LoRaWAN™ CAPACITY TRIAL IN DENSE URBAN ENVIRONMENT

By Ross Gilson (machineQ) and
Michael Grudsky (Semtech)

CONTENTS

1. INTRODUCTION

In December 2017, machineQ and Semtech executed a joint study designed to emulate a full-scale dense network based on these deployment techniques. In the study, multiple sensors hit multiple gateways. The study assessed the value of a dense network by reassuring packet success rate, interference, gateway density, and evaluated performance and quality of service potential in a dense urban LoRaWAN™ network environment. The following paper summarizes the outcomes of this joint study.

2. OBJECTIVES & ASSUMPTIONS

The primary objective of the study was to evaluate the ability of an enterprise-grade LoRaWAN™ network to handle high amounts of sensor-generated traffic in a dense urban deployment model. The intent was to mimic a full-scale dense network deployment to determine the capacity of the network. The following attributes were tested in a real-world setting:

1. Assess whether an average packet success rate of 90% or greater is achievable when sending each packet once, without acknowledgment.
2. Understand LoRa® high-capacity features and prove LoRa's ability to deliver multiple overlapping packets successfully.
3. Understand gateway diversity and how each device communicates with multiple gateways when transmitting data.
4. Stress test cloud infrastructure to verify its readiness for high-volume commercial deployments.
5. Test network server scalability.
6. Refine existing studies on the effect of environment on signal propagation, packet success rate, and gateway diversity.
7. Evaluate the impact of other capacity inhibitors — multipath fading, radio interference and backhaul throughput.

Dense Urban Deployment Environment

In order to complete the goals of the study, the team defined the parameters of a full-scale dense urban deployment using populous data and scalability metrics of the network. The team identified a city area of 1/4 square mile as a reasonable test area for study requirements.

The team focused on identifying traffic volume objectives using the following assumptions:

- There are approximately 100 residences per city block and 1,000 residences in the deployment area of 1/4 square mile.
- Each site in a full-scale deployment would have approximately 10 LoRa-enabled IoT devices.
- Each device would send upstream packets on average once every hour.

10 devices x 1,000 residences x 1 packets/hr x 24 hrs = 240,000 packets/day

These parameters would result in an average of 240,000 packets per day from the equation shown above.

This study opted to round up and evaluate 250,000 packets per day as the normal case to ensure conservative results.

3. METHODOLOGY OVERVIEW

Rather than install ten devices across 1,000 residences, the study mimicked the traffic achieved by 10,000 sensors by deploying ten gateways and 100 sensors transmitting at such packet rates as to achieve the same traffic as 10,000 sensors. The study tested the ability to send and process packets in three phases. 250,000 packets per day as per the normal case predicted in a dense urban deployment environment in Phase 1 (see Table 1 below), 500,000 packets per day in Phase 2 and 1,000,000 packets per day in Phase 3 to test the network under extreme traffic conditions. Phase 3 of the trial — 1,000,000 packets per day — modeled 10,000 sensors in a 1/4 square mile city environment, sending a packet to the cloud on average every twelve minutes. This phase of the study created 25% network load conditions, a very challenging case wherein active traffic was present on each of the eight wireless channels 25% of the time. The packets were sent in randomized intervals to continue to mimic a real-world environment. Table 1 below shows the number of packets transmitted per sensor in each phase of the trial to mimic the 10,000 sensors delivering packets at different rates.

A total of ten gateways was installed during the trial. The study used 8-channel indoor gateways as the primary hardware in dense network deployment due to their economic scalability. A smaller number of channels also results in increased probability of multiple packets occupying the same channel at the same time. This ensured that the study would experience packet overlap.

The machineQ team dedicated considerable time and effort to procure locations and equipment and to make data available for analysis. Semtech helped generate custom sensor configuration for the study and provided theoretical support on LoRa® capabilities.

	Target Daily Packet Volume	Number of Sensors Used	Study Packet Rate (pkt/hr)
Phase 1	250,000	100	104
Phase 2	500,000	100	209
Phase 3	1,000,000	100	417

Table 1. Number of packets transmitted per sensor in the study

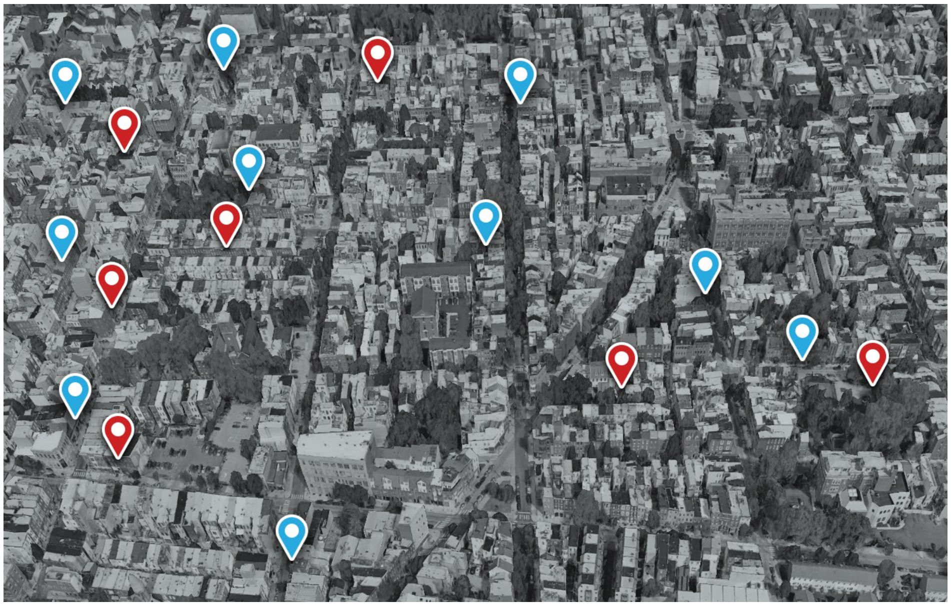

Figure 1. Residences with gateways (blue) and sensors (red)

To protect privacy of the participants, exact locations are not shown.

