

White Paper

The Microwave Reach Rage

Understanding and Debunking Vendors'
Outrageous Reach Claims



DragonWave

www.dragonwaveinc.com

NOTICE

This document contains DragonWave proprietary information. Use, disclosure, copying or distribution of any part of the information contained herein, beyond that for which it was originally furnished, requires the written permission of DragonWave Inc.

The information in this document is subject to change without notice and relates only to the product defined in the introduction of this document. DragonWave intends that information contained herein is, to the best of its knowledge, correct and accurate. However, any/all liabilities associated with the use or accuracy of the information contained herein must be defined in a separate agreement between DragonWave and the customer/user.

DragonWave®, Horizon® and Avenue® are registered trademarks of DragonWave Inc. ©2017 DragonWave Inc. All rights reserved

TABLE OF CONTENTS

| | |
|--|-----------|
| THE MICROWAVE REACH RAGE UNDERSTANDING AND DEBUNKING VENDORS' OUTRAGEOUS REACH CLAIMS | 1 |
| INTRODUCTION | 4 |
| THE MAGIC CIRCLE | 5 |
| TRUE LINK PERFORMANCE IMPROVEMENT..... | 6 |
| UNDERLYING MAGIC CIRCLE PHYSICS | 7 |
| THE THEORY OF INFLATION | 12 |
| STRETCHING THE E-BAND ENVELOPE..... | 14 |
| IN A PARALLEL UNIVERSE..... | 17 |
| CONCLUSION – DEBUNKING THE MYTH..... | 17 |



Introduction

Every now and then, the marketing department of microwave radio vendor X, Y or Z decide that it's time to make some waves in an otherwise dull wireless backhaul world. A proven recipe to draw attention is to claim record-shattering performance for your products.

In our business, knocking down reach or capacity limits works best to turn heads, and that's why we see quite a few claims to fame in those areas.

Unfortunately, the majority of those "record-breaking" reach or capacity claims are misleading, and in fact no performance benchmarks are being shattered at all. In this white paper, I'll try to explain why that is the case and how to debunk outrageous reach or capacity claims.

The Magic Circle

Let me first introduce what I dubbed the Magic Circle, a conceptual tool for visualizing the interrelation of range, capacity and availability in wireless links.

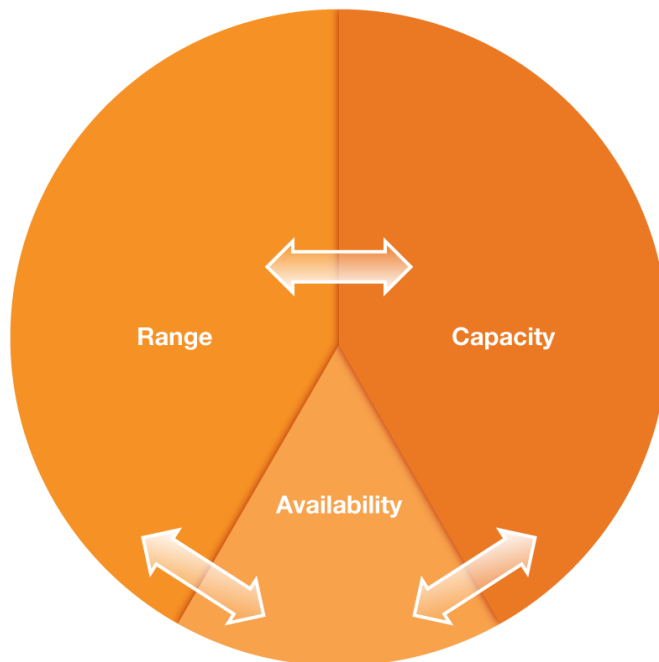


Figure 1: The Magic Circle

The capacity, range and availability of a point-to-point microwave radio link are a zero-sum game once the equipment, antennas and equipment settings are selected.

The size – or circumference – of this circle represents the link budget of a point-to-point microwave radio link. That budget is defined by the transmit power, receiver sensitivity and antenna gain available on that link. Once the equipment and antenna hardware are selected, the radio channel size (bandwidth) is set and the maximum modulation mode is defined, the size of the circle – representing the link budget – cannot change. This means that if capacity is to be increased, either range or availability will suffer. If range is to be extended, there will be a price to pay in terms of availability or capacity. Improvement of any one of those key parameters will result in a deterioration of either one or both of the other parameters: it's a zero-sum game.

True Link Performance Improvement

The true road to progress consists of increasing the **size** of the Magic Circle, or the link budget, and thereby increasing the range, capacity or availability of a link *without* sacrificing its range, capacity or availability. This improvement can be achieved in four distinct ways – assuming that we’re using the most prevalent single-transmitter single-receiver (SISO¹) microwave radio system architecture:

1. **Mounting larger, higher-gain antennas.** This is by far the simplest way to increase the link budget. As a bonus, higher-gain antennas have a more focused, narrow radiation pattern and are therefore less sensitive to off-boresight interference than lower-gain antennas. Unfortunately, it’s not always possible to install larger antennas because of:
 - a. Obtrusiveness – local authorities or site owners tend to restrict antenna size,
 - b. Weight limitations (static loading) of the mounting structure – larger antennas tend to be heavier, too.
 - c. Physical space limitations – there may be insufficient space or access to mount a larger antenna.
 - d. Dynamic loading (wind load) limitations of the mounting structure – larger antennas have a larger area exposed to moving air exerting higher torque on the mounting structure.
 - e. Cost – larger antennas are always more expensive than their smaller equivalents.
 - f. Availability – in most frequency bands, there’s a maximum antenna size – in some instances just 60 cm / 2 foot – that can be had.

