



Predicting LoRaWAN Capacity

Table of Contents

List of figures	2
List of tables.....	3
1. Introduction.....	4
2. Signal Levels	5
3. Orthogonality Factor.....	7
3.1. Interference Levels Observed in Measured Data	7
3.2. Orthogonality Factor Simulation	8
3.3. Interference Levels Inferred from Simulated Data.....	10
4. Prediction of Collisions and Packet Error Rate (PER)	12
4.1. Collision Probability in Random Access (ALOHA)	12
4.2. Single Gateway Packet Error Rate	12
4.3. Network Packet Error Rate of a Frame	14
5. Prediction of Collisions and Packet Error Rate – Simulated Network.....	16
References	18
Disclaimer	19

List of figures

[Figure 1.](#) Signal Level Distribution of Received Frames

[Figure 2.](#) RSSI Comparison - One Gateway and Multiple Gateways

[Figure 3.](#) RSSI Distribution per Data Rate and Traffic Load over Time

[Figure 4.](#) Network Simulation

[Figure 5.](#) Average Interference levels from Out-of-Range Devices

[Figure 6.](#) Packet Count per Hour and Data Rate Distribution

[Figure 7.](#) Channel Load

[Figure 8.](#) Average Reception Redundancy

[Figure 9.](#) Derived Collision Rate versus Load

[Figure 10.](#) Data Rate and Corresponding PER

List of tables

[Table 1.](#) Data Rates and Required Signal-to-Noise Ratios

[Table 2.](#) Orthogonality Matrix

[Table 3.](#) Simulated Orthogonality Matrix

[Table 4.](#) Load Per Gateway, by Spreading Factor

[Table 5.](#) Average Load

[Table 6.](#) Overlap Probabilities

[Table 7.](#) Collision Rate Probability for a Single Gateway

[Table 8.](#) Collision Probabilities for Stage 3 of the Trial

Chapter 1. Introduction

This document is a theoretical analysis of a LoRaWAN® [network trial](#) performed by MachineQ and Semtech in Philadelphia in 2017. The trial consisted of 10 indoor gateways and 108 indoor devices in an urban environment. The goal of the trial was to demonstrate the network capacity of LoRaWAN.

During the trial, devices sent approximately one million frames, over two days. The network load was set to various levels. The success rate evolved accordingly and was always higher than 96 percent.

From the analysis, we show that closed-form formulas can predict capacity accurately. We also show that these formulas can replace long simulations.

Chapter 2. Signal Levels

First, let's analyze the statistics related to the received signal strength. For each frame we have access only to the highest RSSI received. If more than one gateway receives the frame correctly, we know the number of gateways which received the signal, but not the individual RSSI for each gateway. Since LoRa® operates below the noise floor, the RSSI itself can be lower than the noise floor, which is approximately -120dB.

We also know the data rates used. This test was conducted in the US region, so only SF7, SF8, SF9 and SF10 were used (not SF11 nor SF12).

Figure 1 shows the RSSI distribution of correctly received frames. The legend for each curve indicates the proportion of frames which follow this distribution, the average RSSI (in dBm), and the standard deviation of RSSI (in dB). For instance, 30.7% of frames are received by exactly 3 gateways; when this occurs the average RSSI is -103.7dBm and the standard deviation of RSSI is 5.9dB.




Figure 1. Signal Level Distribution of Received Frames

We see that most frames were received by two or three gateways. The number of gateways receiving a signal depends on how the LoRa Network Server (LNS) configures the Adaptive Data Rate (ADR). Different ADR settings can have different redundancy targets leading, for instance, to a lower ratio of single gateway receptions. Here 17% of frames are received by a single gateway.

The next figure is an attempt to model the cases of one, two, three, and four gateways receiving the signal, in comparison to just one. The blue curve shows the mean and standard deviation, and models the single receiving gateway case with a Gaussian distribution with same mean and standard deviation. We then model the case of two receiving gateways with a maximum of two independent Gaussian variables. For three gateways we take a maximum of three, etc.

