Distributed-Input-Distributed-Output (DIDO) Wireless Technology

A New Approach to Multiuser Wireless

Steve Perlman, President & CEO, Rearden Companies
Antonio Forenza, Ph.D., Principal Scientist, Rearden Companies
www.rearden.com

1. Abstract

Distributed-Input-Distributed-Output (DIDO) wireless technology is a breakthrough approach that allows each wireless user to use the full data rate[1] of shared spectrum simultaneously with all other users, by eliminating interference between users sharing the same spectrum. With conventional wireless technologies the data rate available per user drops as more users share the same spectrum to avoid interference, but with DIDO, the data rate per user remains steady at the full data rate of the spectrum as more users are added.

As a result, DIDO profoundly increases the data capacity of wireless spectrum, while increasing reliability and reducing the cost and complexity of wireless devices. DIDO deployment is far less expensive than conventional commercial wireless deployment, despite having vastly higher capacity and performance, and is able to use consumer Internet infrastructure and indoor access points.

The potential of DIDO is to have unlimited number of simultaneous users, all streaming high-definition video, utilizing the same spectrum that a single user would use with conventional wireless technology, with no degradation in performance, no dead zones, no interference between users, and no reduction in data rate as more users are added.

DIDO works indoor/outdoor for urban/suburban applications at distances of several miles, and for rural applications, DIDO works at distances up to 250 miles. Urban/suburban latency is sub-millisecond.

This paper describes how DIDO is dramatically different than conventional wireless technology, how DIDO works, what we have running so far, and the mind-blowing applications DIDO makes possible.

We believe that DIDO wireless will completely transform the world of communications and far more.

2. Background

Whenever there is more than one wireless transmission within range of another, there is the potential for interference between them. You've probably experienced this when listening to AM or FM radio while driving a long distance: as you drive out of the range of one station, you may begin to hear

[1] By "data rate" we mean bits/second (e.g. 10 Megabits/second or 1 Gigabit/second). The term "data bandwidth" or "bandwidth" is the more common term used to mean the same thing in data communications, but because "bandwidth" in the wireless world can also mean the width of the spectrum used in a transmission, we use the term "data rate" in this paper to avoid confusion.

another station that you are approaching. The two stations might duel it out for a bit, where you alternate between one and the other, but at some point you only hear the new station as you drive beyond the range of the first station.

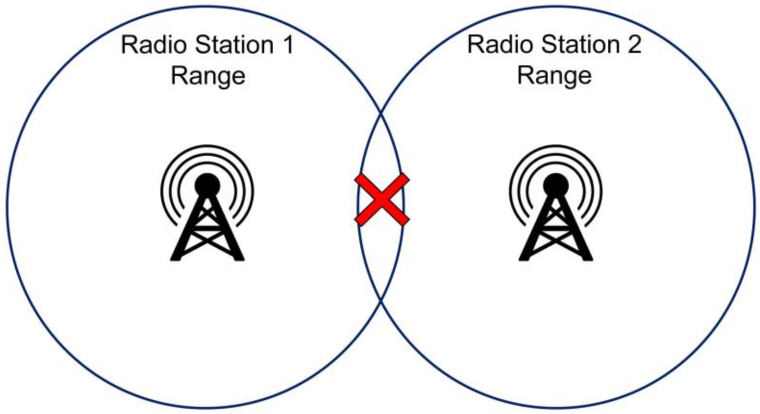

Of course, such radio station interference situations are rare because commercial stations at the same frequency are carefully located far apart with limited transmission power and each is broadcasting one signal to many users. So, during the broadcast era of the previous century, spacing out radio and TV stations and limiting their power output was a good solution to address wireless interference.

With the introduction of cellular phones, the interference situation became far more complicated. Now, with a large number of users simultaneously receiving and transmitting independent communications, it was necessary to subdivide each region into cells, typically a mile or less in diameter, with the cellular towers laid out such that adjacent cells were at different frequencies that did not interfere with each other. As a mobile phone user moved from one cell's range to another's, it would switch frequencies and "hand off" to the new cell. So long as the users stayed within the range of at least one cell tower, they would not lose a connection.

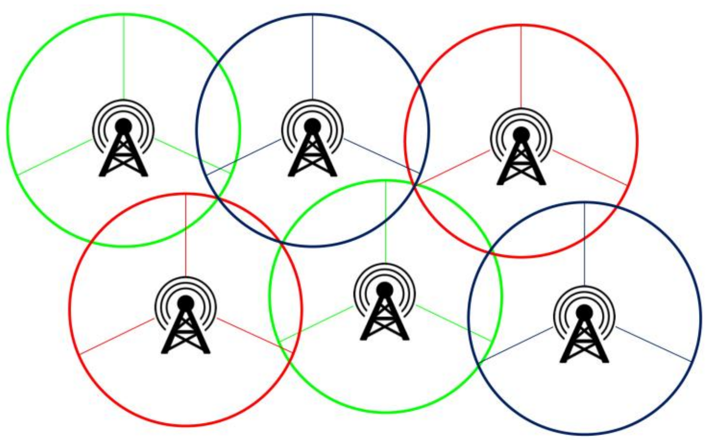

But as mobile usage grew exponentially, both in number of users and the amount of data, it began to tax the cellular system. Improvements were made, such as dividing cells like a pie into sectors, and utilizing more advanced digital transmission techniques such as MIMO. But, despite these advances, a given cell sector had a limited data rate capacity, shared by all users in that cell sector. And, once that data rate capacity was reached, users would not get the data rate they needed, resulting in dropped calls and slow and/or unreliable Internet communications. Although more spectrum and more cellular towers are being added to alleviate the situation, the demand for higher data rates is growing at an even faster pace, particularly as wireless video is becoming the majority of wireless data traffic.

There also has been exponential growth of consumer wireless systems such as Wi-Fi. Unlike commercial cellular systems where placement of towers, sectorization, power and frequency use is carefully planned, most consumer Wi-Fi access points are simply connected to the Internet at arbitrary locations, and then operate largely independently. Wi-Fi is restricted to very low transmission power to limit its range of interference, but nonetheless if there are other Wi-Fi access points (or other users of the same spectrum such a Bluetooth devices) within range, the Wi-Fi access points must do the best they can to co-exist and share data rate among all users. As there are increasingly more Wi-Fi access points, particularly in densely-populated areas, like apartments or offices in cities, and as increasingly more of the traffic is high-definition video, users are unable to consistently get the data rate they need, resulting in unreliable or lower-quality Web, video or videogame experiences.

The problem is mitigated somewhat by more advanced transmission techniques, such as MIMO and beamforming, but in practice, they have not resulted in a consistently reliable experience for applications like video and gaming. Adding more usable spectrum, such as has been proposed in using "White Spaces" would reduce congestion for a time. But, as data capacity needs continue to grow exponentially, and applications like video and videogames require high reliability, it is unclear whether there will be enough usable spectrum to meet user needs. And, once all of the usable spectrum has been allocated to consumer devices, it is effectively impossible to recover that spectrum as new technologies evolve that can use it more efficiently. We run the risk of allocating all of the usable spectrum, but still not having enough.

