loading we used multiple Raspberry Pi devices to emulate industrial endpoints with 50 byte payloads at 50 packets per second. A representative test network configuration is shown below in Figure 5.

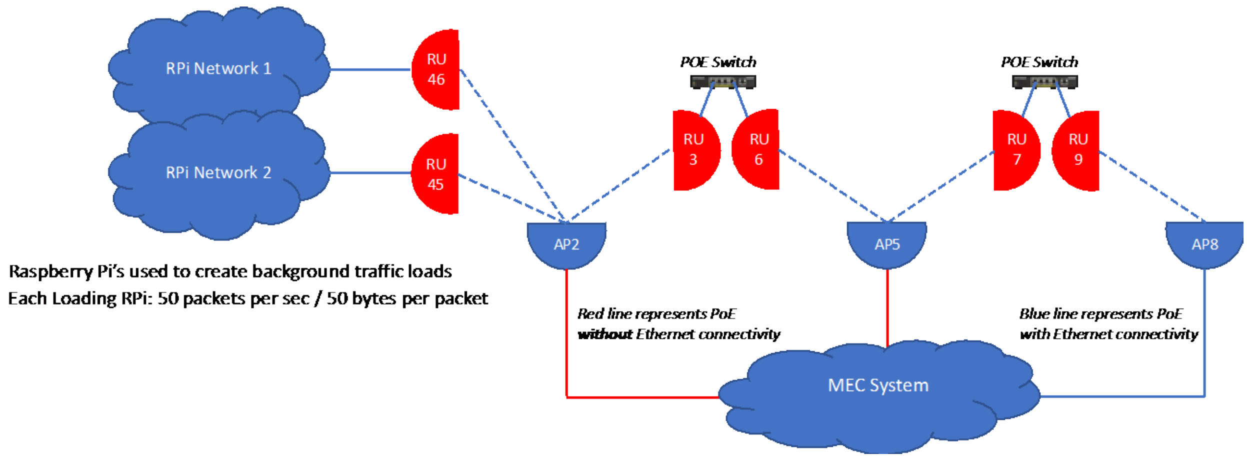

Figure 5 - Example test network configuration

In the example test network of Figure 5 the blue half-circles represent APs and the red half-circles represent remote units (RUs). The only egress point for packets from the MEC system is the blue line to AP8. The red lines to AP5 and AP2 have PoE to power the access points but the Ethernet connection on each switch port has been disabled. So the packets from the MEC system must traverse the path from AP8 to RU46 through 5 airlink hops. This effectively emulates a situation where the traffic from an endpoint would traverse multiple intermediate mesh nodes before arriving at the intended destination.

Several network configurations were employed and round trip delay measurements were collected for different scenarios with different types of traffic loading. In the cases labeled “Case1”, the device under test (DUT) was connected to the same RU as the Raspberry Pis that were emulating industrial endpoints and a varying number of Raspberry Pis were used. In the cases labeled “Case2”, the DUT was connected to a different RU than the Raspberry Pis but both Rus were connected to the same AP. Lastly bidirectional iperf traffic was used in differing amounts to create network loading. The results are shown in Figures 6-8.

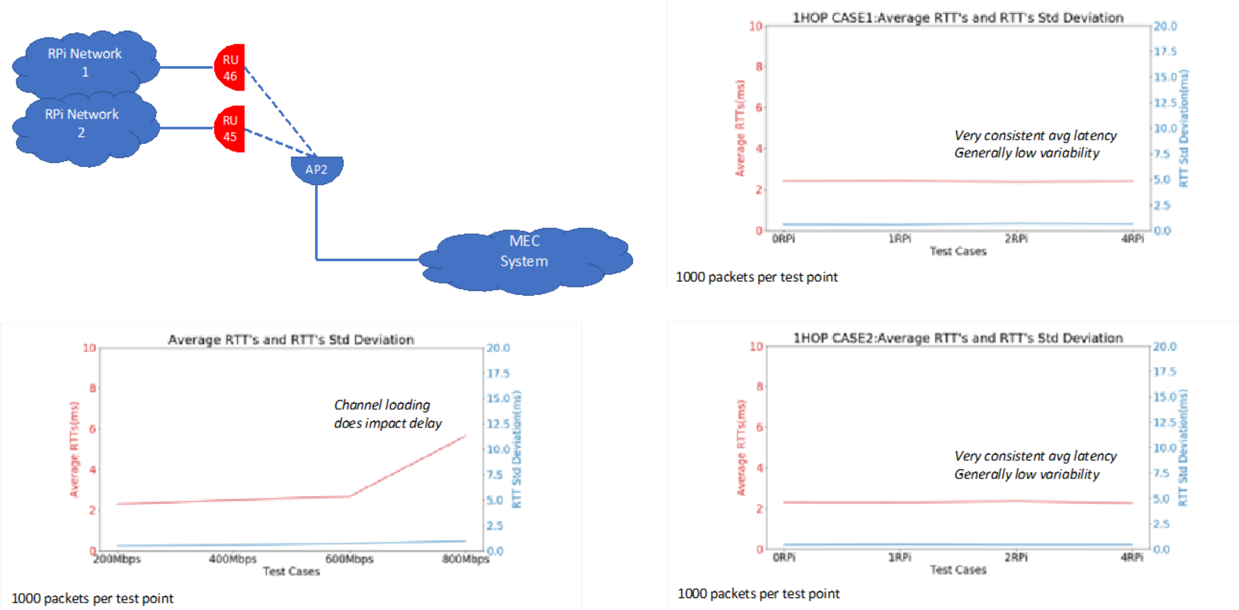

Figure 6 - Round trip delay statistics for the 1-hop test network configuration

As seen in Figure 6 the two sets of data from the Raspberry Pi cases are very similar. There does not seem to be enough loading from the endpoint emulators to really affect the outcomes significantly. This means that in a scenario with a fairly low number of industrial endpoints per access point with a direct connection to the MEC system the delay performance should be very consistent. This is further corroborated by the 3rd case with the iperf traffic. The delay results are very similar to the Raspberry Pi cases until the background traffic gets very high at which point the delay jumps significantly. It has been observed with certain chipsets that airlink errors which require retransmissions can stall the data queue and create a buffering delay which impacts the transmission time. This effect appears to be the cause of the increased delay for the last iperf case data point. But with lower amounts of background traffic the results track those of the Raspberry Pi cases.