2. **Increasing the transmit power of the transmitter.** This can be done by e.g. switching from GaAs² to GaN³ power amplifier circuitry, or by other means. It is typically the least expensive way to improve the link budget. The flip side of this approach is that:
 - a. There could be more interference between adjacent links, negating some of the gain in dense networks. This can be mitigated by using Automatic Transmit Power Control (ATPC) – unleashing the extra power on demand, during deep-fade events – or by using higher performance antennas, for example ETSI Class 4 instead of ETSI Class 3, with a better off-boresight interference rejection.
 - b. Generally the power consumption and heat dissipation of the product will increase as well, requiring larger heat sinks.
 - c. The product cost will generally be higher as well. GaN semiconductors are more expensive than GaAs semiconductors as are the larger power supplies and heat sinks that are required.

3. **Improving the receiver sensitivity.** This is the most subtle and elegant but also the most challenging approach. It implies a lowering of the noise figure and/or the implementation of more performant FEC⁴ algorithms. Both approaches imply a cost impact, because more exotic components and/or cooling sub-systems would have to be implemented. As the point-to-point microwave radio business is very cost-

¹ Single Input Single Output – either end of a link utilizes a single transmitter and a single receiver. MIMO systems (Multiple Input Multiple Output) can achieve better performance (capacity, range, availability) at the expense of (much) higher system complexity and cost.

² Gallium Arsenide – the most common semiconductor used for microwave radio radio-frequency circuitry.

³ Gallium Nitride – a semiconductor used for high-power microwave radio radio-frequency circuitry.

⁴ Forward Error Correction.

sensitive, a material improvement in the noise figure is not feasible unless a technology breakthrough occurs.

4. **Improving the baseband spectral efficiency by deploying payload compression.** With this method, the baseband data that is sent through the link is losslessly compressed *prior* to being transmitted. By doing this, the capacity of the link can be boosted by more than 100%, without compromising the link reach or its availability. Note that the actual link speed on the physical level will *not* increase, but the user will experience higher throughput nonetheless. Obviously, the success of this approach will depend on the compressibility of the baseband traffic pattern. Compression gain will be adversely affected by the presence of pre-compressed data (data already compressed by the user application) or encrypted data in the traffic flow.

All microwave radio vendors have access to the same antennas manufactured by the same vendors, use state of the art receivers with comparable receiver sensitivities and depend on well-known industry-standard FEC algorithms. The main differentiator in terms of system gain and link budget – and thus reach, capacity or availability or the size of the Magic Circle – will therefore mostly boil down to the available transmit power and the implementation of lossless payload compression.

However, a focus on increasing the available transmit power is not necessarily the optimal approach as the overall radiated power (density) is generally controlled by the spectrum regulatory body and there may be monetary penalties for using a higher power or potentially financial incentives for using a reduced power.

Underlying Magic Circle Physics

So why do some microwave radio vendors lay claim to much better ranges for their products than their competitors do for similar product platforms? Assuming that the products that are being compared operate in the same frequency band and channel, use the same (direct mount) antennas and have comparable transmitter, receiver and FEC performance and are not implementing payload compression, those ranges cannot differ that much can they?

Well, yes they can! By playing around with link capacity, availability and range – thus by shifting the borders inside the Magic Circle – and by carefully choosing the target deployment area to select favourable climatic regions and path elevations, the same link can support a very wide range of hop lengths and link capacities.

Atmospheric and rain attenuation impact on link budget

How come? Let's start with the basics. In addition to the transmit power, receiver sensitivity and antenna gain, the following factors will affect the microwave link performance:

1. Atmospheric attenuation. This is strongly frequency dependent as per Figure 2. Atmospheric attenuation is also dependent on air density and therefore a function of altitude (not shown in Figure 2). Atmospheric attenuation generally increases with increasing frequency, with absorption peaks at distinct frequencies like the oxygen absorption peak around 60 GHz, in the so-called V-band. It remains below 1 dB/km for most frequency bands below 100 GHz. Most terrestrial microwave radio links operate in the 3-86 GHz frequency range.
2. Rain attenuation. Rain (and precipitation) strongly affect the propagation of radio waves with frequencies higher than 3 GHz – refer to Figure 3. Rain attenuation tends to increase with radio frequency up to about 100 GHz.




Figure 2: Sea-level atmospheric attenuation of radio waves in (dB/km) on the vertical axis as a function of frequency on the horizontal axis (GHz)




Figure 3: Rain attenuation of radio waves in (dB/km) on the vertical axis as a function of frequency on the horizontal axis (GHz) for different color-coded rain intensities (mm/h)

As can be inferred by comparing Figure 2 and Figure 3, medium rain (12.5 mm/h) generally causes more attenuation than atmospheric attenuation and heavy rain (25 mm/h) tends to cause roughly ten times more attenuation than the atmosphere. This is why rain is the most frequent cause of link degradation or complete, temporary link breakdown.