3. Spectrum capacity and Shannon's Law

Ultimately, the reason we are limited in multiuser spectrum capacity using conventional wireless techniques is due to Shannon's Law, which defines the maximum amount of error-free data rate that can be transmitted through a single communications channel for given amount of spectrum and noise level (the "spectrum capacity" of the channel).

For example, a cellular sector (which may span a kilometer) is essentially a single channel whose maximum data rate is limited by Shannon's Law, and all of the users within that cell sector share that data capacity. As another example, all of the users in overlapping Wi-Fi networks at the same frequency share the Shannon-limited data capacity.

Approaches like MIMO (e.g. 802.11n and LTE) or beamforming can increase the shared data rate limit by a factor of about 3X to 4X in certain situations, but these incremental improvements are overwhelmed by the exponential growth of simultaneous high-data rate users we have seen in the last decade, resulting in increasingly more users sharing limited channel data rate.

As more users in a given area share the same wireless spectrum, the data rate per user declines. So, with conventional wireless systems, Shannon's Law results in a reduction in per-user data rate as more users are added in a given area.

4. What is DIDO wireless?

Distributed-Input-Distributed-Output (DIDO) wireless technology is a new approach to multiuser wireless that allows the number and density of users in the same area to be steadily increased without additional users reducing the data rate of others. In other words, the shared spectrum capacity is not subject to Shannon's law: as more users in a given area share the same wireless spectrum, the data rate per user does not decline. As a result, regardless of how many users are in a given area, each user is able to use the entire Shannon Limit of the channel, despite sharing the same spectrum.

We do not know of a theoretical limitation to how many users we can add to a DIDO system without a degradation in data rate per user. There certainly will be practical limitations with each era of technology evolution, but we have not yet come close to them. So far, as we've increased the number of simultaneous users in the same area to 10 (limited just by the number of hand-built radios we have) we have not seen any degradation in performance. So, while our demonstrated spectral capacity today is 10X the Shannon Limit, we expect we can get to 100X, and are optimistic that 1000X is achievable. But, until we start to see some degradation in performance as we add more users, we will not be able to predict how far it can go.

DIDO wireless has many other advantages and useful properties beyond eliminating data rate reduction among users in shared spectrum. Some are described below after the explanation of how DIDO works, but other wireless behaviors unique to DIDO—including the most intriguing ones—are still being studied. We have observed wireless phenomena that we do not believe have been seen before that will take time to understand and document. Stay tuned…

5. How does DIDO wireless work?

We have found the easiest way to explain how DIDO works is to start with a very simplified description of a wireless system that we are all familiar with: consumer Wi-Fi. We will explain roughly how Wi-Fi works, and then show how DIDO works differently if the Wi-Fi access points are replaced with DIDO access points. Here goes:

a. A simplified example of multiuser Wi-Fi

When a single user (User 1, in blue) with a computer connects to a single Wi-Fi access point 1 (AP 1, in blue) and there are no other users of the spectrum (e.g. no other Wi-Fi, Bluetooth or other devices that share the same spectrum) in the area, if User 1 is within optimal reach (the blue circle) of the Wi-Fi AP 1, the wireless connection process is not only relatively simple, but it can be as reliable as a wired connection, and User 1 receives 100% of the data rate available in the channel.

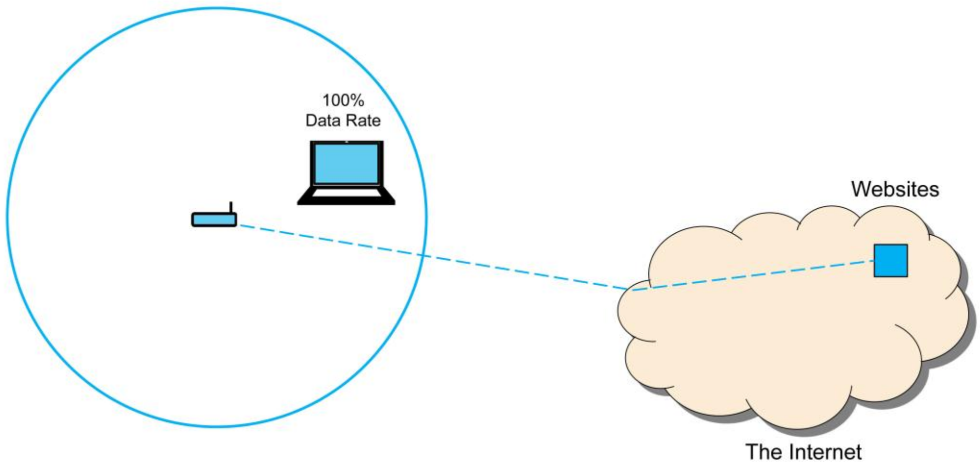

The following is a highly simplified description of how the user watches video from a website on their computer over Wi-Fi: User 1 clicks on a link on a website 1 to request a video stream, the website sends the data for that video to AP 1, then AP 1 modulates the data (i.e. creates a waveform that contains the data) as a radio signal and transmits that radio signal through the air. User 1's computer receives the radio signal, demodulates the data, and plays back the video.

So long as User 1 stays within optimal range of AP 1 and nothing else uses the spectrum, the video will play on uninterrupted, just as if it were going over a wire. Contrary to what we often experience in practice, wireless is actually highly reliable when there is a good connection and no other devices competing for the spectrum. But as soon as others devices start competing for the spectrum, then it becomes increasingly difficult to maintain a reliably high data rate connection.

For example, let's add User 2 (in green) using AP 2 (green, with green circle showing range), both within wireless range of User 1 and AP 1 (e.g. let's say User 1 and User 2 are next door to each other in an apartment building). If both AP 1 and AP2 are used simultaneously at the same frequency, then the Wi-Fi communications process starts out the same, but it ends up very differently:

User 1 clicks a video link on website 1 and User 2 clicks a video link on website 2, the video data from website 1 is sent to AP 1 and the video data from website 2 is sent to AP 2, each AP modulates a radio signal with the video data and tries to transmit that radio signal. But…it's not so simple this time. Each AP quickly discovers there is something interfering with its transmission: The AP using the same spectrum. Consider the following figure.

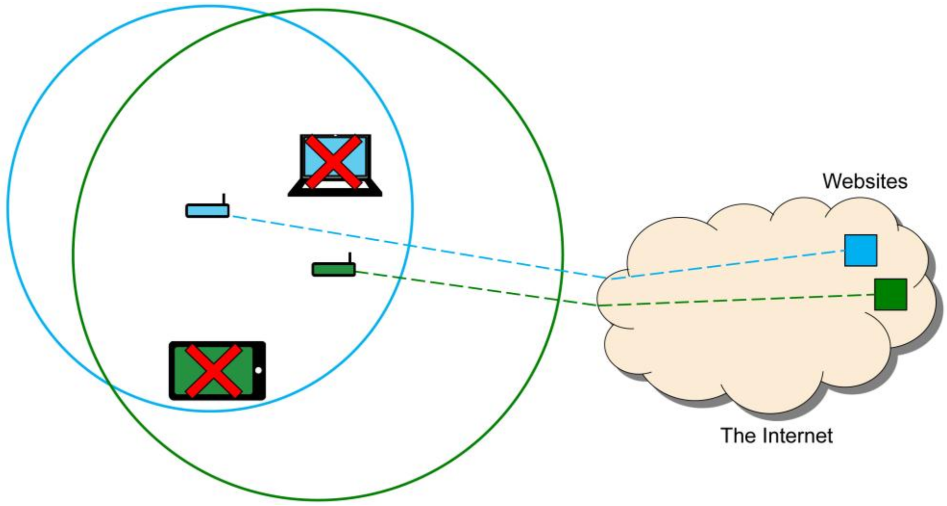

Interference caused by the 2 simultaneous Wi-Fi transmissions is shown by a red "X". Wi-Fi has several techniques to deal with interference, but the net effect is that they all share the channel, and as a result, reduce each user's data rate. For the sake of explanation, we'll show how very simplified spectrum sharing works, with each user taking turns. First User 1 uses the spectrum while User 2 stands by idle:

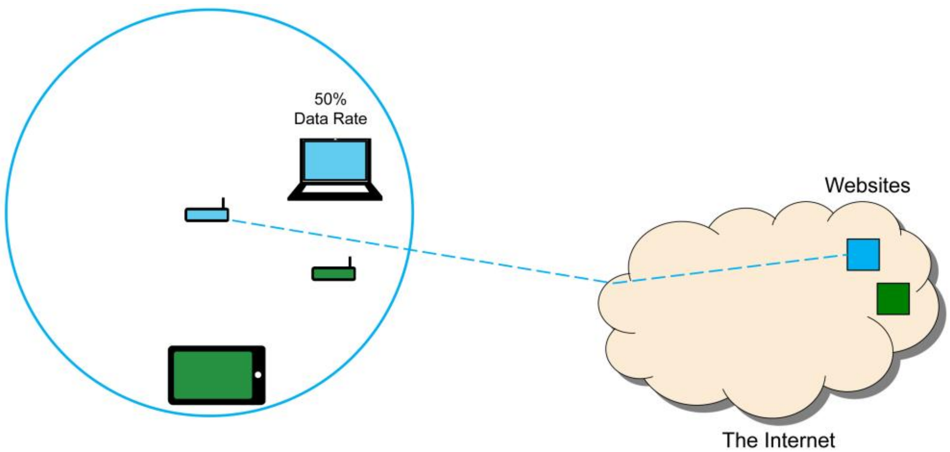

Then User 2 uses the spectrum while User 1 stands by idle:

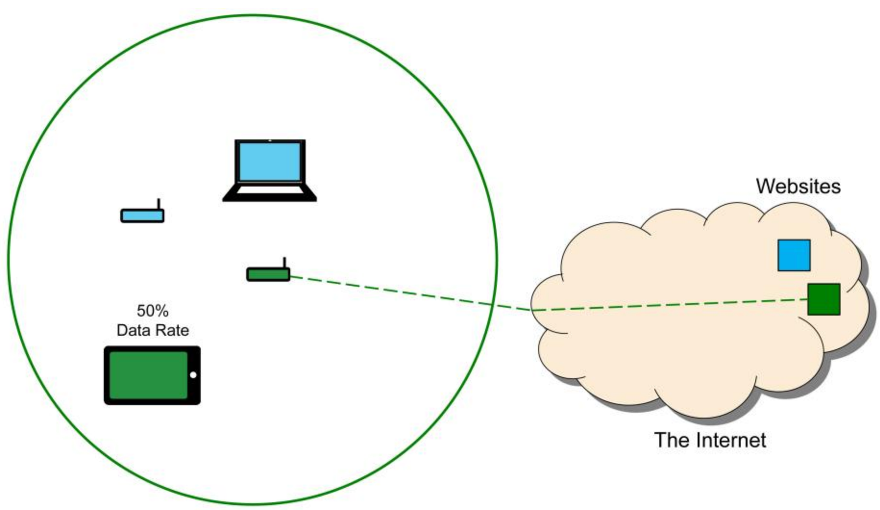

If we assume the protocol divides up the spectrum evenly between the two APs, then each user optimally will get 50% of the available bandwidth. So, when two Wi-Fi users are in range, each user loses half of the data rate they would have had alone. As we add more users and APs at the same frequency within range of Users 1 and 2 and APs 1 and 2, then we see a steady decline in available data rate per user.

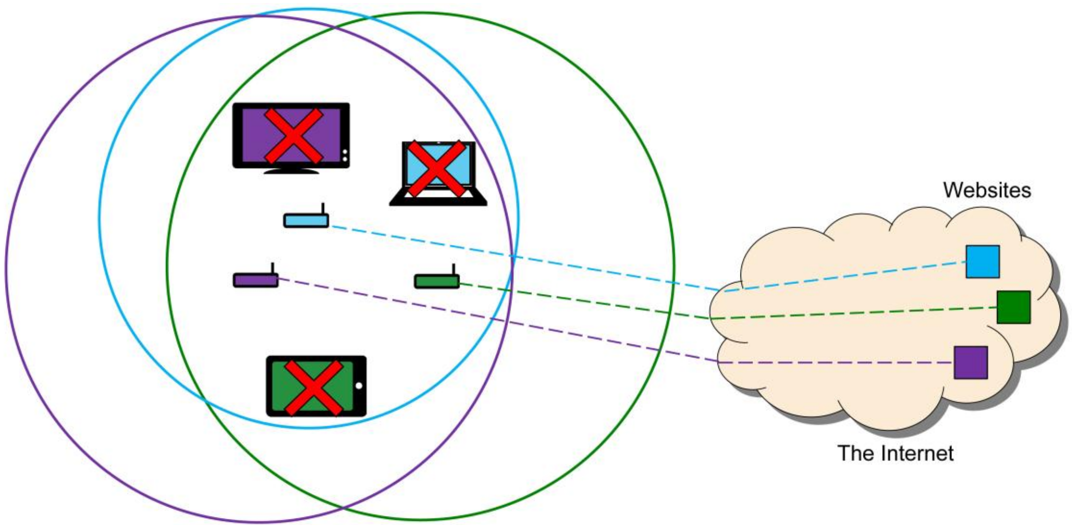

For example, if we add User 3 and AP 3 (purple) near Users 1 and 2 and APs 1 and 2, using the same spectrum, they will all interfere, and so, they will need to each take turns, and each will only get 33% of the data rate. This is illustrated on the next page.