In the test network configuration with 3 wireless hops as shown in Figure 7, the Raspberry Pi case results are largely just an overall increase from the 1-hop case due to the increased amount of network elements. However the last data point in Case 2 seems to show a slight increase from the buffering delay effect mentioned above. The iperf background data results show a more dramatic variation, likely due to the increased number of airlink hops presenting more opportunities for retransmissions and thus a larger impact from the buffering delays.

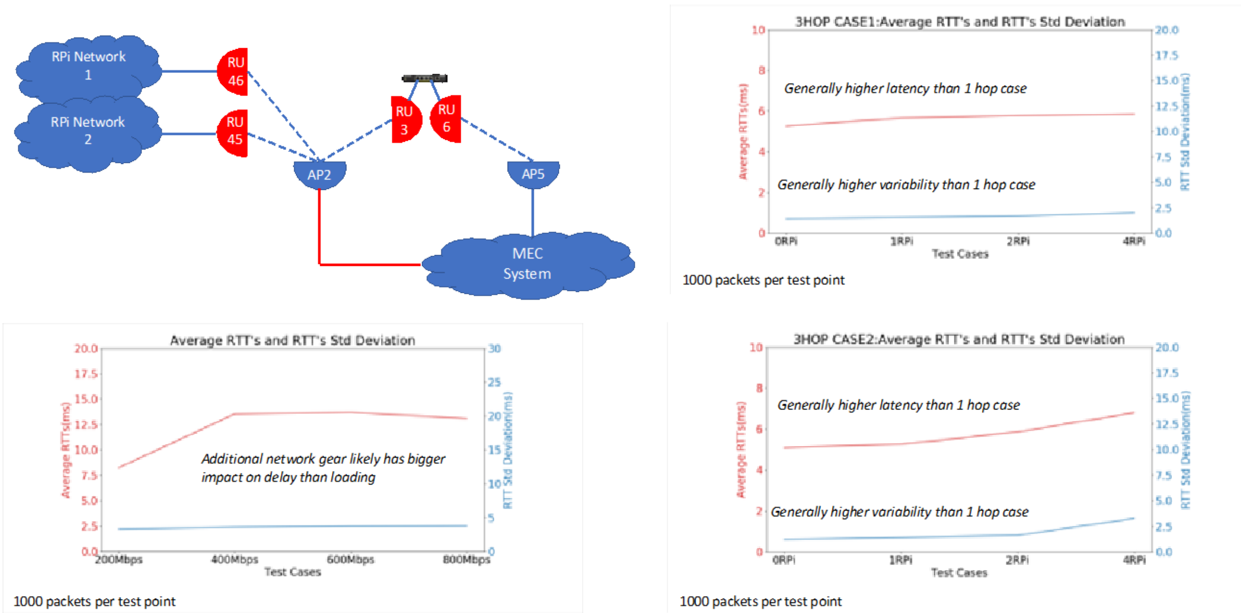

Figure 7 - Round trip delay statistics for the 3-hop test network configuration

For the test network configuration with 5 airlink hops, the results are even more varied as is expected. The Raspberry Pi cases are still generally consistent with more overall delay due to the increased amount of network. Again Case 2 shows more variation than Case 1 with some buffering delays being a likely culprit for the increased delay statistics during the test run.

Perhaps the most interesting results come from the iperf background data case. It is seen in that the test point with 800 Mbps background traffic, the mean delay is significantly lower than that of the test point with 400 Mbps background traffic. I believe this further exemplifies the buffering delay effect as there are now even more opportunities for retransmission events and the results reflect that.

A new aspect of 802.11ay that could be very beneficial in regulating the flow of traffic through the network is the time division duplex (TDD) scheduling feature. 802.11ad uses the typical listen-before-talk (LBT) mechanism which can lead to congestion and collisions is various points of the network, particularly the APs. With TDD scheduling each endpoint has a timeslot on the airlink reserved for its use along with an assigned priority. This new mechanism and its possible benefits will be investigated in an upcoming phase of this project.