Rain zone and availability target impact on link budget

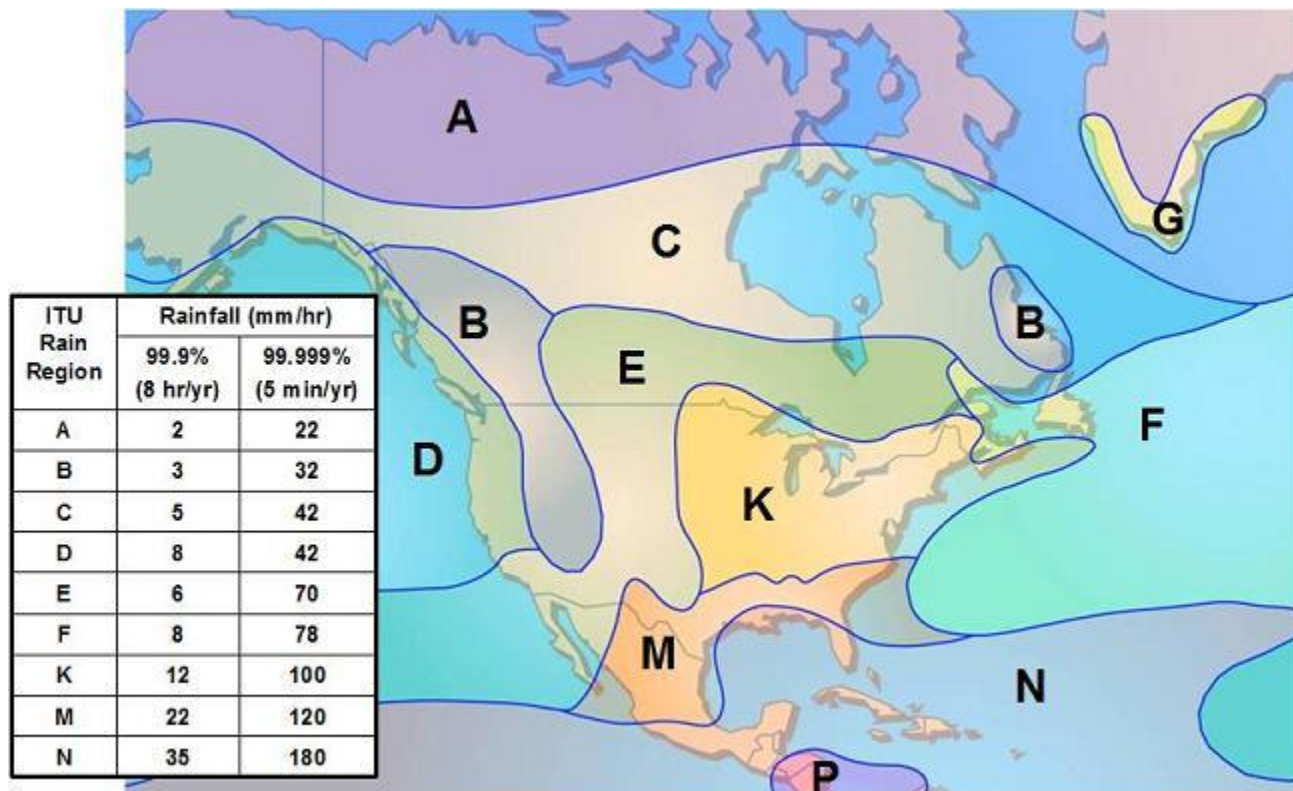


Figure 4: ITU rain intensity regions for North America

How much rain do we have to take into account when designing a terrestrial microwave radio link? Let's have a look at Figure 4. New York City for example happens to be in rain zone K. In that zone, rain intensity exceeds 12 mm/h during 8 hours a year on average, and it exceeds 100 mm/year (tropical rain) during just 5 minutes per annum on average.

If we want our link to have a radio frequency availability of 99.999%⁵ or better – equivalent to less than 5 minutes of downtime per annum – then our link will have to cope with a rain intensity of 100 mm/h. If this link operates in the 10 GHz band, that implies (Figure 3) enduring a rain attenuation (fading) of 3 dB/km. Assuming a 10 km hop length, the fade margin (spare link budget) of the link ought to be at least 30 dB. Note that atmospheric attenuation in the 10 GHz band stands at less than 0.02 dB/km (Figure 2). If we were to use that 30 dB fade margin to *extend* the 10 km hop, at the expense of its availability, we could comfortably extend it *ten-fold* to 100 km, assuming sufficient ground clearance⁶. A 100 km link would incur 20 dB of additional free space loss plus 2 dB of additional atmospheric attenuation, and we have at least 30 dB of fade margin at our disposal.

Capacity impact on link budget

Now how do capacity and range interrelate? The *capacity* of a microwave radio link is largely dependent on the used channel bandwidth and modulation scheme. Doubling the channel size will double the capacity, provided that all other parameters remain the same.

That capacity doubling will however occur at the expense of reducing system gain and link budget by approximately 3 dB, because of the two times higher noise power at the receiving end of the link. That 3 dB corresponds to a bit less than 1 km of range at 10 GHz at a 99.999% availability target in New York (ITU rain region K).