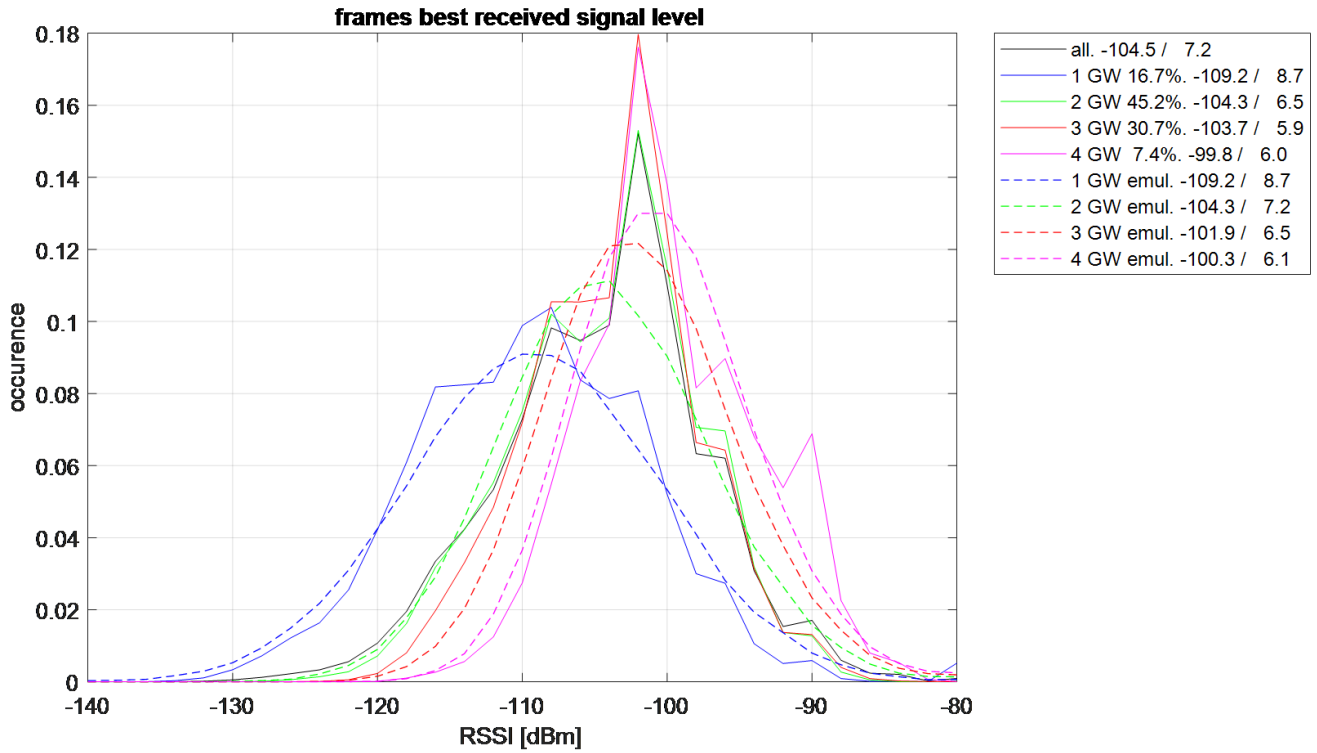


Figure 2. RSSI Comparison - One Gateway and Multiple Gateways

We see that there is a good fit between the observed data and the modeled curves, whatever the number of receiving gateways (i.e. the plain lines match the dotted lines). This means that modeling the RSSI by a gaussian variable is valid, and when a signal is received over several gateways each RSSI is independant. There is a peak in the observation at -102dBm, which corresponds to a measurement bias. This bias comes from a rounding error of the gateway which makes -102dBm twice as likely as -101 or -103.