The following three Figures show how each User and AP takes a turn, as the other two Users and APs sit idle:

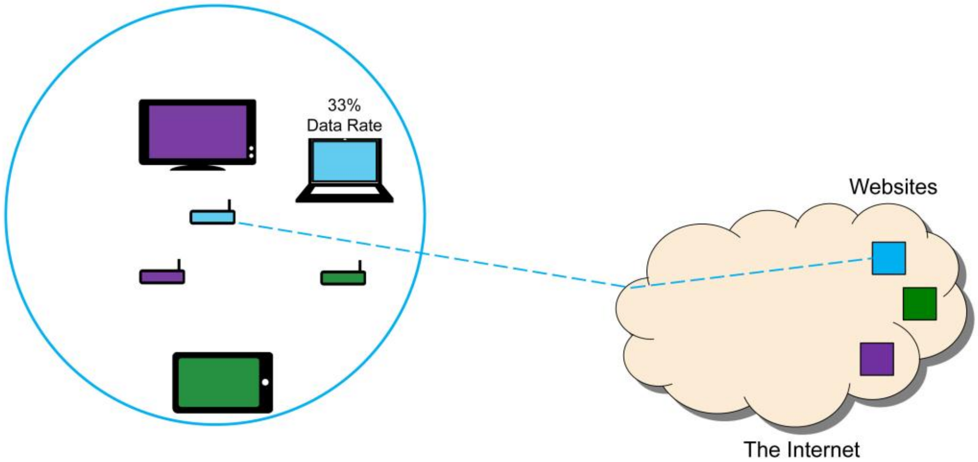

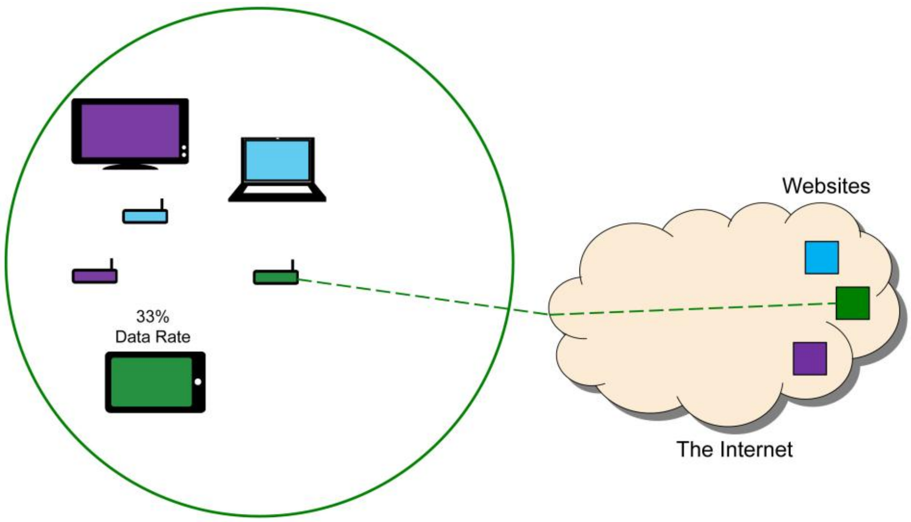

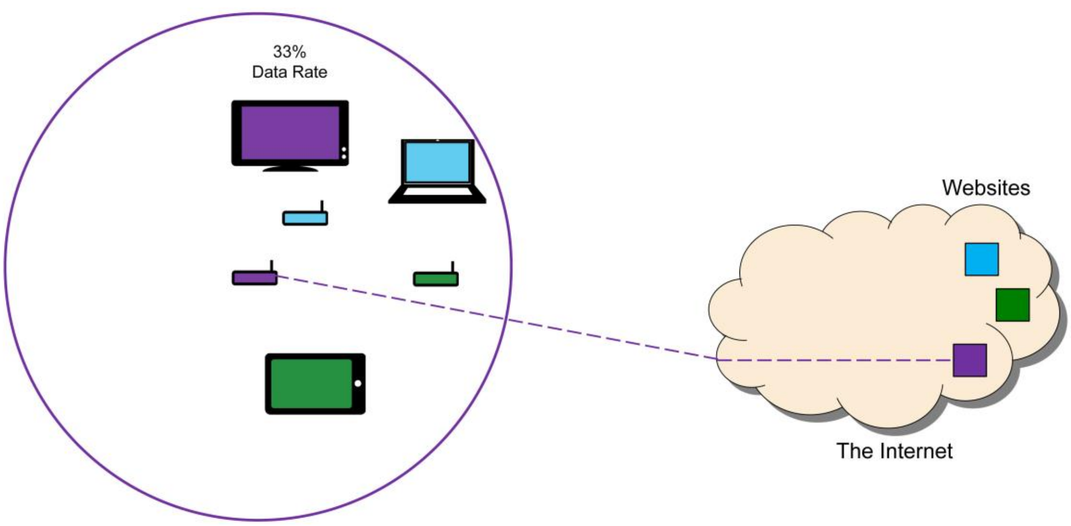

And, the problem only grows worse with more users. In a city building, it is common to see 25 Wi-Fi networks or more listed on a user device. If all of the networks are sharing the same spectrum optimally, each user will only get $1/25^{th}$ (4%) of the data rate. As users are added, the data rate per user steadily decreases.

The reason that this happens is all of the users and APs are sharing the same channel, whose data rate is limited by Shannon's Law. So, if there are 10 users, then optimally, the data rate per user is the Shannon Limit divided by 10. If there are 100 users, then it is the Shannon Limit divided by 100, and so on. (Again, this is a highly simplified description of how Wi-Fi works, but even so, this is a rough estimate of the expected loss of data rate per user as more simultaneous users are added.)

b. A simplified explanation of DIDO

Let's walk through the same example, but rather than use Wi-Fi access points, we will use DIDO access points, starting with User 1 and AP 1. Here goes:

When User 1 (in blue) clicks on a link on a website 1 to request a video stream, the video data from website 1 is NOT sent to AP 1 (in gray—you'll see why all the APs are a neutral color soon). Rather it is sent to a DIDO Data Center (in gray) that supports AP1 and User 1. The DIDO Data Center then processes the data, modulates it into a radio signal waveform and sends the waveform to AP 1, which simply sends the waveform to its antenna and transmits it as a radio signal. User 1's computer then receives the radio signal (small blue circle—why it is small is explained below), demodulates the waveform into data and plays back the video. So, quite differently than Wi-Fi, where AP 1 receives data directly from website 1 and modulates the data into a radio signal waveform, with DIDO, the radio signal waveform is computed in a DIDO Data Center, and all the DIDO AP does is transmit the radio signal waveform.

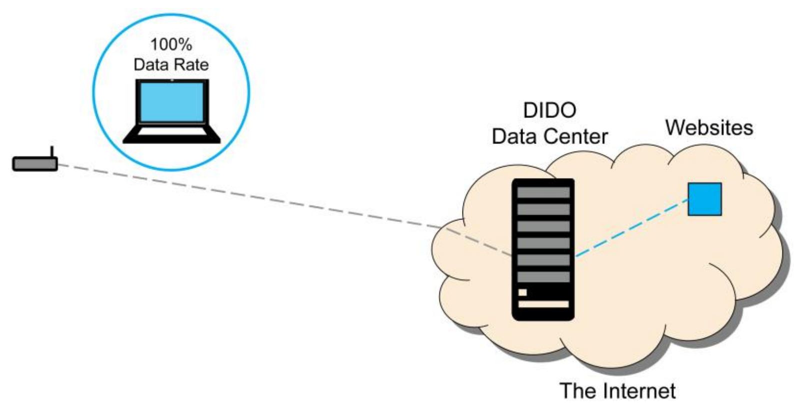

Let's now add in User 2 (green) and AP 2 (also gray—they are all gray because it doesn't matter which user owns which AP). This is where it gets interesting.