4. STUDY SETUP IN PHILADELPHIA

4.1 Gateway Placement

machineQ identified participants in the selected 1/4 square mile urban area to host and support LoRaWAN™ gateway installation in their residences. These locations are identified by blue icons on Figure 1 above. Gateways were self-installed by the participants without any physical assistance from the study team. The study team provided participants with detailed instructions on gateway setup; all gateways were provisioned for the machineQ network in advance of delivery. Participants were advised to place gateways near their Wi-Fi access points or TV cable boxes to mimic real-life future equipment location. Participants shared notes with descriptions of the gateway locations. Here are typical examples:

- Behind TV on stand; 10 feet from the east-facing window and 15 feet from the north-facing window;
- Second floor; two feet off ground, six feet from window with brick exterior;
- Basement; gateway is on desk four feet from a window;
- First floor of the apartment; on the floor next to TV and modem; two feet from the window;
- Next to the wireless gateway towards the center of the house; about 25 feet from the sidewalk;
- Second floor; about three feet from window facing sidewalk;

The MultiTech Conduit 8-channel gateway with wired Ethernet backhaul was used throughout the study. All gateways used machineQ code and were time-aligned by the machineQ network to assure synchronous time-stamping of the packets.

Figure 2. Examples of actual sensor placement locations

4.2 Sensor Placement

Available Airbnb rental locations were booked throughout the selected urban area to host end-devices. Red icons in Figure 1 on previous page identify these sensor locations. Multiple sensors were placed in each residence to represent real-life installations with varying degrees of radio frequency (RF) penetration challenges (in-between appliances, inside kitchen cabinets, in ovens/microwaves, in basements, etc.). See Figure 2 above for indoor sensor placement examples. Additionally, 10% to 15% of the sensors were placed in parked cars within the trial area to represent light indoor use cases. Locations exclusively dedicated to sensors (red icons) were separated at great distances from the locations exclusively dedicated to gateway placement (blue icons) as illustrated in Figure 1.

4.3 Sensor Setup and Configuration

To represent a generic LoRaWAN™ sensor with a five-centimeter rubber antenna (RF-AN0022), the study used the standard Semtech evaluation kit (Nucleo Pack) which is comprised of a SX1276MB1LAS radio board paired with a STM32L073RZT Nucleo board. Kits were powered by generic USB batteries. All sensors were configured for 18 dBm TX power and an 11-byte payload. An 11-byte payload is typical and enables the packets to be sent at all data rates. Both gateways and end-devices used Channels 24-31 defined in the LoRaWAN North American regional specification to separate the study traffic from regular machineQ network traffic as much as possible. Data Rate 4 (DR4), the most efficient modulation rate which results in the shortest packet duration, was excluded from the study. In a real-world deployment, most devices would not use DR4 since it has such short range.

4.4 Traffic

The initial traffic volume target was set at 250,000 packets per day. This target was determined as a goal to serve the theoretical dense network environment of 1,000 residences with ten LoRaWAN sensors per residence, with each sensor sending one packet on average once an hour (see Section 2 for calculation and assumptions). It would not be practical to use 10,000 sensors in the study. Instead, 100 sensors were used at a significantly accelerated packet transmit rate to mimic traffic that would have been generated by thousands of devices. This is a valid substitute as the study objective is to evaluate the ability of the machineQ LoRaWAN network to handle massive packet volume, regardless of the quantity of sensors producing this traffic. 10,000 sensors, sending packets once an hour, create traffic roughly equivalent to 100 sensors sending packets about every 35 seconds. Table 2 below summarizes the objectives of the study, including average packet rate, minimum and maximum inter-packet interval and target daily traffic volume. Packets were sent at random time intervals with no coordination between packets.

	Average Packet Rate per Device (seconds)	Minimum Interval (seconds)	Maximum Interval (seconds)	Daily Packet Volume per Device	Target Total Daily Traffic Volume (Packets)
Phase 1	32	4	60	2,500	250,000
Phase 2	16	4	24	5,000	500,000
Phase 3	6.5	4	9	10,000	1,000,000