Chapter 3. Orthogonality Factor

3.1. Interference Levels Observed in Measured Data

Data rates orthogonality is a shortcut for expressing that different data rates can coexist. Indeed, assuming data rates are different, an aggressor frame will affect a victim frame only as much as white noise would. Since LoRa data rates can be received with negative signal-to-noise ratios (SNR), frames with different data rates can be received at the same time. [Table 1](#) shows the required SNRs for each data rate. These numbers correspond to first-generation LoRa demodulators, the second generation is 1dB better.

Table 1. Data Rates and Required Signal-to-Noise Ratios

Data Rate	SF5	SF6	SF7	SF8	SF9	SF10	SF11	SF12
Required SINR	-2.5dB	-5dB	-7.5dB	-10dB	-12.5dB	-15dB	-17.5dB	-20dB

In addition, for same data rate interference, when there is a signal level difference higher than 7dB, the strongest frame is received correctly. This holds true when the coding rate is 4/5, normal interleaving, which is the LoRaWAN default LoRa PHY configuration.

Next, we plot the measured RSSI distribution per data rate. On the same plot, we will also show the traffic load versus time. While the traffic load varies over time, the data rate ratios remain constant, as does the RSSI distribution per data rate. The set-up is tested with a single data rate (SF10), then the first load level is tested from hours 7 to 18, a higher load from hour 19 to 30, and the highest load from hour 31 to 45.




Figure 3. RSSI Distribution per Data Rate and Traffic Load over Time

From this distribution, and the table above, we can compute the probability that a time overlap of frames yields the loss of a frame. The distribution is simplified to four random Gaussian variables with same mean and Sigma as the observed distribution. We report this in a matrix, with columns representing the aggressing data rate, and lines the victim data rate.

Table 2. Orthogonality Matrix

Data Rate	SF7	SF8	SF9	SF10
SF7 as victim	86%	10%	3%	8%
SF8 as victim	8%	84%	2%	6%
SF9 as victim	11%	9%	85%	8%
SF10 as victim	6%	5%	2%	78%

This table shows that if a frame transmitted using SF8 is overlapped by an SF10 frame while being received by a gateway, there is a six percent probability that the SF8 frame will not be received correctly. Thus, there is a 94 percent probability that the gateway will receive the frame correctly.

If the orthogonality were perfect, there would be a zero percent error rate for overlaps, except on the diagonal.

In this dataset, the RSSI measures are incomplete because we only have access to the best RSSI among all receiving gateways of each frame. It is difficult to determine how different the orthogonality matrix would be with a complete set of measures.

3.2. Orthogonality Factor Simulation

Let's look at a slightly different scenario, one with indoor/deep indoor nodes, and outdoor gateways. The simulation uses 660 gateways with a density of 1.5 gateways per square kilometer (950 m site distance). 100,000 nodes are spread on the grid, and propagation to each gateway is simulated using the Hata model^[1] for a small-to-medium city, with shadowing, fast fading, and indoor penetration losses distributed from 20dB to 40dB. The path loss exponent value here is 3.6.




Figure 4. Network Simulation

Spreading factors SF7 through SF12 are used in this example. Here, we apply an Adaptive Data Rate (ADR) and Transmit Power Control (TPC) strategy with a margin of 8dB to account for fast fading. The range of TPC is 20dB, i.e. once the fastest data rate is reached, the transmitted power can be reduced by up to 20dB.

To compute the orthogonality matrix, we consider the gateways in the center, and gather the distribution of RSSI for each data rate. Then we proceed as with the measured data: for each couple of spreading factors, (victim SF, aggressing SF), we derive the probability that the aggressing SF RSSI is too high for the victim frame to be received. This probability is the probability of $(RSSI_{aggressor} > (RSSI_{victim} - victim_required_sinr))$. For instance, if victim is SF12, and aggressor SF7, this is the probability:

$$P = Proba(RSSI_{SF7} > (RSSI_{SF12} + 20dB))$$

This probability is computed as:

$$P = \int_{x=-142dBm}^{0dBm} proba(RSSI_{SF7} > (x + 20dB)) * proba(RSSI_{SF12} = x)$$