A higher-order modulation scheme will increase capacity as well, but by a fixed amount (in bits per second per Hz) per modulation step, leading to a lower *relative* capacity gain per modulation step as modulation and overall link capacity increases (Figure 5).

⁵ 99.999% as a radio frequency availability target is amongst the strictest in our industry. A 99.995% target is more common, and in some cases 99.99% or less suffice. The availability target depends on the criticality of the application and the availability of topological redundancy for that particular link (routing alternatives).

⁶ That ground clearance is not trivial for long links due to the curvature of the earth. Long hops in sub-10 GHz bands may also suffer from multipath fading, necessitating the deployment of space diversity configurations. Multipath fading is not taken into account in our simplified analysis.

Capacity as a Function of Modulation

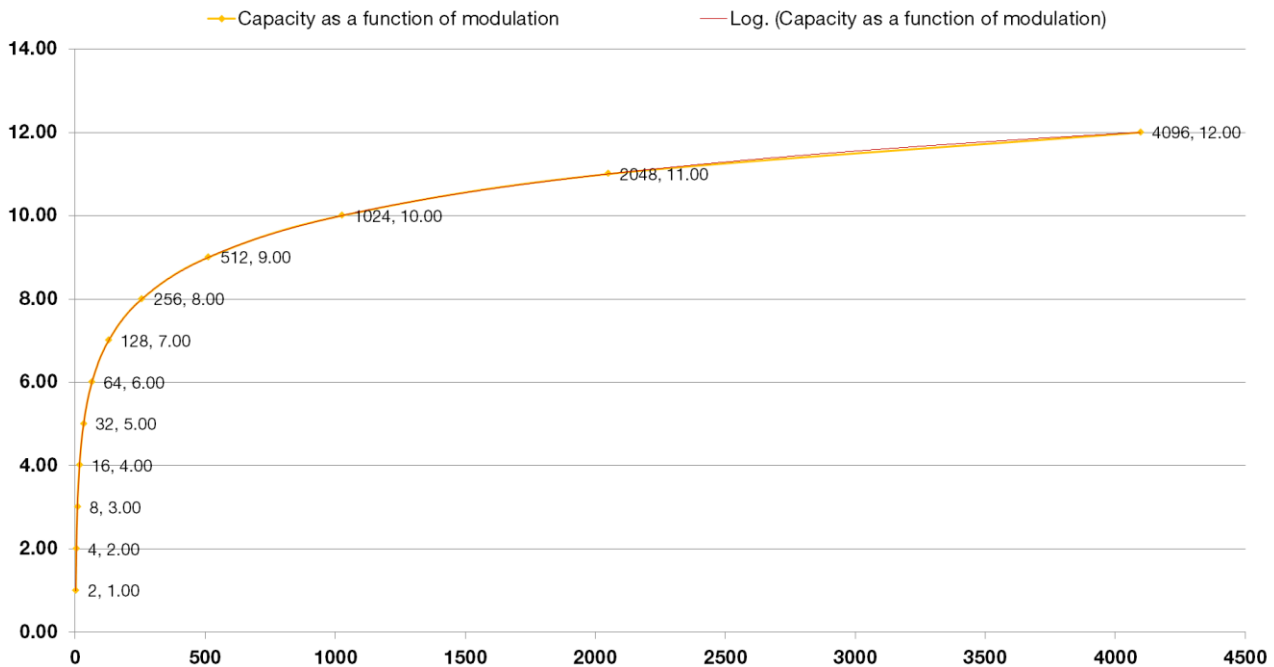


Figure 5: Relative capacity gain (vertical axis) referenced to BPSK⁷ as a function of modulation (horizontal axis)

The modulation in this graph ranges from BPSK (2-state modulation) up to 4096-state QAM⁸. Each consecutive step implies the doubling of the number of modulations states.

Per modulation step, the system gain and link budget will drop by approximately 3 dB as well, because each modulation step implies a doubling of the number of modulation states, which makes signal detection twice as difficult in the presence of noise. As can be inferred from Figure 5, 2048 QAM boosts link capacity 11-fold compared to BPSK, but the required link budget is roughly 30 dB higher as well (10 modulation steps x 3 dB).

The conclusion is that increasing link capacity will decrease link range and/or availability, and if the capacity increase is achieved through increasing modulation, that effect will be stronger than if the same capacity increase is attained by increasing the channel bandwidth.

⁷ Binary Phase Shift Keying.

⁸ Quadrature Amplitude Modulation.

The Theory of Inflation

As we have seen in the previous section, microwave radio link range claims can be inflated tremendously, sometimes by an order of magnitude. It's just a matter of:

- ❖ Selecting a favourable rain zone; an arid place featuring a low probability of high-intensity rain or precipitation. Obviously, most real-life customers and applications won't be lucky enough to reside in such arid places.
- ❖ Lowering the availability target to a ridiculously low level. In our industry, target availabilities are generally 99.95-99.995%. Anything below a 99.95% availability target is suspect, as most customers won't tolerate downtimes that are well past 4.5 hours per annum.
- ❖ Lowering the link capacity (modulation, channel size) to the minimum levels supported. There's nothing wrong with that per se, especially if the lower capacity occurs briefly, during heavy rain, and is restored to nominal levels as soon as possible (adaptive modulation). At the same time, most real-life applications will require higher-than-minimum link capacities at any given time. That means that maximum range won't equal the maximum range at the lowest possible modulation level and channel size. In addition, regulators in certain jurisdictions may define the lowest modulation a link may operate in, imposing a lower bound on spectral efficiency (bps/Hz), and thereby precluding maximum-range links.