When User 1 and User 2 each click on video links, respectively, on website 1 and 2, both websites send the video data to the DIDO Data Center. The data from the two websites is combined and processed in the DIDO Data Center to result in two radio signal waveforms (each quite different than the other), and one waveform is sent through the Internet to AP 1 and the other waveform sent to AP 2. Then—and this is very important—AP 1 and AP 2 both transmit their respective radio signal waveforms simultaneously. Because both APs are transmitting at once, the 2 radio signals collide at both User 1's computer and at User 2's computer, resulting in each computer receiving the sum of the two radio signals at each computer's particular location in space.

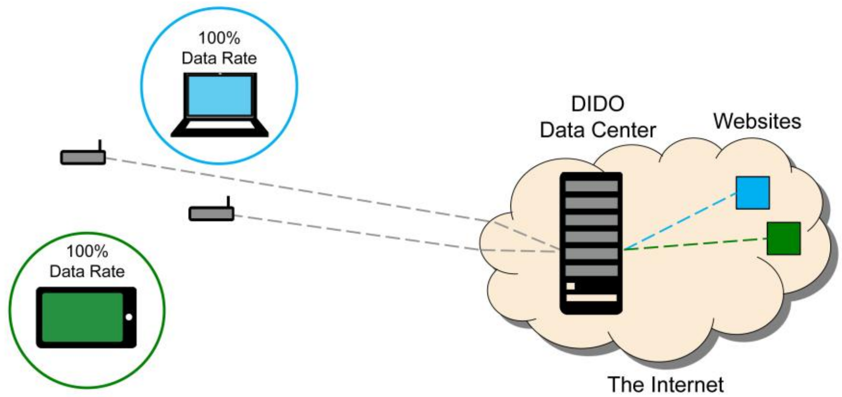

Normally, when you have two very different radio signals colliding with each other, the result is one stronger signal overpowering the other (as in the radio station example given in the Background section), or just indecipherable noise from the two signals interfering with each other.

But, not with DIDO. Instead something rather remarkable happens: the sum of the radio signals at each computer's location results in a clean modulated waveform carrying the data intended for that particular computer. So, the waveform received at User 1's computer contains the video data sent by website 1, and the waveform received at User 2's computer contains the video data sent by website 2. Each computer simply demodulates the waveform and plays the video for its user.

And, here's the really amazing part: what each user receives is what they would have received if they had the channel to themselves, without another user sharing the same spectrum. There is no interference from the other user. Each user is able to utilize the full Shannon Limit of the channel.

If we now go up to 3 users and 3 APs, the same process occurs. All 3 users click on links on 3 websites, and all 3 websites send their data to the DIDO Data Center. The DIDO Data Center processes the data and creates 3 distinct waveforms and sends them to APs 1 through 3. All 3 APs transmit their respective waveforms simultaneously, and all 3 sum together at each of the 3 computers of the 3 users. And, the resulting sum of the 3 radio signals at each of the 3 computers is a clean modulated waveform carrying the data intended for that computer. So, User 1 receives the data from website 1, User 2 from website 2, User 3 from website 3.

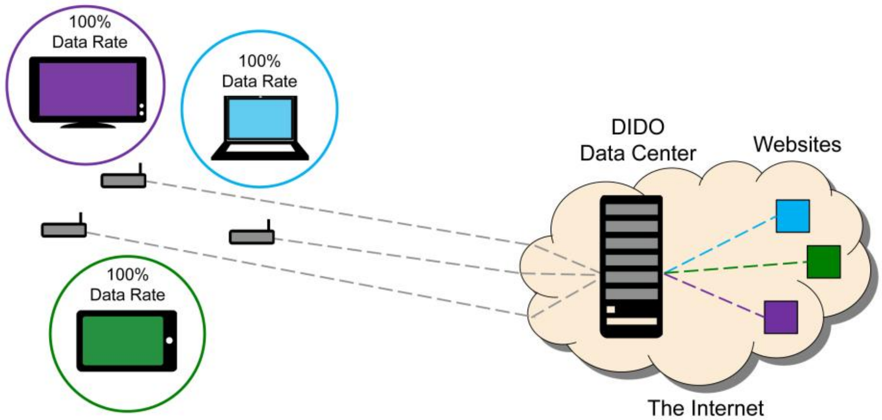

Each user receives what they would have received if they had the channel to themselves, without any other user sharing the same spectrum. There is no interference from the other users, so each user is able to utilize the full Shannon Limit of the channel. Pretty cool.

The exact same process applies to 5 Users and 5 APs, connecting to 5 different websites, as shown in the figure below. There is no interference among the 5 Users, and all 5 Users get the benefit of 100% of the data rate of the channel, and it doesn't matter where the APs are located or which user owns which. Each user gets the data from the website they are connected to through an independent wireless channel.

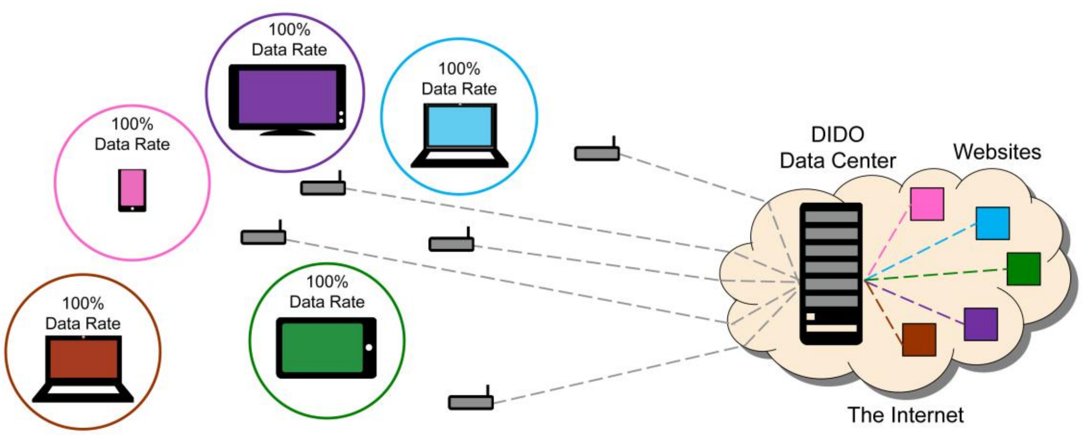

Today, we have 10 Users and 10 DIDO APs working (limited only by how many custom radios we've built so far), and they work extremely well, with no degradation in performance as more users have been added.

While we've focused on how the downstream (from the Internet to the user) process works with DIDO, DIDO is just as efficient in the upstream (from the User to the Internet). All users transmit simultaneously, and all get the full upstream data capacity of the channel. It's more complex to visualize or explain, but essentially the inverse process occurs. Two-way wireless systems are typically TDD (shared up/down channel) or FDD (separate up/down channels), and DIDO works fine in either mode.

There are other benefits from using DIDO that go beyond what can be explained simply, but as some examples, the signal quality and reliability is much higher than that of Wi-Fi or cellular, and for the same range, in general, less wireless power is needed. Also, DIDO was carefully designed so that both the APs and the radios inside of the user devices (whether computers, mobile device or other consumer electronics) are much simpler than Wi-Fi or cellular radios.

So, the big question of course is, how can something like this work? Read on.