Table 3. Simulated Orthogonality Matrix

Data Rate	SF7	SF8	SF9	SF10	SF11	SF12
SF7 as victim	76%	5%	1%	0%	0%	0%
SF8 as victim	17%	80%	1%	0%	0%	0%

Data Rate	SF7	SF8	SF9	SF10	SF11	SF12
SF9 as victim	17%	6%	80%	0%	0%	0%
SF10 as victim	18%	6%	2%	80%	0%	0%
SF11 as victim	18%	6%	2%	0%	81%	0%
SF12 as victim	18%	6%	2%	0%	0%	64%

This matrix is slightly different than that shown in [Table 2](#), which relies on incomplete measures. We can note, however, that if we remove the diagonal, the sum of each line is roughly the same. Given this information, we can see that ADR and power control tend to give devices with a higher RSSI a higher data rate, which means these devices have a better chance of creating errors in case of a collision.

In this MachineQ and Semtech trial, there is no ADR; each physical device emulates traffic from several virtual devices, using the four data rates.

3.3. Interference Levels Inferred from Simulated Data

The measured data is generated from a relatively small network, with only 10 gateways in the same neighborhood. Thanks to the simulation, we can estimate how much additional noise comes from devices that are not in range of a given gateway. Collisions can be classified as one of two types: collisions from frames that can be received by the gateway we consider, or collisions from frames that are out of the coverage area.

Next, using the simulation above, let's consider a central gateway and count the devices in range, along with the data rate at which they operate. We then define three channel load levels: 10 percent, 100 percent, and 400 percent. The 400 percent load means that on average, on each LoRa channel, four transmissions of in-range devices occur. For each channel load, we derive the duty cycle required to achieve the load specified for the in-range devices, assuming an identical rate of frames per device, whatever their data rate. Last, we apply the same duty cycle to the devices that are not in range, and measure the average power received by the central gateway from these devices. We compute this sum over different disk sizes, to check whether the integral converges. The result is plotted in [Figure 5](#). We compare this average interference level to the thermal noise floor. This thermal noise floor is experienced by any receiver, and values at normal temperature are $-114\text{dBm}/\text{MHz} + \text{NF}$, where NF is the noise figure of the receiver. In our case, the bandwidth is 125KHz, and typical NF is 3dB for a gateway, so receiver noise floor is: $-114\text{dBm}/\text{MHz} + 10 \cdot \log_{10}(0.125\text{MHz}) + 3\text{dB} = -120\text{dBm}$.




Figure 5. Average Interference levels from Out-of-Range Devices

First we see that the integral converges.

This convergence is expected, since the path loss exponent (PL) is higher than 3. To explain this point, let us note $itf(r)$, the total interference level from devices within a disk of radius r . It is computed with:

$$itf(r) = K \cdot \int_{d=0}^r \pi d^2 * \frac{1}{d^{PL}} = K \cdot \int_{d=0}^r \pi \frac{1}{d^{PL-2}}$$

where K is a constant that depends on node density, transmit power, and average impact of shadowing and fast fading. Such an integral converges as r grows if $PL-2 > 1$, i.e. $PL > 3$.

In addition to convergence, we also see that a two kilometer radius is enough to gather most of the interference from devices which are out of range of the gateway.

Second, it is clear that the average interference level coming from devices not in range is lower than the thermal noise, even at a high network load. This shows that we can neglect the interference from out of range devices, as such interference does not increase the average noise floor.

Of course, the instantaneous interference power might be higher. Taking a disc radius of two kilometers, we count seven times more un-connected devices than connected devices. This means that at a load of 400 percent there are, on average, 28 simultaneous transmissions from out-of-range devices that contribute to the interference level. The maximum interfering power from a single device corresponds to an SF7 transmission, which would be just below SF7 receiving sensitivity, i.e. -130dBm.