What microwave vendors often do is state that equipment X has Capacity C and range R. More often than not, capacity C is the maximum capacity at a (very) short range and/or in the absence of precipitation and range R comes at the lowest capacity for that particular equipment. C and R can *never* be obtained simultaneously. This is not a problem as long as the buyer understands the tradeoff between range, capacity and availability.

Let's go through some examples now. Please refer to Figure 6. Each of the 2 graphs depicts three different configurations of the very same equipment. In this case, we're examining a narrowband (50 MHz channel) V-band link (56-58 GHz) deployed in Moscow, Russia.

- ❖ The blue line is a plot of range versus capacity (modulation) for integrated 20 cm antennas,
- ❖ The center line is a plot for the same equipment using 30 cm (1 ft.) parabolic dishes,
- ❖ The rightmost line is a plot for the same equipment using 60 cm (2 ft.) parabolic dishes.

Two things are obvious:

1. **Antenna gain** – related to antenna area (aperture) – has a significant impact on hop length,
2. **Modulation** – has a strong impact on range, and an obvious impact on capacity. Doubling capacity though modulation would rob us of approximately one third of the range.

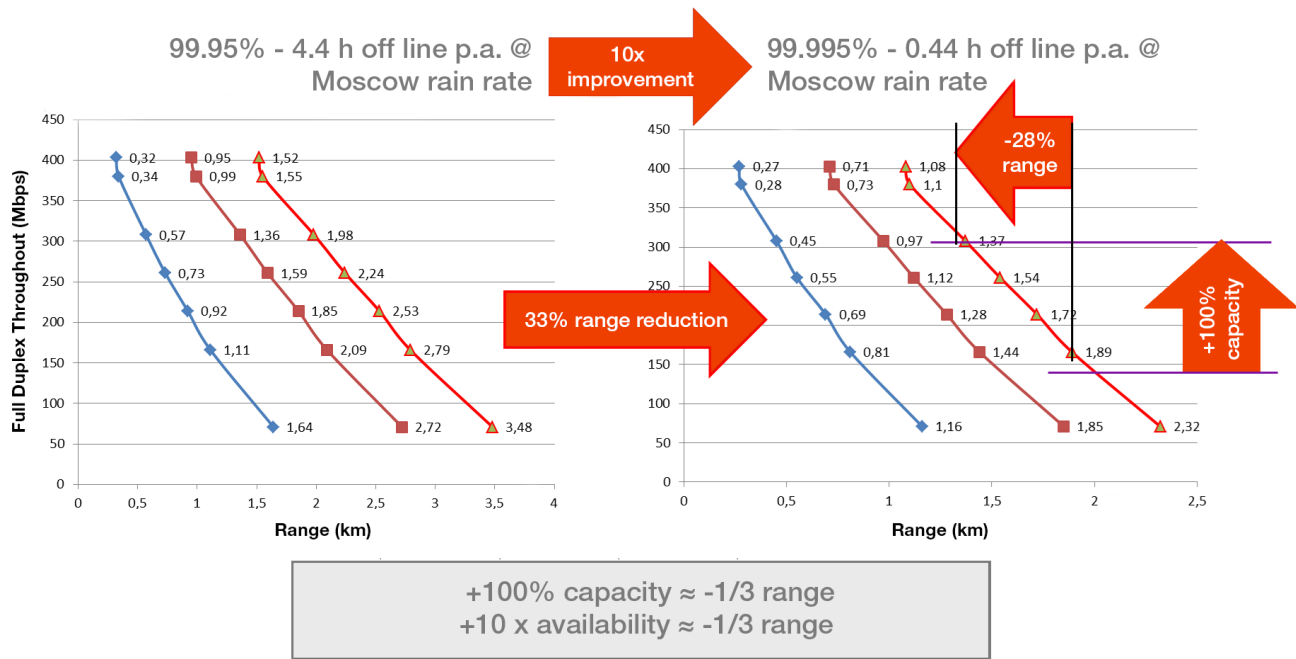


Figure 6: The effect of availability, antenna size and modulation (capacity) on range

What's perhaps even more striking is the impact of a subtle variable – availability – on range. If we aim for 99.995% availability, or 26 minutes of downtime p.a., then our maximum range at maximum capacity with 60 cm dishes in Moscow for this particular equipment is a mere 1.08 km. If we just relax our availability target to 264 minutes of downtime p.a., then our range grows to 1.52 km, a 41% improvement, and at no obvious cost at all – except perhaps customer satisfaction.

As can be inferred from Figure 6, we *could* claim that this V-band link has a capacity of 400 Mbps and range of 3.5 km. Actually, as the link is a DragonWave link, it supports bulk compression (Bandwidth Accelerator). Therefore, provided that it's fed with compressible payload, we could claim a 1 Gbps capacity and a range of 3.5 km. Quite remarkable for a V-band link that operates inside the 60 GHz attenuation spike depicted in Figure 2.

