RF planning for broadband wireless access networks

The need for broadband connectivity is serious. Before deciding on an approach to BWA, you need to know the issues.

By Jay Klein

With the growing importance of the Internet and data communications for business, demand for high-bandwidth data connections has skyrocketed.

One of the most viable ways to meet this bandwidth demand, particularly for small and mid-sized businesses and residential areas—and especially for the local loop (or "last mile" as it is often called)—is through fixed wireless connectivity. Fixed wireless makes it possible for carriers to provide broadband wireless without relying on existing telecommunications infrastructure. It also has a lower cost of entry, with much faster deployment times than wired systems. This is important in servicing the final destinations of telecommunications infrastructures where the cost of service can often be exorbitant.

There are a number of approaches to providing broadband wireless access (BWA), and at least as many issues as well. One of the foremost issues is that of frequency reuse.

In the United States, the spectrum range of 24 to 40 GHz is being used for local multipoint distribution service (LMDS), and more broadly, BWA. Two system types are being employed to build out these networks: point-to-point (PTP) systems historically used for backhaul of telecom traffic, and a new generation of point-to-multipoint (PMP) systems that can cover entire service areas with a single central hub. Given the long history of PTP radios, RF planning for these deployments is straightforward and well-known, especially given the narrow beam antennas used in this technology.

Point to multipoint

PMP systems use wide beam antennas and are deployed in a cellular pattern to cover an entire city. Consisting of hundreds of customer premises equipment (CPEs) being serviced simultaneously by a single centralized radio (hub or base station), they present the RF engineer with a greater challenge in planning a fixed wireless network deployment, as seen in Figure 1. A PMP radio system is in contrast with point-to-point (PTP) radio systems, which have only two endpoints with communications occurring in dedicated channels between those two endpoints. For PMP systems, two-way communications occur between the hub and each CPE on a burst-by-burst basis.

Figure 1. A next-generation BWA point-to-multipoint network architecture.

As the deployments of PMP broadband wireless access systems evolve from field trials to commercial deployments, it is necessary to examine RF planning and frequency reuse issues associated with current BWA systems. Deploying service on frequencies allocated for LMDS in a real-world scenario is different from laboratory settings. With multiple base stations creating multiple cells, there are interference issues among base stations resulting from complex interactions. And the type of duplexing scheme used has been drawn into this debate. Some common misconceptions about these issues surrounding duplexing have resulted in misunderstandings about the complexity and robustness of each approach.

Duplexing time division and frequency division

Given its history of carrying analog traffic, wireless duplexing has been implemented traditionally by dedicating two distinct frequency bands: one for upstream transmissions and one for downstream. This technique is commonly referred to as frequency-division duplexing (FDD). The two bands are typically of equal size, given the symmetrical nature of telephone communications. The upstream and downstream bands are separated by a "guard band," an unused frequency block that can be large, relative to the upstream and downstream band sizes. The guard band provides the necessary up/down isolation to make a FDD system operate properly.

Duplexing to provide simultaneous, two-way communications services can also be performed in time—as in time-division duplexing (TDD)— rather than frequency. In this approach, transmit and receive operate on the same frequency, but at different times. Because of the ability to buffer traffic and the relative speed of switching between the two functions, simultaneous two-way communications are preserved.

TDD is an established technique that has been successfully deployed in communications systems throughout Japan and Europe. As the technology basis for personal handyphone service (PHS) in Japan and digital European cordless telecommunications (DECT), TDD has been commercially proven to give service providers a low-cost and low-power wireless platform for both base stations and CPEs. Shown to perform in the rigors of a mobile telephone environment, TDD is now making its way into fixed wireless access as a more flexible and cost-effective method of delivering data, voice, video and communications.

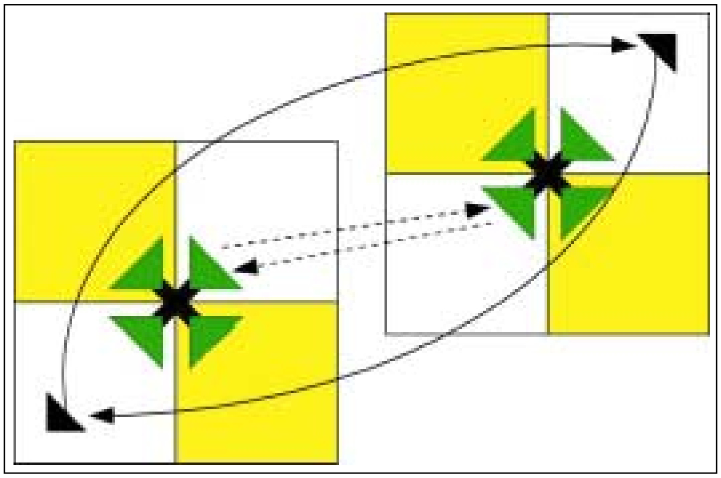
Figure 2. TDD-only interference issues: if frame-synchronized, these scenarios are eliminated.