Chapter 4. Prediction of Collisions and Packet Error Rate (PER)

4.1. Collision Probability in Random Access (ALOHA)

The usual formula for calculating the probability of collision is:

$$P_{col} = 1 - \exp(-load * 2)$$

Where load is the offered load. This assumes that when a time overlap occurs, both frames are lost. This also assumes that all frames have equal length.

A more generic formula can be used, assuming that the length of victim and aggressor frames are different. We call it the probability of time overlap, rather than collision, because a time overlap does not always result in lost frames.

$$P_{ovlap} = 1 - \exp(-load_interferer * (1 + length_victim / length_interferer))$$

When **length_victim** is null, we go to the known formula of probability that channel is occupied, at any instant.

When **length_victim** equals **length_interferer**, we go to the classic formula.

Last, if **length_victim** tends to infinity, the probability of overlap tends to 1, whatever the load of the interferer might be.

4.2. Single Gateway Packet Error Rate

Let's compute the load per gateway, i.e. the traffic each gateway receives.

The frames all have 8 bytes of LoRaWAN payload, Table 4 below shows their duration.

Table 4. Load Per Gateway, by Spreading Factor

Data Rate	SF7	SF8	SF9	SF10
Time on Air	57ms	103ms	185ms	377ms

We cannot directly observe the offered load for each gateway because some packets are lost and because the logs only provide the number of receiving gateways, rather than the gateway IDs. Figure 5 shows the total number of received frames, without duplicates, over time. From these curves, we then derive the spreading factor distribution, to check whether it is relatively constant over time.




Figure 6. Packet Count per Hour and Data Rate Distribution

Next, we derive the load for the network, noting that the load is evenly spread over eight channels. Figure 6 shows the load for each channel.




Figure 7. Channel Load

Now we compute the average load for the three steps of the trial, as shown in [Table 5](#).

Table 5. Average Load

Test Phase vs Load	Phase 1	Phase2	Phase 3
SF7 load	0.7%	1.3%	2.8%
SF8 load	0.9%	1.9%	4.0%
SF9 load	1.8%	3.4%	7.1%
SF10 load	1.1%	3.2%	6.6%
Total load	4.4%	9.8%	20.4%

From the load, we compute the probabilities of overlap. For each phase, there are four probabilities for each data rate. Here, we take the example of Phase 3, and report the overlap probabilities in [Table 6](#).

Table 6. Overlap Probabilities

Overlap %	SF7 interf.	SF8 interf.	SF9 interf.	SF10 interf.
SF7 as victim	5.4%	5.8%	8.5%	7.2%
SF8 as victim	8.1%	7.7%	10.1%	7.9%
SF9 as victim	13.1%	11.3%	13.2%	9.4%
SF10 as victim	22.3%	18.1%	19.2%	12.4%

Not all time overlaps result in reception errors. We can now use the pseudo orthogonality matrix to multiply, term by term, the overlap probabilities. This gives the collision probability. The sum of each line is the total probability of collision, from all data rates. We report the collision rate, for a single gateway, in [Table 7](#). We see that, compared to [Table 6](#), the diagonal terms are slightly reduced, while the non-diagonal terms are reduced a considerable amount, thanks to the orthogonality of the various spreading factors.

Table 7. Collision Rate Probability for a Single Gateway

Overlap %	SF7 interf.	SF8 interf.	SF9 interf.	SF10 interf.	Total
SF7 as victim	4.7%	0.6%	0.2%	0.6%	6.1%
SF8 as victim	0.7%	6.5%	0.2%	0.5%	7.8%
SF9 as victim	1.4%	1.1%	11.3%	0.7%	14.5%
SF10 as victim	1.4%	1.0%	0.3%	9.6%	12.4%

4.3. Network Packet Error Rate of a Frame

To determine the network PER, we compute the gateway redundancy from the received packets. We then assume that this redundancy is very close to what it would be without collisions. (This is backed by the fact that the redundancy is stable during the three stages of the trial.)