Stretching the E-band Envelope

Most microwave radio links operate inside the 3-42 GHz spectrum. In addition to those “mainstream” links, there are fewer links operating in the V-band (around 60 GHz) and the E-band (around 80 GHz). Several vendors are experimenting with links around and beyond the 100 GHz mark, to quench the thirst for ever higher capacities.

The E-band in particular has attracted lots of attention lately, because a vast amount of spectrum is available, enabling throughputs in the 1-10 Gbps range, and because atmospheric attenuation (Figure 2) is modest compared to the V-band and even higher frequencies.

The challenge with E-band, and in fact all radio frequencies beyond roughly 30 GHz, is their susceptibility to rain (Figure 3). And as rain attenuation is much stronger than atmospheric attenuation in almost all parts of the world (Figures Figure 3 and Figure 4), rain rates and availability targets will be the limiting factor for E-band range.

So what can we expect from E-band links in terms of range? Let’s provide some examples – refer to Figure 7.

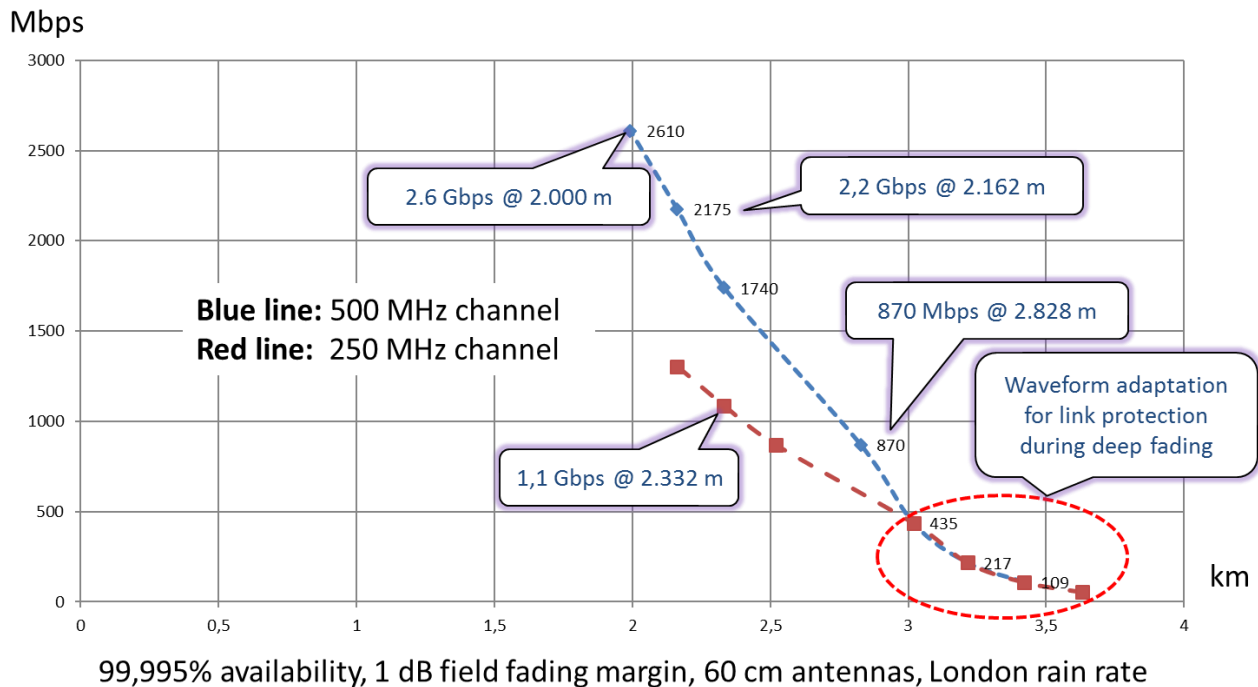


Figure 7: Range (horizontal axis, in km) versus capacity (vertical axis, Mbps) for an E-band link in London

Utilizing 60 cm (2 ft.) parabolic dishes at both link ends, at carrier-grade 99.995% availability

As can be gleaned from the figure, this link can achieve full capacity (2.6 Gbps) at a 2 km range. If 870 Mbps of capacity is acceptable, and we use a 500 MHz channel, the link can be stretched to 2.8 km. Using a 250 MHz channel, we can hit 1.1 Gbps at up to 2.3 km. This is the kind of performance envelope (1-3 Gbps at a 2-3 km range) to be expected from an E-band link in a real-life deployment theater and supporting carrier-grade (99.995%) availability targets.

So why do some vendors claim to support links in excess of 15 km (10 mi.) in the E-band? Do they have access to superior technology? Most likely not.