Table 2. Target traffic density and packet volume, excluding packet duration

Figure 3. Phase 1 and phase 3 traffic density and inter-packet delay variance

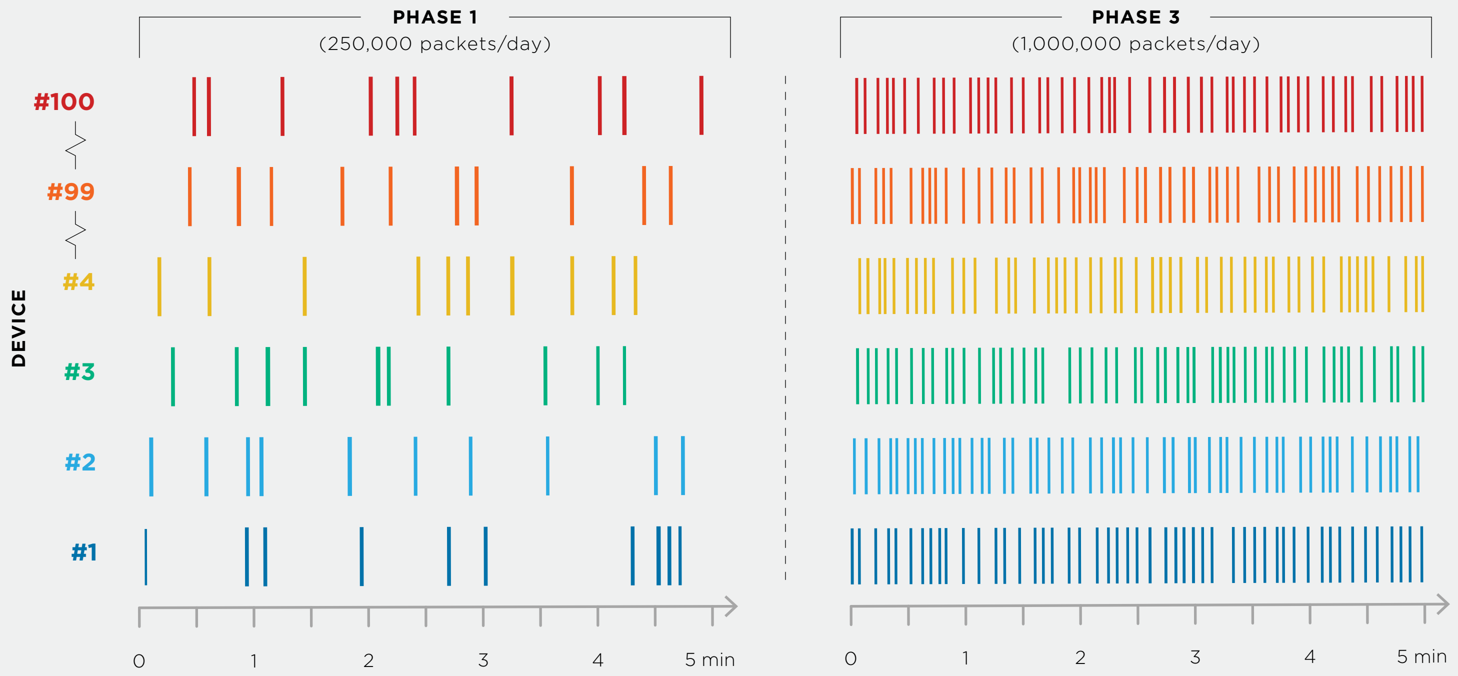

Phase 1 of the study was conducted at the initial target of 250,000 packets per day. Each sensor in the study generated packets, on average every 32 seconds, with a randomized inter-packet interval of +/- 28 seconds. Therefore, time between packets sent by the devices randomly varied from a minimum of four seconds to a maximum of 60 seconds while maintaining overall average packet rate of once every 32 seconds. Packets were sent randomly while ensuring total packet count hit 250,000. Figure 3 shows a theoretical example of inter-packet delay variance at 250,000 packets per day for a five-minute period under these randomized delay conditions.

While the target volume of 250,000 packets per day is the theoretical prediction of packet transmission in the estimated dense network environment, Phases 2 and 3 created heightened stressful traffic conditions. Phase 2 doubled the amount of generated traffic. In Phase 3, traffic volume was quadrupled to a throughput of 1,000,000 packets per day. These increments were implemented in order to evaluate network behavior with a significantly high probability of on-air packet collisions. Phase 3 was executed with the intent to identify a maximum sustainable network load. A secondary intent was to stress test cloud infrastructure to verify its readiness for upcoming high volume commercial deployments. Figure 3 also shows a theoretical example of inter-packet delay variation in Phase 3 over a five-minute interval.

Packets were counted in sequence and each packet was sent once without a confirmation acknowledgment. Success rate was evaluated by counting missing sequence numbers, also known as Packet Error Rate. If a packet wasn't received, it was missing in the incremental sequence. This study did not retransmit missing packets.

5. STUDY RESULTS

During the study, depending on the indoor location of the sensor and its proximity to the gateways, devices settled on a LoRaWAN™ configuration as highlighted in Table 3. Additionally, LoRaWAN gateways can receive transmissions at different data rates simultaneously. Table 3 shows the breakdown of devices per data rate across all three phases of the study. Time on Air is shorter for higher data rates yielding less time for packet interference.

Although some downlink traffic was present, the study's intent was to evaluate the network's ability to handle massive amounts of uplink packets. Please note that downlink traffic was not a focus of this study. Sensors did not require acknowledgment, and the limited downlink traffic was comprised of MAC commands from the machineQ network to the sensors.

Six of the sensors did not launch their intended configurations at the start of the study due to hardware issues; therefore, all results are based off of 94 active devices. See the yellow icon in Figure 4 for the location of the six devices. Out of the set of ten gateways, two units did not receive a significant amount of traffic. See the white icons on Figure 4 identifying these gateway locations. The remaining eight locations (green icons) received 99% of the traffic volume and provided sufficient gateway density for a successful capacity study.