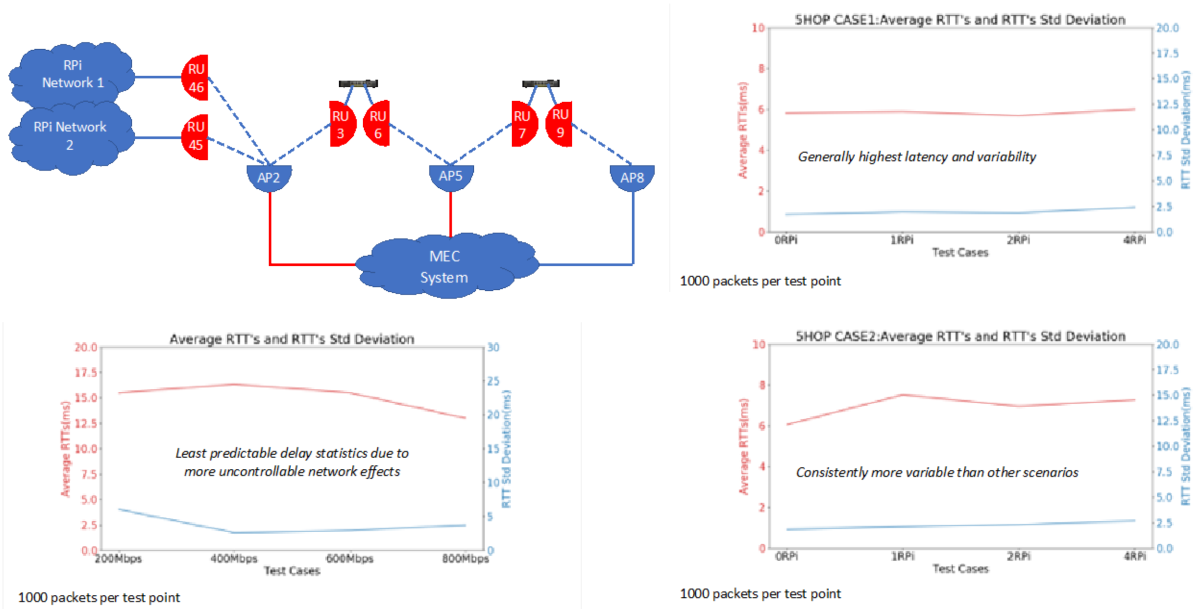

Figure 8 - Round trip delay statistics for the 5-hop test network configuration

Other considerations that have been observed to affect delay statistics are the sizes of the packets and the number of packets per second required for an endpoint as seen below in Figure 9 and Figure 10. The data in both of these figures was captured with the network in the 5-hop configuration along with 500 Mbps of background traffic.