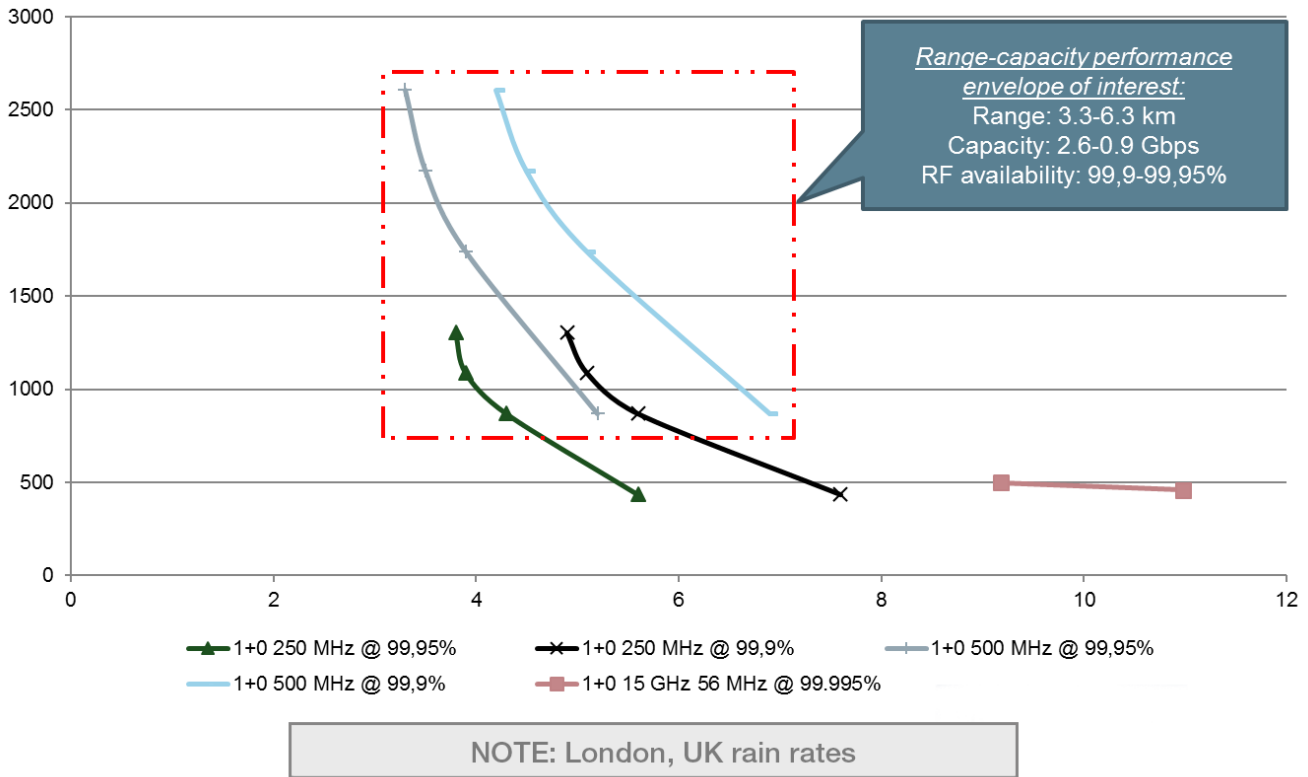


Figure 8: Range (horizontal axis, in km) versus capacity (vertical axis, Mbps) for an E-band link in London

Utilizing 60 cm (1 ft.) parabolic dishes at both link ends, at 99.9% and 99.95% availability

Let's have a look at Figure 8. All parameters, except for the availability target, equal the parameters in Figure 7. As can be seen, this leads to a considerable extension of the range at full-speed (2.6 Gbps), from 2 km in Figure 7, to 3.3 km (grey line) or 4.2 km (light blue line) in Figure 8. If we accept a capacity in the order of 500 Mbps, we can get to a range of 7.6 km. So what's the range of this equipment 2.0 km? 2.8 km? 3.3 km? 4.2 km? or 7.6 km? It all depends on the rain zone, availability target and capacity target.

Still, 7.6 km is a far cry from 15 km... Let's see what happens if we whittle down the availability target to well below 99.9%.