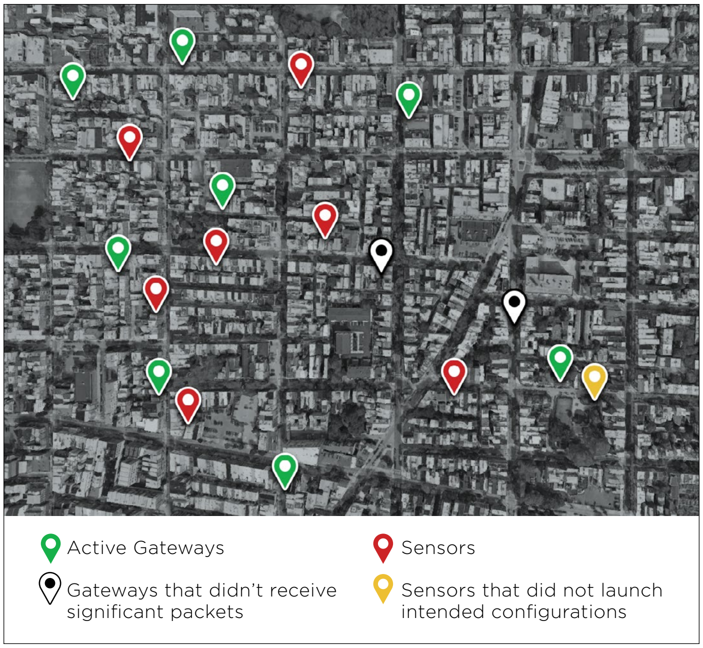

Figure 4. Active gateways (green) and sensors (red) locations

LoRaWAN Data Rate	LoRa® Configuration	Gateway RX Sensitivity (dBm)	Packet Duration (seconds)	Percentage of Devices (pkts)
Data Rate 0	SF10 / 125 kHz	132	0.37	13%
Data Rate 1	SF9 / 125 kHz	129	0.185	26%
Data Rate 2	SF8 / 125 kHz	126	0.103	27%
Data Rate 3	SF7 / 125 kHz	123	0.057	34%

Table 3. Breakdown of devices per data rate in all three phases

Figure 5. Shows distance between sensors and a gateway in the office that received packets

Interestingly, several thousand packets (out of 915,000 generated by the study) were also received by multiple gateways at the machineQ lab, located 1.4 miles away, which is obstructed by significant urban infrastructure — see Figure 5 above.

Table 4 on the next page provides a snapshot of the actual packet log from the study, documenting multiple packet overlap instances where two or more packets survived the overlap and have been delivered error-free to the cloud.

Figure 6 below provides further illustration of the 3-packet overlap instance highlighted in blue on Table 4. On Channel 24, the packet from Device with Serial number (DevAddr) 00A57720 (DR0) lasted 371 milliseconds. It finished its transmission and was time-stamped by the network at 10:00:22.706. This packet coincided with packet from Device Serial number (DevAddr) 00B34DAE using DR3 and with packet from DevAddr 0145243D (DR2, 103 milliseconds in duration). All three packets survived the overlap event.

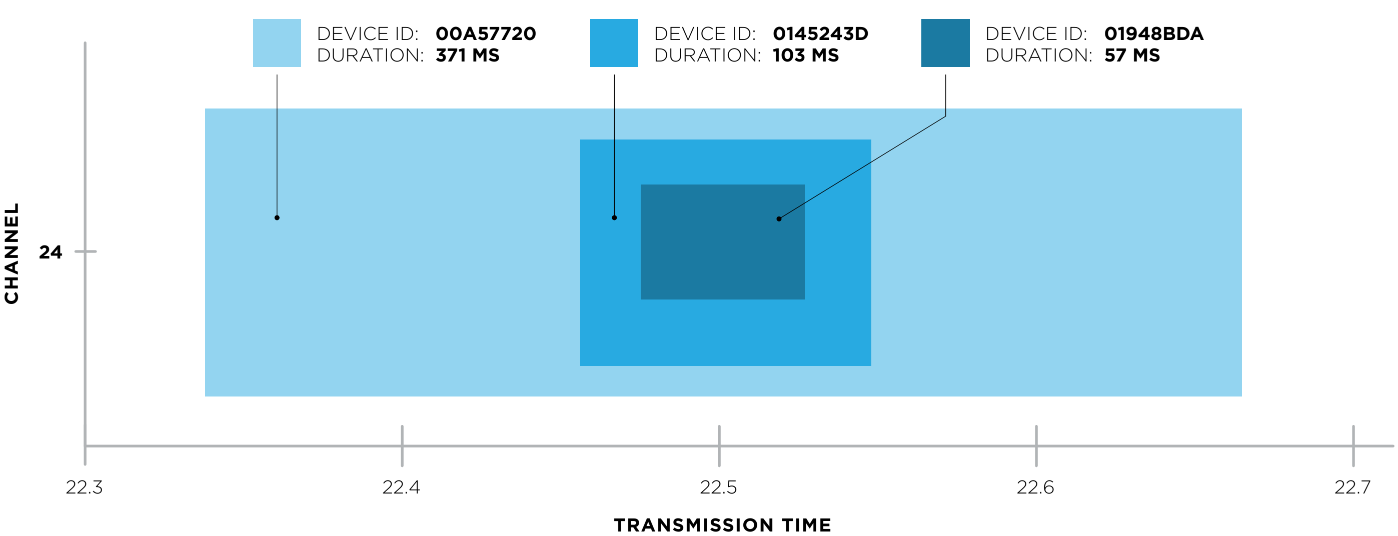

Figure 6. Example of recorded packet overlap, also known as packet collision. All packets survived the overlap occurrence.