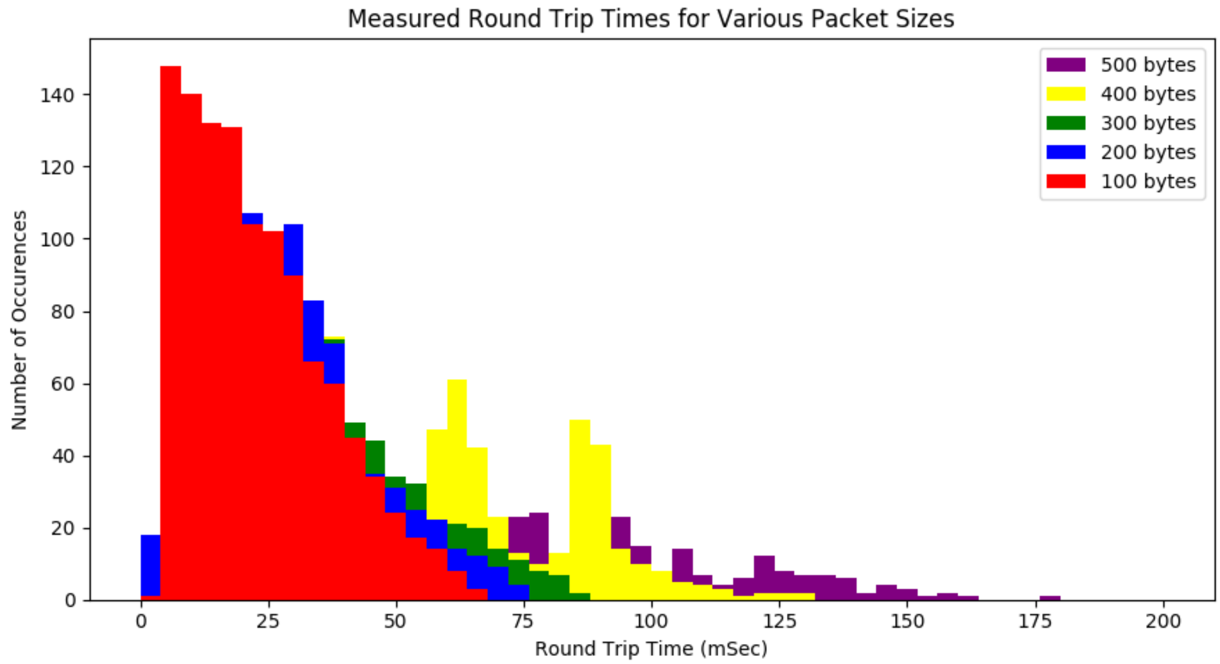

Figure 9 - Measured round trip times for various packet sizes

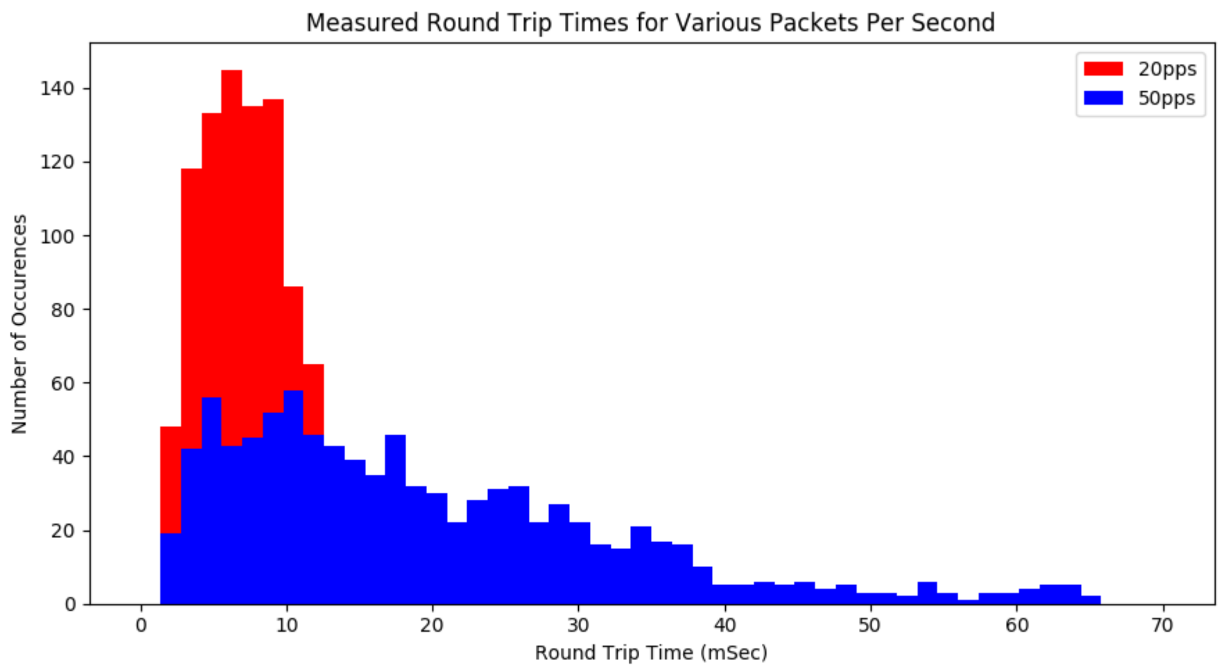

Figure 10 - Measured round trip time with 20 and 50 packets per second

The data runs for both figures was comprised of 1000 packets. The packet rate for Figure 9 was 10 packets per second. The packet size for Figure 10 was 50 bytes. As can be seen from the figures, packet size and packet frequency can both have significant effects on the measured round trip times, especially in the multi-hop network configurations. These parameters need to be carefully considered when deploying networks for industrial automation.

4.2. Second Indoor Trial

In another trial at a different location, 60 GHz indoor nodes from a different vendor were modified and deployed with different antenna modules to evaluate their suitability for IA applications. Hemispherical antennas were determined to be the best choice at both APs and clients. The use of hemispherical antenna modules are recommended for indoor IA applications which can provide performance improvements compared to the multiple input multiple output (MIMO) antenna modules.

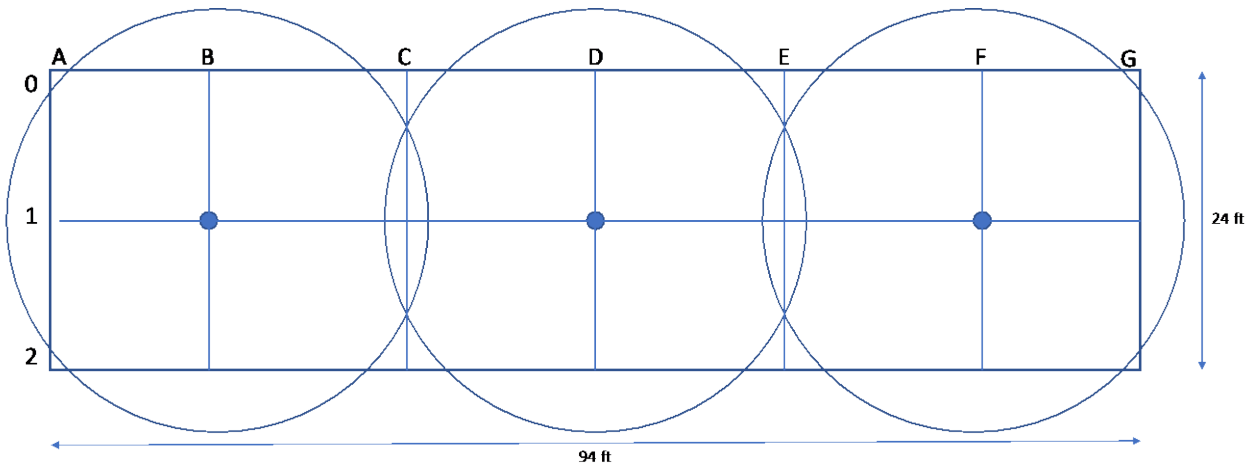

Table 6 – Measured RSSI (dBm)

Row	A	B	C	D	E	F	G
0	-55	-59	-57	-57	-60	-55	-59
1	-55	-55	-61	-55	-59	-55	-58
2	-59	-57	-58	-54	-60	-55	-57

Table 7 - Measured round trip delay (microseconds)

Row	A	B	C	D	E	F	G
0	282	333	397	380	345	339	311
1	390	363	351	468	452	471	532
2	410	384	354	385	418	369	380

Delay measurements seen in this second trial are much lower and more consistent, so equipment selection is critical for IA network performance. Even though vendor #2 uses the same chipset as vendor #1, the performance characteristics, especially delay, are substantially different due to the specific equipment implementations. As with any network application, equipment selection should be carefully considered for IA networks to achieve the required level of performance.

Use cases that require endpoint mobility along with real-time and/or constant control cannot readily achieve the required QoS with standard 802.11ad equipment. Of course there are many factors that govern this situation, but the amount of time required for endpoint reassociations between APs can easily exceed the required control packet interval resulting in performance issues for the more demanding use cases such as motion control. Further enhancements are needed for mobility situations and robustness against connection issues due to objects moving within the environment if real-time control is required. Improvements for these scenarios are being investigated as an ongoing phase of this project.

5. Use Cases and Demonstrations

In order to assess and characterize the performace of 60 GHz networks for industrial automation applications, several demonstrations were developed each employing several IA use cases.

5.1. Demonstration System #1

The first demonstration involved a user-controlled mobile robot car that is piloted through a series of gates in a timed course. A live video stream from the robot car is displayed on a computer monitor and the user is told by the MEC system which gate to navigate to next. The position of a game controller joystick is read by the system and converted to servo motor parameters which are sent to the robot car to control its speed and direction. The user is then able to pilot the car using the displayed video and the joystick control. The elapsed time is also displayed for the user. Infrared (IR) sensors in each gate detect the presence of the robot car and the user is then given the next gate in the course on the screen by the system. If the wrong gate is detected by the system, this notification is displayed on the screen. The course continues until the last gate is reached at which point the elapsed time is recorded.

5.1.1. Use Cases Demonstrated

The industrial automation use cases demonstrated in Demo #1 are listed below in Table 8.