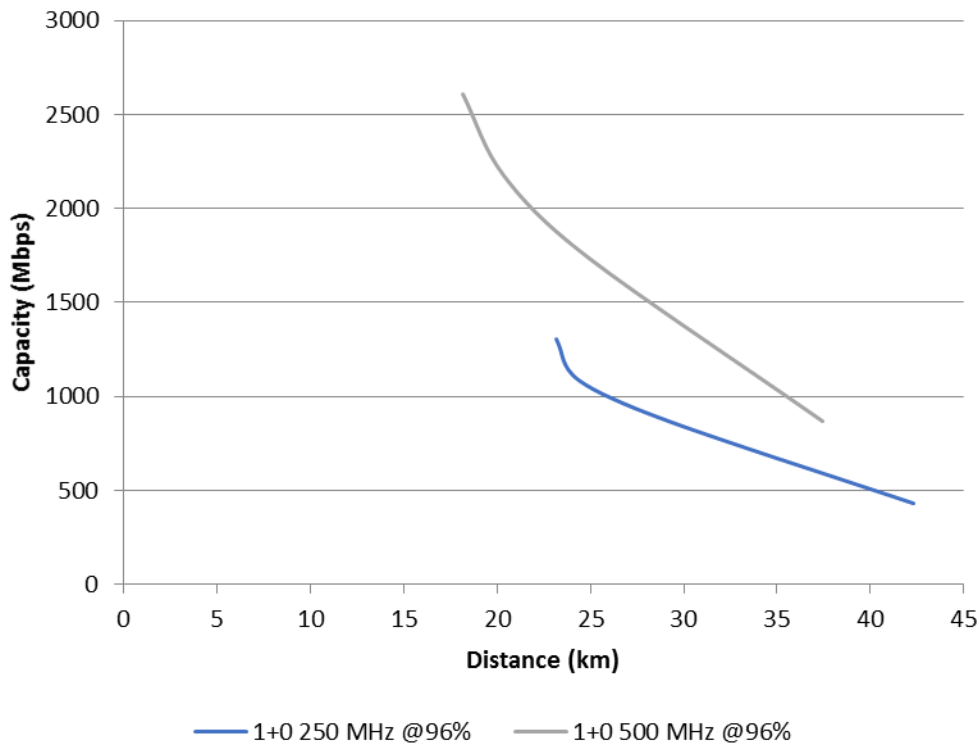


Figure 9: Range (horizontal axis, in km) versus capacity (vertical axis, Mbps) for an E-band link in London

Utilizing 60 cm (1 ft.) parabolic dishes at both link ends, assuming no rain fade margin at all, and a 250 MHz and 500 MHz channel respectively

Please note that Figures Figure 7 and Figure 8 depict the performance of exactly the same hardware in exactly the same place, London, using the same availability calculation methodology. The sole variable is the availability target. In Figure 9, we've removed the constraints that precipitation imposes on us altogether. In Figure 7, hitting 3 km at full speed (2.6 Gbps) appears to be elusive. We can do 2.8 km at 0.87 Gbps at best. In Figure 9, we shatter the 15 km (10 mi.) limit with ease, and do that at "full throttle".

Can we build a 15 km E-band link in London? Yes we can! Can we build a *reliable* 15 km E-band link in London? That depends on what you define as "reliable". A link featuring e.g. 98% availability will be down for about 2% of the time. That corresponds to more than a calendar week – 7 days – per annum on average. It depends on the application whether or not that is acceptable, but for 98% of the applications it definitely won't be.

Yet on rare occasions, even 96% link availability will be good. The author once met with a customer in Russia who wanted to deliver Internet connectivity to remote villages using sub-20 GHz long range (30 km up) microwave radio links. He wanted to implement those links "cheap and cheerful", minimizing expenditure, therefore limiting antenna cross-section and constraining the link budget. It seemed mission impossible.

Then I asked about the availability target for those links. Surprisingly, the customer mentioned 96%⁹. When I asked whether 2 weeks of downtime would be acceptable, he retorted: *why not? They don't have Internet for 52 weeks per annum now, they won't be having Internet for 2 weeks per annum in the future, so what?*

In a Parallel Universe

Several vendors propose to use a sub-30 GHz link *in parallel with* an E-band (80 GHz) link. The sub-30 GHz link provides for good link availability, typically 99.995%, as it's less susceptible to rain, at a comparatively low capacity, typically in the 200-500 Mbps range. The E-band link provides for very high capacity, typically between 1-3 Gbps, and a relatively low availability, typically down to 99.9%.

In Figure 8, we can see that a 15 GHz parallel link (close to the bottom right corner) can easily provide in excess of 400 Mbps at up to 11 km range (using 60 cm dishes) at a good 99.995% availability. Having such a link in parallel with our E-band link, we can get decent throughput between 3.3 and 6.3 km – depending on the desired capacity.

In this way, we can get the best of both worlds. Most of the time (99.9%) we get 1-3.1 Gbps of throughput, during approximately 0.5% of the time we get reduced throughput (down to the lowest throughput of the low-frequency link) and during just 0.005% of the time we get no service at all. Link engineering and defining the availability target(s) becomes more cumbersome though.

Conclusion – Debunking the Myth

We conclude that it's impossible to unambiguously define “the” range and “the” capacity for any given microwave radio link. Range, capacity and link availability are all interrelated, remember Figure 1, the Magic Circle. The rain rates in the deployment area and antenna sizes will play a major role, too.

Most of the wild range claims that we see around us are misleading. They are based on deployment in extremely arid areas or on low or very low link availability targets at the lowest link capacities. All three assumptions are utterly unrealistic in the vast majority of real-world deployments.

Faced with an extreme range claim, one should ask for additional information like:

1. Deployment rain zone,
2. Availability target,
3. Link capacity at the stated range,
4. Used antenna sizes,
5. The used radio channel size.

⁹ The ITU model is validated for availability targets between 99% and 99.999%. If link modeling beyond those bounds is required, a different model needs to be used, e.g. the Crane model.

There is also a workaround. The easiest and best way to compare microwave radio system radio frequency performance and its range potential is in fact *not* through direct range comparisons. What is needed is the guaranteed (not nominal or typical) system gain (in dB) in a defined channel (e.g. 50 MHz, 56 MHz, 500 MHz) at a given modulation (e.g. 2048 QAM, 64 QAM) and a defined BER¹⁰ (e.g. 10^{-6}). This ensures a fair “apples to apples” comparison, and it’s breathtakingly simple to perform.

For more information on this approach, please refer to the following white paper: <http://bit.ly/2mlxDwl>

¹⁰ Bit Error Rate.