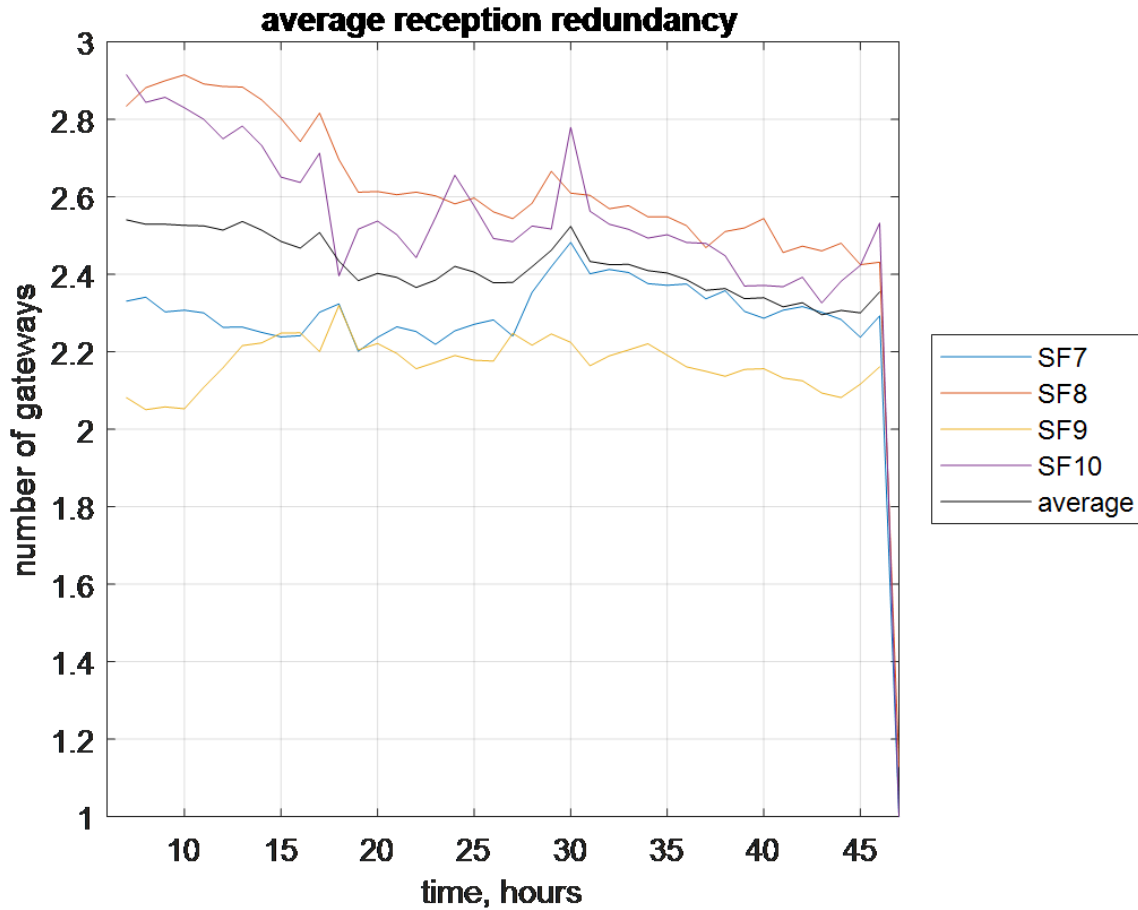


Figure 8. Average Reception Redundancy

Instead of taking the redundancy directly, we subdivide it by the single receiving gateway case, the two gateways case, the three gateways case, etc. The final results come from the weighted average of the probability of these distinct cases.

To give a simplified example, let's assume that 30% of frames are received by a single gateway, 50% by two gateways, and 20% by 3 gateways. The network level probability of error per packet PER_{nwk} , from a gateway level probability PER_{gw} , will be: $PER_{nwk} = 0.3*PER_{gw} + 0.5*(PER_{gw})^2 + 0.2*(PER_{gw})^3$.

This gives the following collision probabilities for each SF for the third stage of the trial.

Table 8. Collision Probabilities for Stage 3 of the Trial

Collision	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%

We see that this closely matches our observations.