Table 8 - Demo #1 use cases

Use Case	Description
Motion Control	Robot servos updated 20 times per second according to the joystick position
Control-to-Control	Live video stream (4 Mbps) from robot to desk client
Closed Loop Process Control	Next gate indication and elapsed time display
Process Monitoring	Gate status messages including wrong gate detection
Mobile Control Panel	Real time progress and status displayed on desktop

5.1.2. Demonstration System Architecture

The system architecture of the demonstration system is shown below in Figure 11.

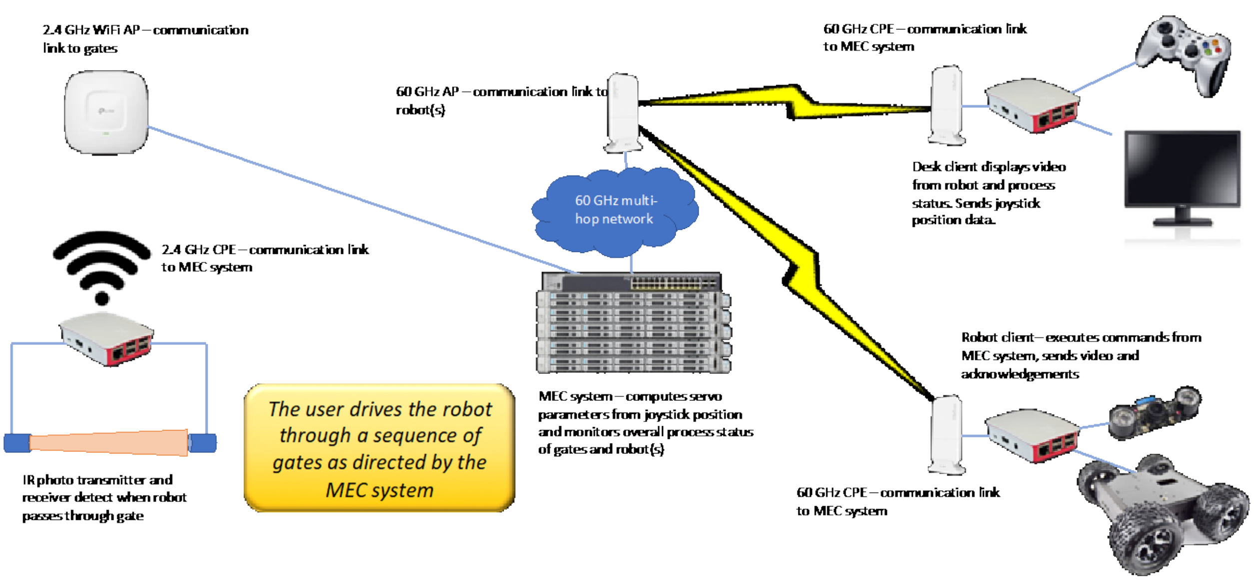

Figure 11 - Demo #1 system architecture

This demonstration system features a mobile robot which is controlled by a game controller joystick. The mobile robot includes a camera which streams video back to the user. The user then moves the joystick to control the speed and direction of the robot. With this interface the user directs the car through a series of gates which is monitored by the MEC system. The elapsed time is displayed on the output screen, and the overall goal is to have the robot go through the correct sequence of gates in the shortest amount of time.

5.1.3. Demonstration System Data Messaging Architecture

The sequence of messages that flow between the various functional entities of Demo #1 is shown below in Figure 12.

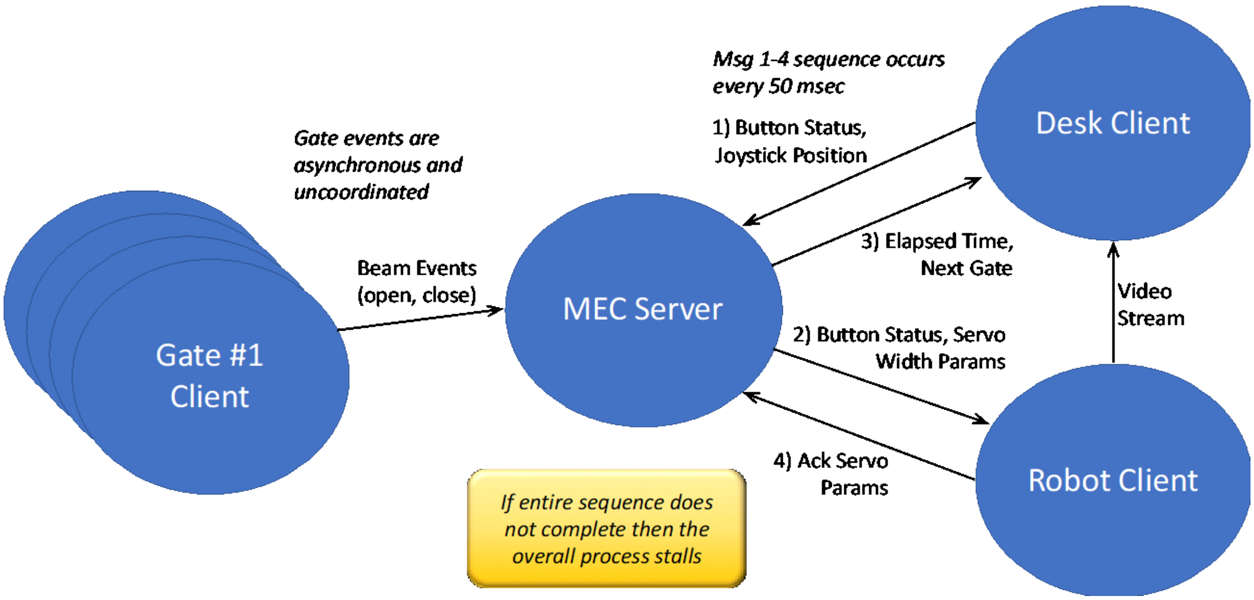

Figure 12 - Demo #1 data messaging architecture

The data architecture and message flows depicted in Figure 12 is implemented rather poorly on purpose. If any of the messages in the continuous sequence 1-4 are lost then the entire process will stall. Admittedly this is poor software design, however for our puposes this will readily illustrate when a packet loss creates a critical problem for a use case. During the normal course of the demonstration this situation never occurred. It can occur if the robot car is driven beyond the coverage area of the access point. This would be considered a mobility situation and as stated above, mobility enhancements are being investigated in an ongoing phase of this project.