6. How is DIDO able to create waveforms that, when they sum together at each user results a clean waveform with the data for that user?

The complete answer to this question is very long, involving immensely complex mathematics, very carefully designed software and hardware, and new data communications and modulation techniques. The following is a highly simplified explanation:

DIDO is a cloud wireless system. All of the intelligence of the DIDO system is in a DIDO Data Center, which then communicates to all of the users at once through all of the APs at once. So, you can think of the DIDO APs as a vast random array of antennas extending out from the DIDO Data Center for miles, but instead of running long wires from the Data Center to the antennas, DIDO uses the Internet to connect to each DIDO AP, allowing each DIDO AP to be placed anywhere there is an Internet connection, whether indoor or outdoor, much like a Wi-Fi AP could be placed anywhere there is an Internet connection.

DIDO communication begins with the DIDO APs exchanging brief test signals with the DIDO user devices. By analyzing what happened to these test signals as they propagate through the wireless links, the DIDO Data Center determines precisely what will happen when it transmits data signals from the APs to users, and how the simultaneously transmitted signals will sum together when received by each user device. Then, the DIDO Data Center uses this analysis, along with the data each user is requesting (e.g. video from a website), to create precise waveforms for all of the APs that, when transmitted at once will sum together at each user device to create a clean, independent waveform carrying the data requested by that user. So, if there are 10 APs and 10 users all within range of each other, then 10 radio signals will sum together at each antenna of each user's device to produce an independent waveform for each device with only that device's data.

It is important to note that DIDO was designed to be a mass-market product, not a lab experiment. As a result, many of design decisions were specifically made to enable DIDO to be built practically and inexpensively, and to scale to any size. For example, much of the work has been focused on making sure the DIDO Data Center computing requirements scale linearly, not exponentially, as more users are added. DIDO also uses very advanced mathematics, not switching, to enable users to move among adjacent DIDO Clusters (the DIDO APs controlled by each Data Center make up a DIDO Cluster) seamlessly without requiring any "handoff", allowing adjacent Clusters to share the same frequency.

7. With so many users sharing the same spectrum at once, why doesn't Shannon's Law result in the reduction of the data rate per user as more users are added?

The reason is Shannon's Law is not about *spectrum* data rate limits, it is about *channel* data rate limits. We are used to thinking of spectrum and a channel as the same thing, because in conventional multiuser wireless system such as Wi-Fi or cellular, when you have a high density of users in the same area (e.g. an apartment building or a cell sector), all the users are sharing the same spectrum *and* the same channel. So, as you add users, Shannon's Law applies because they are all sharing the same channel, and it is incidental they are all sharing the same spectrum. The reason techniques like MIMO and beamforming can increase data rate by 3X-4X using the same spectrum is they are able to create a few (perhaps 3 or 4) independent channels in a densely-shared area, and indeed, if the circumstances allow for it, they are able to achieve 3X-4X the Shannon Limit in shared spectrum.

DIDO is a general solution that creates an independent channel for each user, even in densely-shared areas. Since each user has an independent channel, Shannon's Law does not apply, despite the fact that all users are sharing the same spectrum.

The reason why DIDO channels are independent is very hard to understand the first time you hear it, so let's start with a different wireless arrangement that is not as general as DIDO, but is easier to comprehend.

Let's say you had 10 highly directional antennas (e.g. dish antennas) spread out in a line and 10 users spread out in a line parallel to the dish antennas. If each dish antennas is aimed directly at user in the same position in the other line (i.e. all 10 beams are parallel and do not overlap), then there will be 10 beams transmitted to 10 users, and each will be an independent channel that does not interfere with the other channels. Suppose that instead of keeping all 10 beams parallel you pair up the dish antennas and users randomly, so that each dish antenna only points at 1 user, but now the beams are all crossing each other. We still have 10 independent channels because, despite the fact the beams cross each other in space, they do not cross each other where each is received by a user, and as a result do not interfere (e.g. just as theater spotlights that cross in the air, but hit the stage at different points do not interfere). So, in this very specific arrangement, with 10 antennas and 10 users, we can have 10 independent channels in the same area. This tells us the concept of 10 simultaneous channels is physically *possible*. Unfortunately, in practice with a typical multiuser wireless arrangement, the antennas and users are distributed randomly both among and around each other, with walls and other obstacles all about. Even if we very carefully aim the dish antennas, even trying to bounce the radio signals off walls, we'd find very few completely independent paths between antennas and users, and as a result, in typical wireless networks we can't come anywhere close to creating 10 independent channels in an area. As a result, users need to share channels and Shannon's Law applies, limiting the data rate per user as more users are added.

Unlike the dish antenna example, DIDO has almost no restrictions as to the arrangement of antennas and users, walls or other obstacles. In fact, the more widely *distributed* the DIDO APs and more widely *distributed* the DIDO users, the better. Hence the name: "Distributed-Input-Distributed-Output". In wireless parlance, widely distributing both AP and user antennas achieves exceptionally high "diversity". Diversity makes each antenna "statistically independent" (its signal paths are different than other antenna signal paths), which is how the DIDO Data Center distinguishes each APs signal from those of the many APs that reach a given user. This allows the DIDO Data Center to figure out precisely what waveforms it needs to generate so that all the waveforms sum together into a clean waveform for each user. Each of these clean waveforms is an independent channel.

So, as more users sharing the same spectrum are added to a DIDO network, Shannon's Law does not apply. Equally importantly, DIDO achieves this goal in a completely general wireless system: APs are located where it is convenient to place them, and users can be located anywhere.

8. A directional antenna channel can be visualized as a beam, like a theater spotlight. How does one visualize a DIDO channel?

A DIDO channel can be roughly visualized as a 3D sphere that surrounds the antenna of a user device. We call it the "area of coherence". Within that sphere, the channel for that particular user exists. Outside of that sphere, it does not. So, if there are 10 users in a DIDO network, there are 10 spheres, one around each user device. As a user device moves, the sphere moves with it.

Strictly speaking, the area of coherence is not exactly shaped like a sphere, especially when there are only a few APs (and strictly speaking, a wireless beam is not exactly shaped like a spotlight, especially when there are only a few antennas), but the analogy is good enough for visualizing what is going on in a practical DIDO deployment with many APs.

9. At what frequencies does DIDO operate? What is DIDO's range?

We've tested DIDO at frequencies from 1 MHz to 1 GHz, and it works great at all frequencies. We are confident it will work well at higher frequencies as well. As with any wireless system, different frequencies exhibit different properties such as ability to efficiently penetrate certain materials or ability to reflect off the ionosphere.

In practice we've found that DIDO systems are much more reliable than conventional wireless systems because of the enormous diversity achieved by having many distributed antennas. For example, a Wi-Fi or cellular network may have dead zones because of some obstacles that block the path from a Wi-Fi AP or cellular tower to the user (e.g. when you walk around the corner of a building and your call is suddenly cut off). With DIDO, because the antennas are scattered about, even if some antennas are obstructed from reaching a user, it is likely that some other antennas will reach the user without a problem.