Chapter 5. Prediction of Collisions and Packet Error Rate – Simulated Network

We now show that this method can be applied to a simulated network. We start with a propagation and ADR simulation, then we compare the results from an explicit collision simulator and the closed formulas described above.

The explicit collision simulator creates a time grid, adds frames to this grid, computes received signal levels for all gateways at each time instant, then computes the PER from each frame's SINR. The load is varied, so the end result is the system PER as a function of load. The load unit is the number frames per hour per gateway. This number is the total number of frames divided by the number of gateways, so that multiple receptions have no impact on our definition of load.

The method using closed form formulas starts with the same propagation and ADR simulation. This gives signal levels and data rates for all simulated devices. From these we compute the orthogonality matrix. Then, for each load, we compute overlapping probabilities and collision probabilities at the gateway level. Last, from the reception redundancy for each spreading factor, which is also an output of the propagation and ADR simulation, we derive the PER at the network level.

The comparison of two methods is shown in [Figure 9](#).




Figure 9. Derived Collision Rate versus Load

The dotted green line indicates that the ADR is worse. The ADR margin is raised from 5dB to 8dB, meaning a higher average redundancy (1.8 compared to 1.4 gateways), and the TPC is reduced from 20dB to 15dB. The primary effect of this is to increase the load of the lowest data rate. The propagation model here assumes placement of devices indoors/deep indoors (20dB to 40dB penetration losses). We also expect the devices to be static, so a low margin of 5dB is reasonable. SF12 devices have a higher redundancy than average: 1.9 gateways at 5dB ADR margin.

[Figure 10](#) shows the result of another simulation, where we see that the PER depends on the data rate. Low data rates tend to have a higher PER.