5.2. Demonstration System #2

A second system was developed to demonstrate a fully autonomous scenario involving 2 robots performing interactive tasks as directed by the MEC system. The demo consists of a robotic arm loading a number of widgets into bins mounted on a mobile robot. The MEC system directs the mobile robot to go to the first loading position and monitors its progress during the trip. Once the mobile robot is in the correct position the MEC system directs the robot arm to load widgets into the first bin mounted on the mobile robot. The robot arm movements consist of a series of poses that are directed and monitored by the MEC system and require quite a large number of messages to perform. Once the first set of widgets is loaded, the MEC system directs the mobile robot to turn around so a second set of widgets can be loaded into a second bin. Once the mobile robot is in the new position the robot arm performs a similar series of actions as before to load another set of widgets into the other bin. Once the second set of widgets is loaded the mobile robot is sent off to its next location and its progress is monitored.

5.2.1. Use Cases Demonstrated

The industrial automation use cases demonstrated in Demo #2 are listed below in Table 9.

Table 9 - Demo #2 use cases

Use Case	Description
Mobile Robots	Mobile robot with 3 pre-programmed waypoints and position status checks every 100 msec. ~1100 messages exchanged during the demo.
Motion Control	Robot arm with 58 separate poses controlled by the MEC system. ~144,500 messages exchanged during the demo.
Process Monitoring	5 interlinked segments monitored and coordinated by the MEC system during the demo.

5.2.2. Demonstration System Architecture

The system architecture of the second demonstration system is shown below in Figure 13.

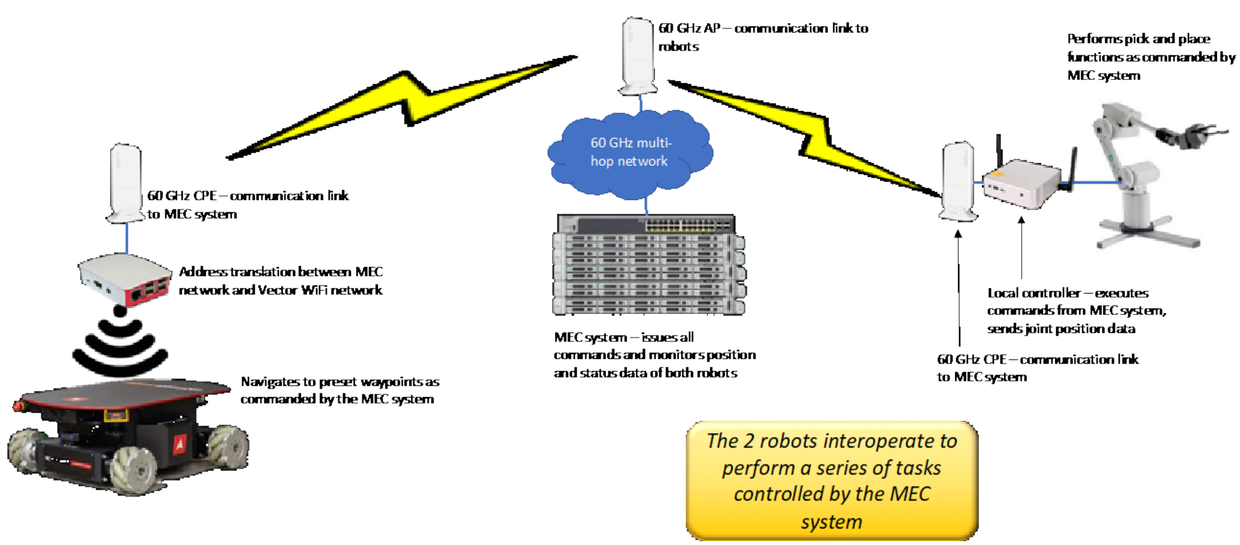

Figure 13 - Demo #2 system architecture

In this demonstration everything is controlled and monitored by the MEC system. The robots are programmed to perform certain tasks which requires both to operate appropriately in order to complete the overall job successfully. The overall job is for the mobile robot to collect a number of widgets which are loaded onto it by the robot arm. So there are five distinct job phases that are controlled and monitored by the MEC system.

6. Economics and Ecosystem

In this section I will define and illustrate an example industrial space for comparative purposes (50 m x 20 m). The following sub-sections give the expected equipment for deploying networks based on 60 GHz nodes vs a 5G system.