Time stamp HR:MIN:SEC	AirTime (seconds)	DevAddr	SF	Data rate	Channel	GW Count	GW[1] Id	GW1 LoRa® RSSI	GW[2] Id	GW[2] LoRa RSSI	GW[3] Id	GW[3] LoRa RSSI
10:00:45.267	0.103	0180C116	8	2	LC24	4	004A2XXX	-105.13	004A2XXX	-111.41	004A2XXX	-118.12
10:00:45.252	0.185	017A243D	9	1	LC24	2	004A2XXX	-102.68	004A2XXX	-120.07		
10:00:31.770	0.057	00E4EA89	7	3	LC24	2	004A2XXX	-105.71	004A1XXX	-110.32		
10:00:31.759	0.185	00674DAF	9	1	LC24	2	004A2XXX	-106.41	004A2XXX	-117.28		
10:00:22.706	0.371	00A57720	10	0	LC24	1	004A2XXX	-111.46				
10:00:22.555	0.103	0145243D	8	2	LC24	2	004A2XXX	-103.32	004A2XXX	-109.89		
10:00:22.524	0.057	01948BDA	7	3	LC24	3	004A2XXX	-96.39	004A2XXX	-109.77	004A1XXX	-119.46
10:00:16.415	0.057	005AC116	7	3	LC24	3	004A2XXX	-103.32	004A2XXX	-109.28	004A1XXX	-119.58
10:00:16.372	0.185	00FAD159	9	1	LC24	1	004A2XXX	-114.07				
10:00:16.252	0.057	0142932D	7	3	LC24	3	004A2XXX	-101.76	004A2XXX	-106.40	004A1XXX	-116.10
10:00:34.800	0.103	012E243D	8	2	LC25	3	004A2XXX	-100.37	004A1XXX	-108.92	004A2XXX	-109.76
10:00:34.797	0.057	011CB0D4	7	3	LC25	3	004A2XXX	-101.71	004A2XXX	-114.64	004A1XXX	-115.54
10:00:03.821	0.371	0190EA89	10	0	LC25	1	004A2XXX	-120.27				
10:00:03.796	0.103	000746CA	8	2	LC25	3	004A2XXX	-106.85	004A2XXX	-110.86	004A2XXX	-116.54
10:00:56.838	0.185	00FAD159	9	1	LC26	2	004A2XXX	-108.46	004A2XXX	-114.41		
10:00:56.770	0.103	00324DAF	8	2	LC26	2	004A2XXX	-101.75	004A1XXX	-106.83		
10:00:02.813	0.371	018D4DAE	10	0	LC26	1	004A2XXX	-93.14				
10:00:02.756	0.185	01C54DAE	9	1	LC26	2	004A2XXX	-107.88	004A2XXX	-116.00		
10:00:58.728	0.185	01EF72A9	9	1	LC27	2	004A2XXX	-100.85	004A2XXX	-120.79		
10:00:58.711	0.057	0107D5CF	7	3	LC27	2	004A2XXX	-101.61	004A1XXX	-112.39		
10:00:24.417	0.057	005AC116	7	3	LC31	3	004A2XXX	-102.02	004A2XXX	-114.00	004A1XXX	-116.44
10:00:24.390	0.103	007046BD	8	2	LC31	5	004A2XXX	-99.49	004A2XXX	-101.94	004A2XXX	-109.17
10:00:24.250	0.185	000482EB	9	1	LC31	2	004A2XXX	-106.54	004A2XXX	-112.64		

Table 4. Examples of multiple packets concurrently coexisting on the same channel

	Target Daily Packet Volume	Actual Generated Packets per 12 Hours	Projected 24 Hour Packets (normalized to 100 sensors)	Network Load	Average Success Rate of First Transmission
Phase 1	250,000	114,305	247,779	7.45%	97.73%
Phase 2	500,000	233,712	509,095	13.22%	97.25%
Phase 3	1,000,000	491,705	1,081,714	25.68%	96.23%

Table 5. Phase by phase traffic volume, network load and success rate

Throughout the entire study, the network server and cloud infrastructure supported all of the traffic successfully. The success rate remained high despite the real-world urban environment with various obstructions from different materials and infrastructure. This success proves that propagation through a residential urban environment is not problematic for an enterprise-grade LoRaWAN™ network. Other capacity inhibitors, multipath fading, radio interference and backhaul throughput, were not directly studied, but due to high success rates it can be assumed that they will not have a significant impact on real-world deployment.

The study proved that at each packet volume of 250,000, 500,000, and 1,000,000, the success rate was over 95%. Table 5 above shows the success rate and actual packets transmitted per phase.

Phase 1 of the study lasted 12 hours and produced satisfactory results. Total generated packet count was 114,305 with an average success rate of 97.73%. Normalized to a 100-device study set, this resulted in a 247,779-daily packet volume. 97.73% success rate indicates that prospective IoT customers can deliver their traffic with a high degree of confidence.

Phase 2 of the study lasted 12 hours and generated 233,712 packets with an average hourly success rate of 97.25%, resulting in normalized 509,095 daily packet volumes. While traffic density doubled, success rate remained high. Table 6 below shows the average success rate per transmission for each phase of the study.