Also, DIDO APs are able to transmit far longer distances than Wi-Fi APs or cellular towers. The reason is they can transmit at higher power, if necessary, without the concerns that Wi-Fi or cellular has of overlapping APs or cells at the same frequency and causing too much interference. In fact, with DIDO, the more overlap the better. As a result, indoor DIDO APs can (as needed) safely transmit a mile or more, and outdoor APs can safely transmit much farther.

One particular DIDO configuration that with a range of approximately 250 miles in radius (500 in diameter) is what we refer to as "DIDO Rural". Certain HF frequencies (e.g. 3-7 MHz) will bounce off the ionosphere to allow transmission beyond the curvature of the Earth (or over mountains or into canyons). It is a transmission property known as "Skywave". Transmitting straight up is called "Near-Vertical Incidence Skywave" (NVIS) and will result in the reflected signal coming back to Earth to cover an area of about 500 miles in diameter. The following figure illustrates how groundwave transmissions can cover about 40 miles in diameter before being blocked by the curvature of the Earth, but NVIS transmissions are not blocked by the curvature of the Earth at all, allowing for about 500 miles of coverage.

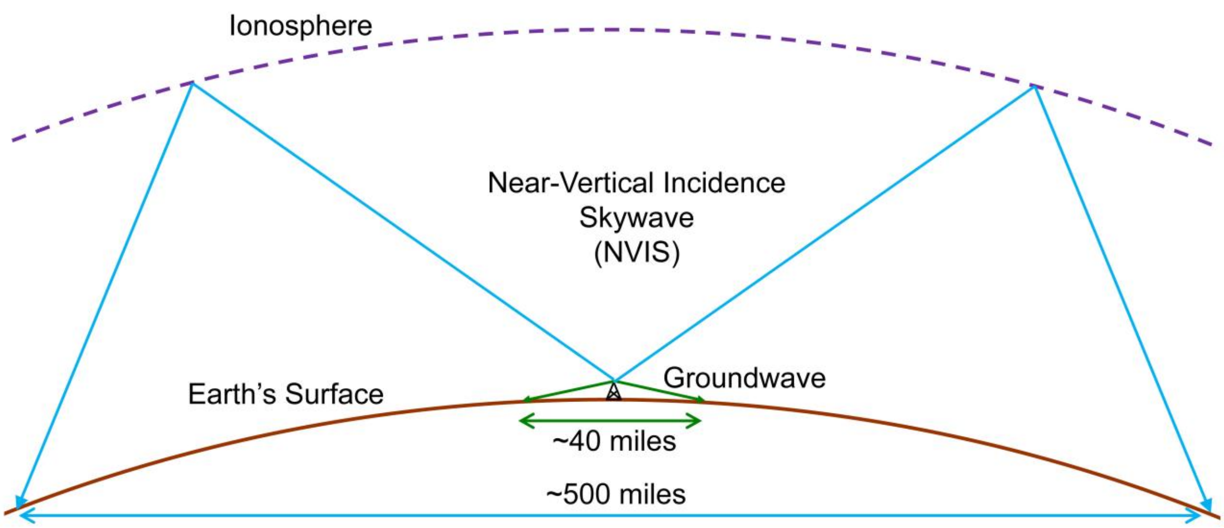

A key limitation for data communications using the HF bands that work with NVIS is that they are very narrow and a large percentage of the HF band would have to be used to provide sufficient data rate for broadband applications to even one user in the entire 500-mile area. Of course, DIDO would be an ideal technology to overcome this limitation because a large number of users can simultaneously use all of the available spectrum, each receiving the full data rate available in the HF channel.

We have successfully built a DIDO-NVIS system that achieves exactly this result. From 3 transmitting antennas, up to 32.9 miles away and beyond the curvature of the Earth, we have simultaneously delivered independent signals to 3 user antennas that are as close to each other as typical rural homes.

The 3 APs were located in Pflugerville, Lake Austin and Austin, Texas, all transmitting on the same frequency and connected by consumer Internet connections to a DIDO Data Center.

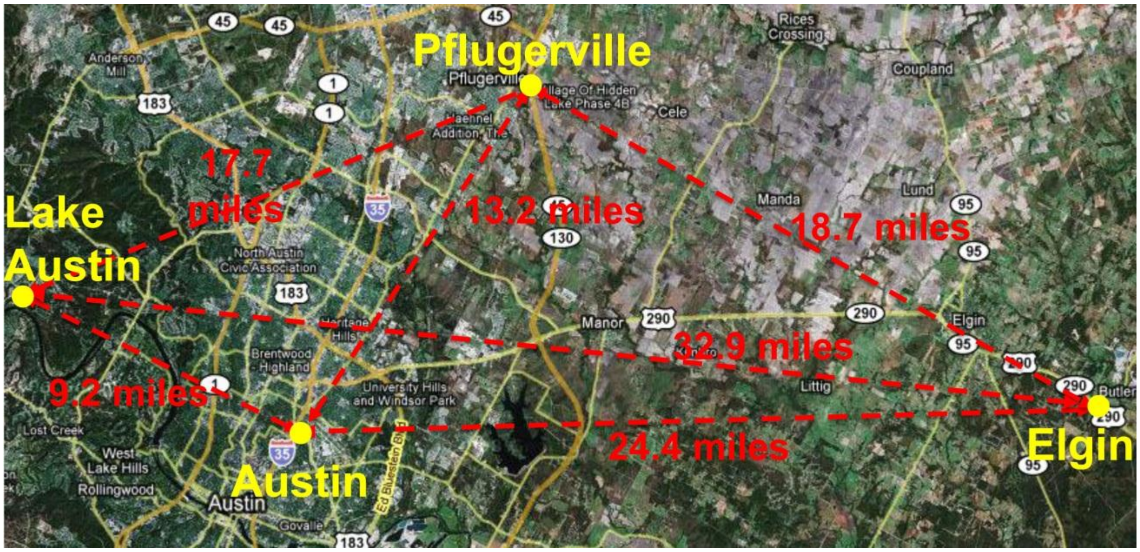

3 rural Users near one another in Elgin, Texas, each received data on independent channels:

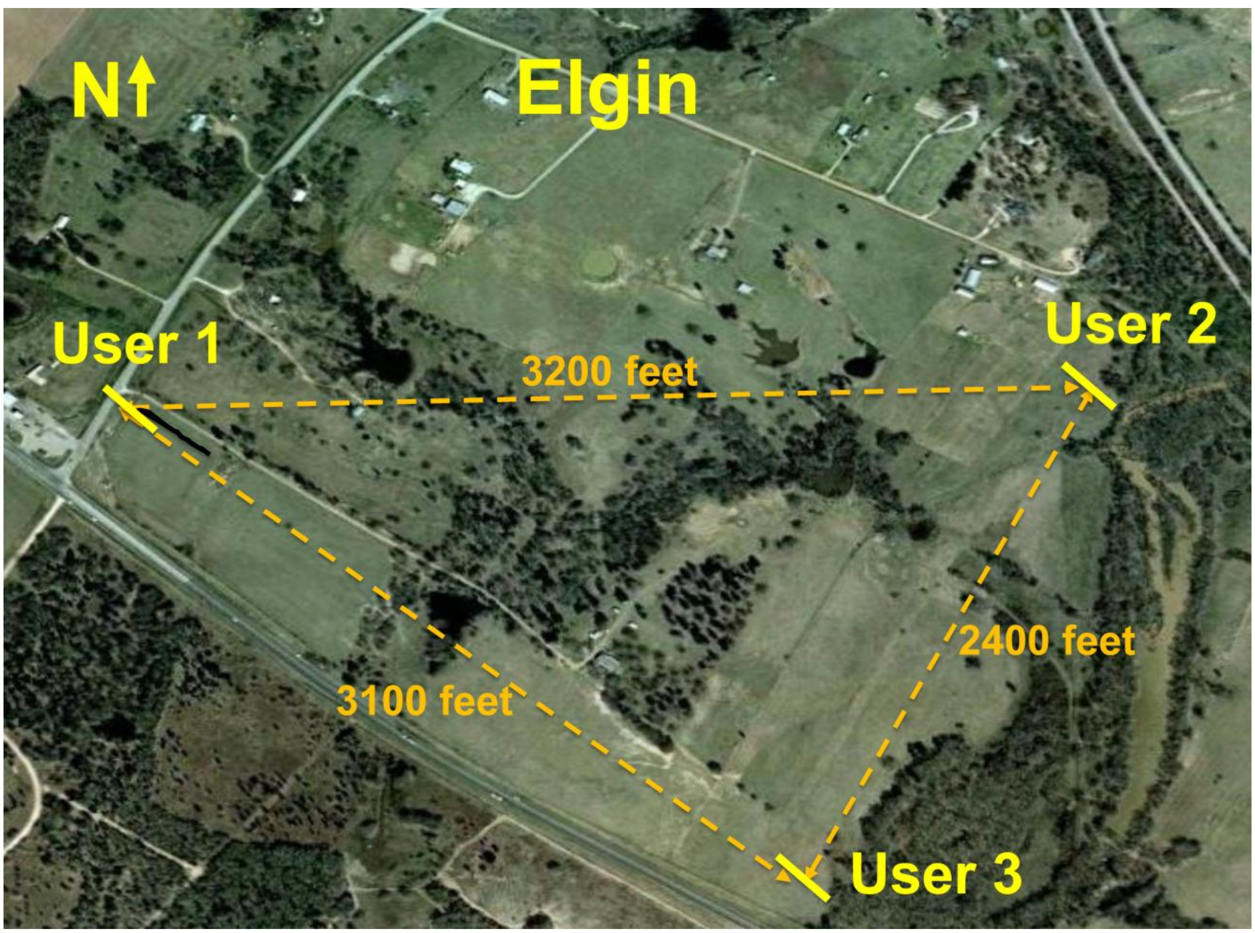

We stopped at 3 users simply because of the time required to set up the experiments and travel between test sites (and in all seriousness, we had issues with cows chewing through cables and amorous bulls knocking over antenna masts), but the wireless characteristics we saw were the same that we were seeing in our urban/suburban experiments at higher frequencies, and we expect we'd be able to expand the number of simultaneous users accordingly. As such, DIDO-NVIS provides a very inexpensive and efficient way to deliver broadband to rural areas, or remote areas in the wilderness.

10. How does DIDO compare to conventional wireless systems in terms of latency (lag)? How does DIDO relate to OnLive?

Cellular systems have long range, but have very high latency (up to 150 milliseconds or more for 3G). Wi-Fi systems are fairly low latency (a few milliseconds), but are limited in range. There is no inherent reason for long-range wireless networks to be higher latency than wired or fiber networks, but the complex wireless protocols and distribution systems used in multiuser wireless systems add a great deal of latency. Even LTE (4G) incurs 40ms or more of latency.

DIDO is extremely low latency, whether used short range or long range, because it can provide an independent channel for each user without having to rely on complex sharing protocols. In urban and suburban applications, DIDO will typically incur less than 1 millisecond of latency. DIDO Rural (using NVIS) incurs about 2-3 milliseconds of latency. As such, DIDO is exceptionally well-suited to support OnLive (www.onlive.com), a cloud gaming/cloud application technology that hosts videogames and applications remotely in data centers, and then using patented low-latency video compression, streams the video from the games or applications to the user's device. The round trip delay from the moment a user presses a button or touches a screen, to the point where OnLive updates the user's screen with video streamed from the OnLive Data Center, is so short that perceptually it seems like the game or application is running locally. While OnLive works over conventional Wi-Fi and cellular wireless, DIDO not only provides exceptionally low latency for OnLive, but provides for much higher (and much more reliable) bandwidth than convention wireless systems.

Another advantage of combining DIDO and OnLive is that they both are cloud-based systems, hosted in data centers. When OnLive servers are co-located with DIDO servers, there is unprecedented low latency, to the point where remote-hosted OnLive gaming and computing is lower latency than the same games and applications running on a local device. This is due to the much faster response achievable in high-performance servers in a data center than in a local device, and the minimal latency incurred when using DIDO for transmitting the video. Effectively, a DIDO/OnLive Data Center would deliver DIDO wireless waveforms directly from videogame and application servers. The latency would be so low and the wireless reliability so high, it would be creating a ubiquitous entertainment and computing resource that is available wirelessly everywhere to any device. User devices would not even require local computing resources because they'd always be connected everywhere (even in rural areas or on airplanes) to OnLive's high-performance servers. Essentially, local devices would simply be used for display/sound and user input. (Eventually—and unlike everything else in this white paper, this remains to be demonstrated, but I'm confident it is achievable—there will be no physical display screen, speakers or input device, but rather we'll see, hear and interact with a panoramic 3D experience that can be all around us whenever and wherever we want it to be, all created and transmitted reliably and instantly from the cloud.)

DIDO, OnLive and MOVA (a company that captures photorealistic human facial and body motion) were all developed together in the Rearden (www.rearden.com) incubator in what is now a decade-long development effort. The unveiling of DIDO is the latest piece in the puzzle (there are more coming…), enabling the delivery of constantly available, instantaneous, photorealistic entertainment, computing and communications through a seamless merger of cloud computing, cloud gaming and cloud wireless.

As an example, Geni4 is a photorealistic computer-generated character created by MOVA Contour, hosted as a live character by OnLive, and can be delivered reliably throughout the world through DIDO wireless technology. Geni4 is a computer-generated character that is far too complex to be hosted by a local computing device, and is only possible through OnLive cloud gaming technology delivered through live, low-latency video streams. Interactive characters this real will be commonplace on all devices, whether in the home, office or mobile. And, DIDO wireless will provide enough spectrum capacity to support as many simultaneous, low-latency live video streams as there are users.

11. How long did it take to develop DIDO wireless technology?

DIDO has been in development for about a decade. While the theory behind DIDO is immensely complex by itself, the practical implementation of a DIDO has been a matter of overcoming one challenge after another. A small team of dedicated engineers, scientists and mathematicians has worked day and night to bring DIDO technology to life.

We have only scratched the surface of the potential of this technology. We believe there are not only far more applications in communications, but we believe that the unprecedented control and capacity that we have with radio signals will lead to a wide range of applications in other fields, such as medicine, imaging, manufacturing and alternative energy. It seems that with every new prototype we get working, there are new phenomena and potential applications that are uncovered. It is incredibly exciting.

We believe DIDO wireless will completely transform the world.