6.1. Example Network Comparison

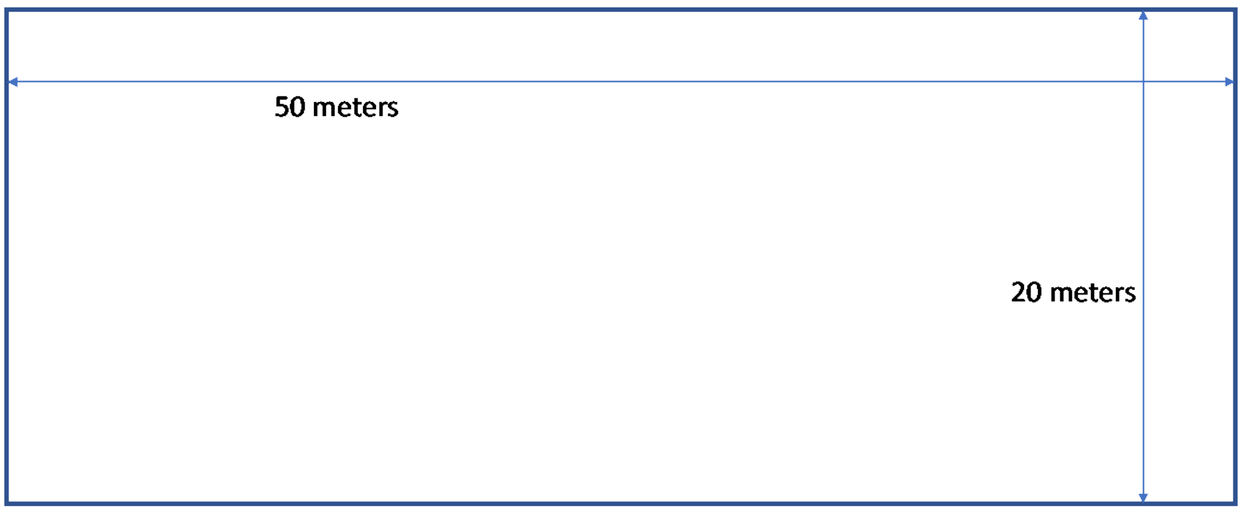

Figure 14 - Example industrial space

6.1.1. 60 GHz Example Network Costs

Based on the network deployed in the 2nd indoor trial described above, it is reasonable to assume that 2 rows of 4 access points would be needed to cover the example space defined above. Although fewer APs might serve to provide adequate coverage, this number will provide the desired overlapping coverages to allow endpoints to reassociate as required to meet the necessary level of performance.

The costs of capital expenses (CapEx) for a 60 GHz network to address the example space are listed below in Table 10. These include both equipment and installation costs.

Table 10 - Example 60 GHz network CapEx costs

Element	Unit Cost	Quantity	Subtotal
Access Points	$500	8	$4000
AP Installation	$250	8	$2000
STA Endpoints	$200	20	$4000
STA Installaion	$50	20	$1000
Edge Compute System	$5000	1	$5000
Total 60 GHz (CapEx)			$16000

The example 60 GHz network would not require any additional components beyond what is already part of the customer's IT infrastructure. Operational expenses for the network equipment (e.g. power) would be covered by the customer. The CapEx costs would be factored into a services offering by the MSO as part of a larger bundle which could include value-add services such as analytics and predictive maintenance. The relatively low network expenses, both capital and operational, could make this approach quite attractive for MSOs and their service integrator partners.

6.1.2. 5G Example System Costs

For the example space the most appropriate 5G radio units would be industrial picocells. Based on the dimensions of the space it is likely that three units would be required to allow sufficient coverage overlap. Assuming the 5G core network (5GC) is located in the multiple-system operator (MSO) network, an on-premises user plane function (UPF) unit is required to handle local Ethernet traffic between the MEC and the picocells.

Table 11 - Example network 5G CapEx costs

Element	Unit Cost	Quantity	Subtotal
Industrial Picocells	$5000	3	$15000
Picocell Installation	$250	3	$750
Enterprise UPF Unit	$2000	1	$2000
STA Endpoints	$250	20	$5000
STA Installation	$50	20	$1000
Edge Compute System	$5000	1	$5000
Total 5G (CapEx)			$28750

The 5G system does require connectivity to a 5GC which for this exercise could be located in the MSO network. The MSO would likely factor the operation and maintenance cost of the 5GC into the offered services package which would be an additional monthly operational expense for the customer. If the customer requires or does not wish for their internal traffic to be exposed to the MSO they could be provided with an on-site 5GC to enable a private network for an additional CapEx cost plus it is likely the MSO would charge an operations and maintenance charge for the local 5GC.

Another aspect of the 5G system is that of usable spectrum. Currently 3GPP 5G systems require licensed spectrum over which to operate. The standards for 5G NR in unlicensed spectrum (NR-U) are expected to be ratified in September, 2020, and this will allow for operation of 3GPP 5G systems in unlicensed frequency bands once equipment is commercially available in 2021. However, most 5G networks are non-standalone (NSA) meaning they require a 4G anchor channel in order for the 5G channel to operate. As networks evolve these limitations will go away, but they will be a major deployment consideration for the next few years and should be factored into any service offerings and target markets.

6.2. Ecosystem Comparisons

6.2.1. 60 GHz Equipment Ecosystem

Currently the 60 GHz equipment ecosystem is not large. In 2018 the global millimeter wave technology market was $289.2M with the frequency range 57-86 GHz having the primary revenue generation for the market [10]. This market is primarily focused on outdoor equipment with few offerings for indoor gear. The 802.11-based 60 GHz equipment market is expanding due to Facebook's Terragraph initiative, but that equipment is largely targeted for outdoor access and backhaul. Indoor is becoming of more interest, but it will take time for new players and offerings to come to market.

The largest 802.11ad chipset supplier is Qualcomm, but there are several others as well. Most notably Sivers and Peraso supply 802.11ad radio ICs along with Peraso and Blu Wireless who supply 802.11ad baseband ICs. Intel produces 802.11ad ICs for endpoints, but they seem to be used only in Intel modem products. With the advent of 802.11ay it remains to be seen how the chipset and corresponding equipment landscapes may change. Qualcomm is producing an 802.11ay chipset that is being used extensively in the Terragraph equipment, but it is not yet available in mass market quantities. Peraso has also announced an 802.11ay chipset and others will as well as the equipment market develops further.

Much of the currently available equipment has similar specifications due to limited number of chipset suppliers and equipment manufacturers vying for largely the same target markets. Differentiation and competitive advantages seem to come largely from support systems e.g. management, configuration, etc.

Bottom Line: The 60 GHz chipset and equipment markets comprise both large and small players, but the equipment target markets have not taken off so overall there is not a lot of muscle behind the current offerings.