	First Transmission	Second Transmission*	Third Transmission*
Phase 1	97.73%	99.94847%	99.99997%
Phase 2	97.25%	99.92438%	99.99994%
Phase 3	96.23%	99.85787%	99.99980%

Table 6. Average success rate if acknowledgment mode enabled

*Only if subsequent transmissions are necessary

Figure 7. Hour by hour packet count and success rate

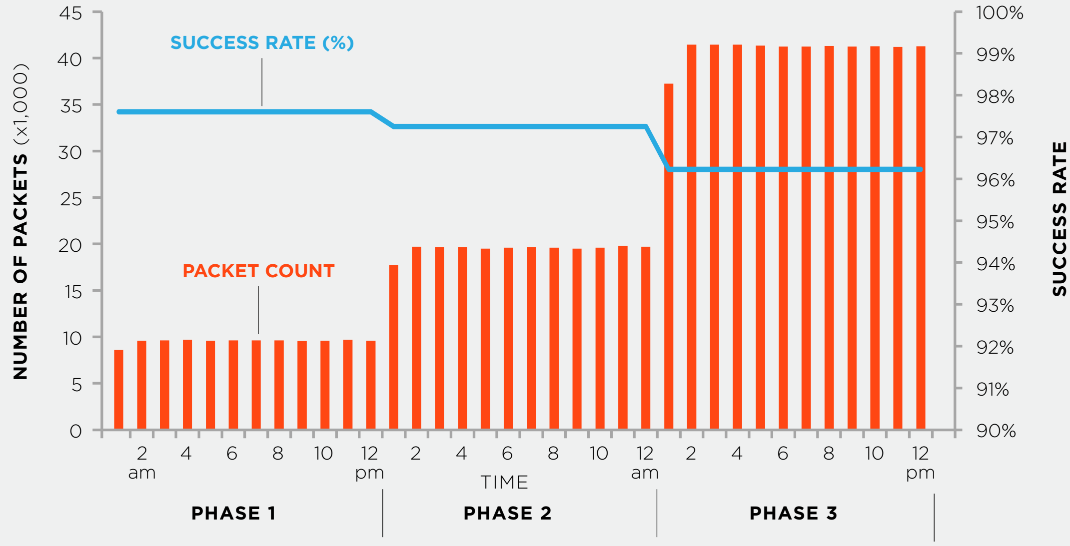

Figure 8. Gateway count (avg/max) and packet success per data rate for all three phases of the study

In Phase 3, the final and most challenging 12 hours of the study, we recorded 491,705 packets. This is equivalent to 1,081,714 packets per day (normalized to a 100-device study set). The traffic was handled by the eight active 8-channel gateways with an average hourly success rate of 96.23%. During this phase, traffic density was persistently maintained at an average rate of 12.7 packets per second. That is 1.6 packets per second per channel. This phase of the study created 25% network load conditions. Under these extreme load conditions, the study recorded a packet loss of 3.77%. Phase 3 demonstrated the enterprise-grade potential of a professionally installed and managed dense LoRaWAN™ network.

Figure 8 provides additional details on the specifics of how traffic was handled by the network in all three phases of the study. Packets were seen by 2.5 gateways on average.

Figure 9. Gateway diversity recorded throughout the study. Gateway diversity is defined as the % of packets received by multiple gateways.