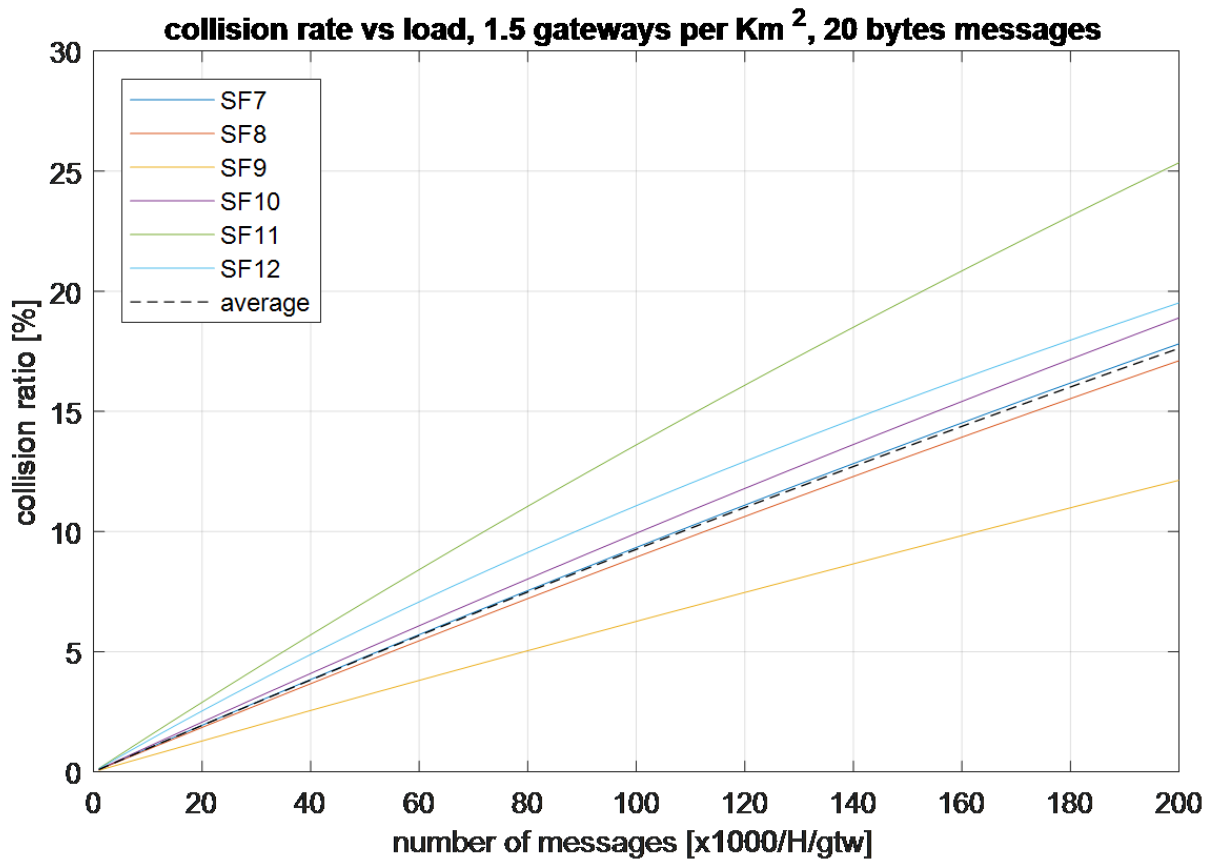


Figure 10. Data Rate and Corresponding PER

From this curve, we can see that from a capacity perspective, message repetition is a good strategy. For instance, looking for a PER=1%, the maximum load is 10,000 messages per hour. This means that with a single transmission, 10,000 messages per hour can be transmitted to achieve a PER of 1%.

To achieve 1% PER after two transmissions, the single transmission PER target becomes 10%, this is because frames are independent from the point of view of collisions, therefore the probability of losing 2 frames is $(10\%)^2 = 1\%$. 10% PER for single transmissions corresponds to 110,000 messages per hour, or 55,000 unique messages with double transmission. Last, triple transmission corresponds to 21.5 percent PER, as $(0.215)^3 = 1\%$, and a unique load of $\sim 240,000/3 = 80,000$ messages per hour.

Repetition, therefore, increases capacity for a given Quality of Service (QoS) target. Of course, it is possible to adjust the repetition rate to the device data rate. Delay between successive transmissions does not matter.

It is also possible to use a Forward Error Correction on groups of messages, which is a generalization of messages repetition. This trades PER for delay, and considerably reduces the overhead when compared to systematic repetition. For applications that are not delay-sensitive, this is the solution of choice.

References

- [1] From Wikipedia, the free encyclopedia, Hata model

Disclaimer

IMPORTANT

Information relating to this product and the application or design described herein is believed to be reliable, however such information is provided as a guide only and Semtech assumes no liability for any errors in this document, or for the application or design described herein.

Semtech reserves the right to make changes to the product or this document at any time without notice. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

Semtech warrants performance of its products to the specifications applicable at the time of sale, and all sales are made in accordance with Semtech's standard terms and conditions of sale.

SEMTECH PRODUCTS ARE NOT DESIGNED, INTENDED, AUTHORIZED OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT APPLICATIONS, DEVICES OR SYSTEMS, OR IN NUCLEAR APPLICATIONS IN WHICH THE FAILURE COULD BE REASONABLY EXPECTED TO RESULT IN PERSONAL INJURY, LOSS OF LIFE OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE. INCLUSION OF SEMTECH PRODUCTS IN SUCH APPLICATIONS IS UNDERSTOOD TO BE UNDERTAKEN SOLELY AT THE CUSTOMER'S OWN RISK. Should a customer purchase or use Semtech products for any such unauthorized application, the customer shall indemnify and hold Semtech and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs damages and attorney fees which could arise.

The Semtech name and logo are registered trademarks of the Semtech Corporation. All other trademarks and trade names mentioned may be marks and names of Semtech or their respective companies. Semtech reserves the right to make changes to, or discontinue any products described in this document without further notice. Semtech makes no warranty, representation or guarantee, express or implied, regarding the suitability of its products for any particular purpose. All rights reserved.

© Semtech 2024