The essential difference between TDD and FDD for RF planning considerations is that TDD relies on one channel for both upstream and downstream transmissions, dividing transmissions by time. FDD uses two channels, one for upstream and one for downstream transmissions, with a guard band between the frequencies to protect against interference between the two channels. FDD is a legacy approach based on analog technology, while TDD has been developed for digital transmission from the ground up. Critics have dismissed or attacked TDD based on its complexity and its deployment challenges, but a deeper analysis shows that TDD has clear advantages over FDD.

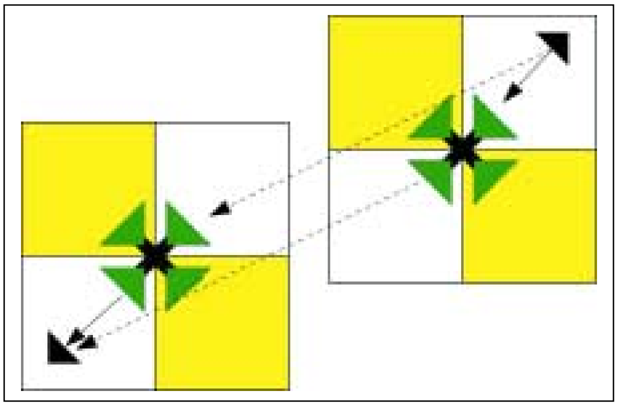
Figure 3. FDD and TDD interference scenarios: base station to CPE (downlink) and CPE to base station (uplink).

Since the inception of two-way, digital LMDS systems in 1995 and, more recently, the U.S. auctions of 28 and 31 GHz spectrum, there have been no major deployments of PMP networks by carriers. The reason for this delay is the high cost and low functionality contained in first-, and even second-, generation systems—carriers could not justify deploying on a mass scale. As a result, BWA carriers are using PTP systems for the majority of their build-outs. TDD is poised to increase the attractiveness of PMP systems, but the consideration of robust, widespread build-outs for PMP networks requires a look at the RF planning for these new systems and how it compares with the legacy FDD approach.

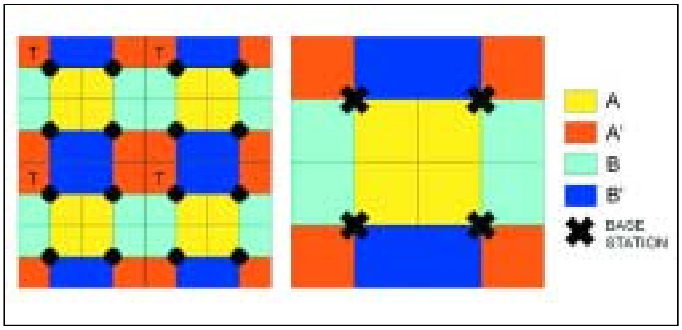

Figure 4. Frequency reuse of 2: Synchronizing frequencies to prevent interference.

The remainder of this article explores the RF planning considerations of deploying a BWA network, while considering the use of synchronization to reduce interference and improve robustness.

Addressing interference

The focus of RF planning for TDD and FDD systems is the avoidance of interference–unwanted signals that disrupt communications. The result of interference in BWA systems is errors and dropped packets. Ultimately, interference reduces the capacity of a network because bandwidth must be used to correct errors.

The two main categories of interference are adjacent-channel interference (ACI) and co-channel interference (CCI). ACI involves interference from signals at a different frequency on a nearby channel. ACI is primarily a power and separation issue. By adjusting the power output so that each hub and CPE sends signals only strong enough to reach the other, this interference can be reduced. Likewise, spatial separation, including front-to-back isolation, is used to reduce interference. As co-channel interference involves other channels at the same frequency, it can be much harder to eliminate. Due to the nature of PMP build-outs, these types of interference can be felt from both the other channels in the same hub and those from other hubs. This is why RF planning is so critical to a successful deployment.

One criticism often leveled against the TDD approach is that it allows for additional sources of interference not found with FDD deployments. Because the same channel is used for both upstream and downstream transmissions, it's possible for CPE to pick up a signal intended for a base station sent by other CPE. It's also possible for base stations to receive signals from other base stations meant for CPE (see Figure 2). FDD does not experience problems with interference between base stations or between CPE. Because different radios use different channels for sending and receiving, a

CPE can only “speak” to base stations, and vice-versa. Therefore, TDD systems can experience four sources of interference (base station-CPE, CPE-base station, base station-base station, and CPE-CPE), while FDD systems only experience two types (base station-CPE and CPE-base station) (See Figure 3). In any case, with the implementation of TDMA technology (something that is true of all current generation BWA systems), only one source of interference is experienced at any given time.