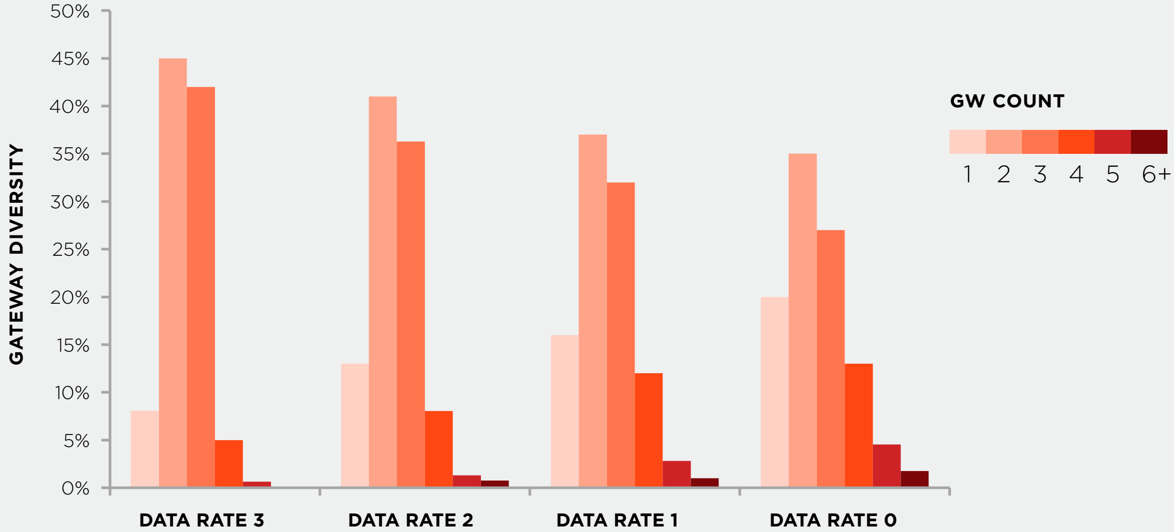

Figure 10. Hour by hour overlap ratio by data rate

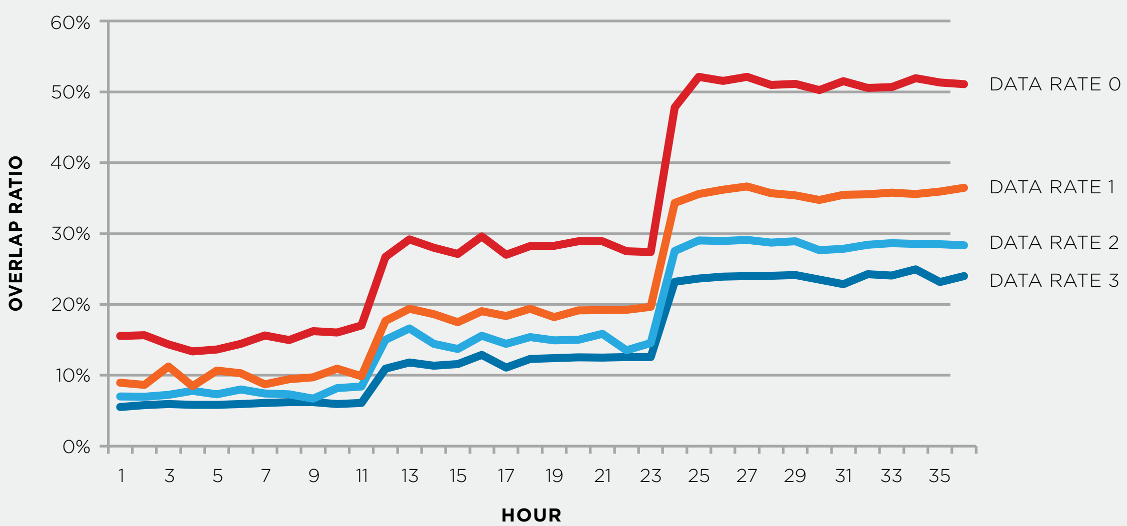

The majority of the traffic was seen by two to three gateways or more, as shown in Figure 9. Gateway diversity will be essential to achieving high packet delivery success rates. In a dense network, each packet will hit multiple gateways.

Additionally, the study recorded that a significant percentage of all packets were overlapping with traffic from other sensors. In high-density traffic conditions (as illustrated by Figure 3), there is a high probability of packet overlap, which occurs when two or more packets are received at the same time.

Figure 10 provides hour by hour overlap statistics. In Phase 3, more than 50% of the Data Rate 0 packets were overlapping and delivered successfully. LoRa® Technology allows overlapping packets to survive potential collision and to arrive at the network server error-free.

6. HOW WAS SUCCESS ACHIEVED?

How did we achieve high success rates under extremely stressful traffic conditions?

One of the most beneficial features of LoRa® Technology is data rate orthogonality. Orthogonality refers to the idea that LoRaWAN™ sensors are guided to various data rate configurations by the network depending on their proximity to the gateways. LoRa Technology uses a complex algorithm based on Time of Arrival (TOA) of a packet to a gateway and RSSI to determine which gateways to prioritize for certain packets. These different spreading factors are orthogonal or nearly invisible and non-interfering with each other. This feature allows for multiple LoRaWAN packets to simultaneously coexist on the same channel at the same time as seen in the results of the study. In any other viable IoT technology with similar adoption levels, such an event of two or more packets concurrently appearing on the same channel would have resulted in packet loss. LoRa orthogonality is a highly advantageous feature delivering increased network capacity. Combining LoRa orthogonality with a dense gateway deployment model results in enterprise-grade, high-capacity implementation.

What is gateway density and why is it important?

Gateway density is an important measure of a full-scale network. It represents the number of gateways that receive each packet. The success of the orthogonality of the LoRaWAN network depends on high gateway density. The team designed a model that predicts expected success rate based on gateway density. Out of all gateways in the study, one received the heaviest amount of traffic. The amount of traffic per data rate on that gateway was DR3 – 5.5%; DR2 – 4.5%; DR1 – 7% and DR0 – 6.6%. This gateway saw a total of 23.6% of all traffic. We used this number to extrapolate a theoretical success rate by applying three parameters to this traffic sample. First, we compensated for RSSI range distribution during the study. Then we used a set of theoretical spreading factor orthogonality formulas. Finally, we applied standard theoretical collision rate formulas. Based on these parameters, we can calculate a theoretical expected success rate for this traffic sample in a single gateway environment. This is reflected in the first row of Table 7. From the single gateway success rate, we calculated the effects of gateway diversity using recorded statistics (average and maximum gateway diversity plus data rate-specific success rates).
This is reflected in row two of Table 7.

	Data Rate 3	Data Rate 2	Data Rate 1	Data Rate 0
Single GW Expected Success Rate	93.40%	91.60%	85.70%	86.50%
GW Density Expected Success Rate	97.00%	96.70%	95.40%	96.10%

Table 7. Expected success rate from first transmission in single vs. dense gateway deployments

	Data Rate 3 (SF7)	Data Rate 2 (SF8)	Data Rate 1 (SF9)	Data Rate 0 (SF10)
GW Density Expected Success Rate	97.00%	96.70%	95.40%	96.10%
2.5 GW Density Actual Success Rate	97.80%	96.72%	96.50%	96.40%
Delta	+0.80%	+0.02%	+1.10%	+0.30%

Table 8. Expected vs. actual results

Gateway success rate is slightly higher at higher data rates, where Time of Arrival (TOA) of each packet is shorter, yielding less interference from other packets. The difference is < 2% and deemed not statistically significant. A secondary observation from this analysis is that available modeling and capacity estimation tools are very accurate. Comparing calculated and measured results reveal that actual results are within 1% of the estimate. See Table 8 above. The modeled rates were within 1.1% or better between predicted and actual performance figures. This shows machineQ has the ability to model capacity requirements to address specific density requirements. Table 9 above shows the average success rate for each transmission when acknowledgement mode is enabled.