6.2.2. 3GPP 5G Equipment Ecosystem

The 3GPP ecosystem is very robust, but sadly the 3GPP ecosystem is no longer as diverse as it once was due to many consolidations amongst vendor companies. Last year the 5G global infrastructure market was nearly $1B and by 2026 it is predicted to be over $50B [11]. The market is supported by heavyweight equipment and chipset vendors that can be considered telecom institutions. The variety of offerings is staggering with a range of radios that run a gamut from ones that suppport very large cells to ones used in very small cells in a wide variety of frequency bands and combinations of frequency bands. The radio platforms can be based on anything from standard chipsets to purpose-built designs to commercial off-the-shelf (COTS) hardware.

The 5G standards targeting industrial automation operation are still expected to be fully ratified in September of 2020. This operation is called ultra-reliable low-latency communication (URLLC), and equipment that supports this feature is expected to be available near the end of 2021 at the earliest.

The 3GPP 5G standardization activities are very active and highly dynamic with 370 total 3GPP members and 1267 5G related standards in Release 16. The 5G promotional activities as well are highly spirited with many companies touting the virtues of the new technology in creative and interesting ways.

Bottom Line: The 3GPP 5G chipset and equipment markets are dominated by heavyweight telecom giants offering a wide selection of gear for many target markets. This includes Industrial Automation with the new URLLC feature once equipment is available late next year.

7. Conclusion

60 GHz radio technology can provide robust connectivity for indoor environments. Line of sight is not necessarily required due to the tendency for the radio signals to undergo a considerable amount of scattering from objects and surfaces in the environment. As such, precise antenna orientation does not seem to be as critical for the target environments. There is naturally some variability in the radio signals, but generally the radio links have been seen to be very stable and provide robust connectivity. The emerging 802.11ay standard will provide true mesh capabilities which will further enhance performance and robustness of the target scenarios compared to the equipment and network configurations that were used for this study so far.

In general the observed performance of the tested representative network configurations can support 3GPP use cases with the sole exception of the most stringent motion control case. Other sub-categories of motion control require 3 and 6 msec round trip performance which can readily be achieved with proper network planning and equipment selection. For these use cases networks should be limited to the 1-hop configuration to ensure that the required performance is achieved.

Multiple ingress/egress points are critical for fault tolerance and for system performance as well. Mesh is best suited for this, but a nearly equivalent situation can be achieved with multiple Point-to_Multi-Point (PtMP) deployments provided the clients can reassociate as required to meet the system performance criteria. IA networks should be designed for service delivery aspects, but this can lead to complications if reconfigurations are needed. Packet sizes and repetition rates are key parameters for IA network designs.

As observed in our testing some equipment can exhibit uncontrolled behaviors. This may limit the addressable target applications in the short-term, but with new chipsets and equipment (e.g. 802.11ay) these limitations are likely to disappear. Another aspect of 802.11ay that could be very beneficial in this regard is the TDD scheduling feature which should regulate the traffic flows through the network much more than the listen-before-talk behavior of 802.11ad.

In conclusion, 60 GHz networks can deliver < 2 msec round-trip time (RTT) latency and meet 3GPP use case performance requirements today with proper equipment selection and network design for substantially less total cost of ownership (TCO) and will offer better capabilities for addressing these use cases in the near future.

Acknowledgements: The author would like to thank Pooja Shankar for conducting much of the testing done in the first indoor trial.

Abbreviations

AP	Access Point
DUT	Device Under Test
GHz	GigaHertz
HMI	Human-Machine Interface
Hz	Hertz
IA	Industrial Automation
IT	Information Technology
LBT	Listen Before Talk
Mbps	Megabits per second
MEC	Multi-access Edge Computing
MIMO	Multiple Input Multiple Output
msec	millisecond
MSO	Multiple-System Operator
PtMP	Point to Multi-Point
RTT	Round Trip Time
RU	Remote Unit
TDD	Time Domain Duplexing
UE	User Equipment
UPF	User Plane Function
URLLC	Ultra-Reliable Low-Latency Communication
3GPP	3rd Generation Partnership Project
5GC	5G Core

Bibliography & References

[1] Instructor comments from *Implementing Industry 4.0: Leading Change in Manufacturing and Operations*, MIT Sloan Executive Education, July, 2019

[2] Course materials from *Implementing Industry 4.0: Leading Change in Manufacturing and Operations*, MIT Sloan Executive Education, July, 2019

[3] 3GPP TS 22.104: *Service requiremens for cyber-physical control applications in vertical domains*, Annex A

[4] 3GPP TS 22.261: *Service requiremens for the 5G system*, Annex C

[5] IEEE 802.11r-2008 has been folded in as Section 13 of IEEE 802.11-2016

[6] *The Digital Shopfloor - Industrial Automation in the Industry 4.0 Era: Performance Analysis and Applications*, John Soldatos et al (editors), River Publishers, May, 2019

[7] Network Computing, *WiFi Fast Roaming, Simplified*, https://www.networkcomputing.com/wireless-infrastructure/wifi-fast-roaming-simplified, retrieved August 6th, 2020

[8] Chen et al, *Millimeter-Wave Fixed Wireless Access Using IEEE 802.11ay*, IEEE Communications Magazine, Vol. 57, Issue 12, December 2019

[9] RCR Wireless News, *Bringing 5G NR to unlicensed spectrum*, https://www.rcrwireless.com/20181102/5g/5g-nr-unlicensed-spectrum, retrieved August 13th, 2020

[10] Research and Markets, *Global Millimeter Wave Technology Market Size, Market Share, Application Analysis, Regional Outlook, Growth Trends, Key Players, Competitive Strategies and Forecasts, 2019 To 2027*, March, 2020

[11] Fortune Business Insights, *5G Infrastructure Market Size, Share and Industry Analysis By Component (Fibers, Cables, Antenna, Transceiver, Wireless Backhaul, Modem, Router), By Communication Infrastructure (Small Cell, Macro Cell, Radio Access Network (RAN), Distributed Antenna System (DAS)), and Regional Forecast 2019-2026*, July, 2019