Another significant benefit of gateway diversity is survivability of overlapping packets at the same data rate. When two packets with the same data rate arrive at the same time at a single gateway, one or both of these packets will not survive. See Figure 11, reflecting hour by hour study statistics on same Data Rate packet overlap. In a single-gateway deployment,

	First Transmission	Second Transmission	Third Transmission
Data Rate 0 2.5 GW Density Success Rate	96.40%	99.87%	99.9998%

Table 9. Expected success rate if sensors in acknowledgment mode

Note: At Data Rate 0, only 3.9% of devices would retransmit after initial transmission.

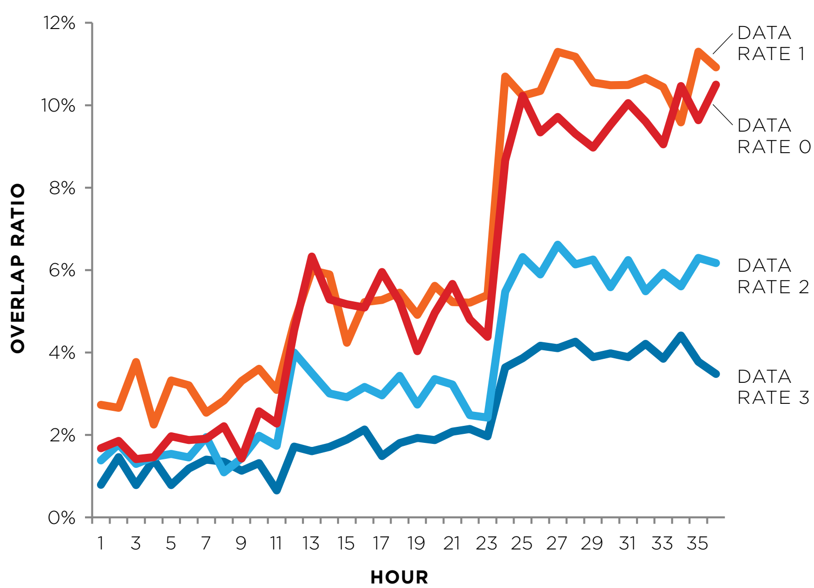

Figure 11. Overlap ratio at same rate

Time stamp HR:MIN:SEC	AirTime (seconds)	DevAddr	SF	Data rate	Channel	GW Count	GW[1] Id	GW1 LoRa® RSSI	GW[2] Id	GW[2] LoRa RSSI
10:00:10.443	0.185	003D2916	9	1	LC27	2	004A2XXX	-106.79	004A1XXX	-122.64
10:00:10.301	0.185	01C54DAE	9	1	LC27	2	004A2XXX	-107.61	004A2XXX	-114.28

Table 10. Example of successful reception of two same data rate overlapping packets

at least half of this packet volume would not have survived collisions resulting in significantly lower success rates. This is not the case in a dense cellular environment. While colliding with another packet on one gateway due to comparable signal strength, these same data rate packets may survive and be successfully received by other gateways favorably positioned. Table 10 provides a snapshot of the recorded packet log from the study, illustrating the same data rate packets surviving the overlap event and being successfully received by separate gateways.

7. CONCLUSIONS

machineQ and Semtech conducted a study with the objective of modeling a full-scale dense network deployment and testing different success metrics to ensure that a LoRa® based system would perform at capacity. Overall, the study showed that a full-scale LoRa deployment could handle more than the expected number of packets/day in a dense network environment.

The study revealed the following about each of the initial objectives:

Achieved a packet success rate of 90% or greater in a dense scenario. The study concluded that in each phase, the success rate was above 95% and therefore indicates that LoRa can transmit data reliably at scale. The 25% Network Load in Phase 3 was easily handled by a dense network with a 96% success rate on first transmission. We were not able to conclude a maximum network load from this study; however, the calculation tool predicts that this would be two million packets per day.

Confirmed the ability for multiple overlapping packets to be delivered successfully. There were multiple instances of overlapping packets in the study that were delivered successfully. The dense network deployment packet loss due to overlapping packets is minimal in full-scale deployment.

At full-scale deployment, the network successfully utilized gateway diversity and showed that each device communicates with different gateways when transmitting data. Our study revealed that, on average, each sensor communicated with 2.5 gateways confirming the success of orthogonality in full-scale deployment, which is a key advantage of a multi-tenant network.

The cloud infrastructure and network server performed successfully in the full-scale deployment scenario. During this study, the network server performed to expectations and was able to easily handle traffic from the increase in packet rate across phases.

The study was able to provide an initial understanding of the effect of environment on signal propagation, packet success rate, and gateway diversity. Although this particular metric was not studied in depth, the success rate was high despite the real-world environment of the study. Further exploration is needed to detail how specific environmental factors can affect propagation, packet success rate, and gateway diversity.

The team was able to confirm that other capacity inhibitors — multipath fading, radio interference, and backhaul throughput — did not significantly affect capacity in this scenario. Although multipath fading, radio interference, and backhaul throughput were not directly evaluated, the packet success rate was extremely high. Similar to exploring the environment's effect on the different success metric, further investigation is required to determine exactly how these interferences may impact a LoRaWAN™ network.

Looking ahead, it's clear that network capacity is not an obstacle for ubiquitous LoRaWAN coverage. By allowing end-devices to transmit at any time, complexity is shifted from the devices to the network, significantly reducing the amount of energy required at the end-device to negotiate and maintain network coordination. This joint study with Semtech confirmed that machineQ has executed a reliable, enterprise-grade LoRaWAN network for a variety of IoT applications.