Unsynchronized TDD

During initial PMP deployments, interference of all types—for both TDD and FDD networks—is less of an issue, simply because networks will not be used to full capacity. Given even allocations of 100 to 150 MHz per carrier, available spectrum will be adequate to serve the customer base without the need for frequency re-use (FR) within each base station. As a result, the ease of compensating for interference of ACI will only be prevalent within each base station. CCI across base stations can be minimized through proper channel placement.

With proper RF planning, based on the requirements of TDD, these interference scenarios are improbable. So with FR up to1 (FR=1 being defined as using all available frequency within each base station), interference issues including those unique to TDD can be resolved adequately through careful RF planning. Carefully selecting channels to be used in each sector of a base station deployment, as well as physically locating base station antennas some distance from each other—as on separate corners of a rooftop—can eliminate interference concerns completely.

The tilt of the hub antenna to adjust the elevation of the flat beam targeted toward CPE can keep the signal from reaching unintended CPE, and the CPE antennas use narrow beams and are targeted precisely toward the specific hub antenna to which they are communicating. In particular, CPE-to-CPE communication is unlikely because CPE antennas are highly directional and targeted precisely at base stations, and because line-of-sight conditions between CPE seldom exist. If they do, a slight adjustment in the placement of one of the antennas eliminates this problem.

Polarization of the signal as a key technique for reducing interference can also be used to create further isolation between wanted and unwanted signals. If the beam is sent in the vertical plane, only CPE configured to receive a vertical wave will be able to receive it. CPE configured for horizontal waves will only pick up a small amount of the signal, even if the base station is aimed directly at the CPE. Proper alignment of base station and CPE antennas in conjunction with polarization can eliminate almost all CCI. Any remaining interference can be eliminated by attenuating the power output of the CPE so it is just great enough to reach the right base station.

Continued on page 60

TDD Advantages

Given that the additional sources of interference can be eliminated through these RF planning techniques, where there is no reuse of channels within each base station (FR), TDD deployments are on par with FDD solutions interference-wise. And yet, TDD presents the carrier with clear market advantages over FDD. First and foremost, the benefit is the ability for the system to adapt in real time to the up and down bursts of users. In contrast, FDD systems use fixed-channel allocations for up and down bandwidth; they are always inefficient in their utilization of equipment and capacity.

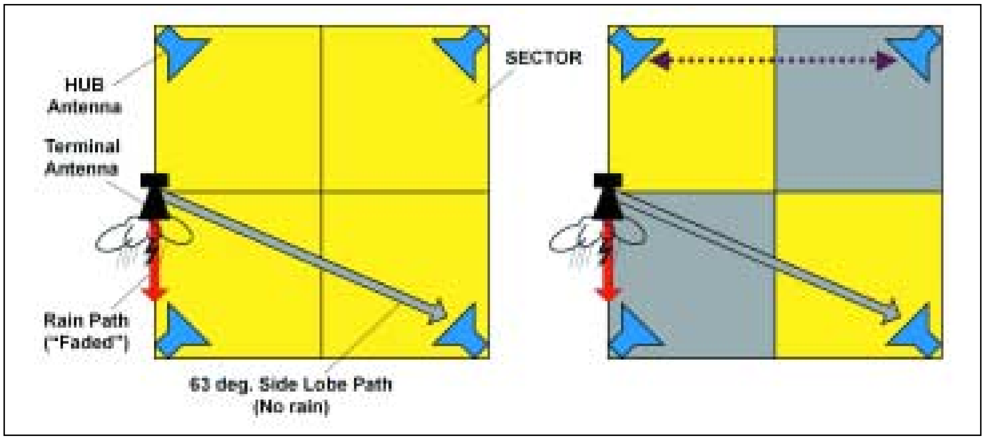

Figure 5. Uncorrelated rain problem: In scenario 1: In cases of rain fade, APC will boost terminal signal, which balances antenna discrimination on an upfaded link. In scenario II, nearby sectors use same frequency and polarization, but are time discriminated. In fade, the power boost would be ignored by neighbor hub, because it has not time associated with it.

Because the radio frequency and modem parts of the transmit and receive chains are operating on the same frequency at different times, a TDD system reuses certain elements for both chains (such as filters, mixers, frequency sources, and synthesizers) and eliminates isolation complexity completely.

In addition, with TDD, no guard bands are required to separate upstream and downstream frequency traffic. Usually as much as 200 to 300 MHz frequency separation is needed between transmit and receive frequencies for cost-effective modem designs in FDD. The result for the service provider is either the loss of that spectrum to serve customers or increased system costs because of expensive and inefficient duplexers. Subsequently, the use of TDD results in a substantial additional frequency savings that is most noticeable in block allocations such as LMDS A band with 850 MHz or LMCS bands of 500 MHz. For allocations such as the LMDS B Band where there is not enough guard band for FDD, TDD is the only solution.

The absence of guard bands also makes a TDD system flexible enough to be applied to different frequency allocations. For example, vendors have designed FDD systems for specific transmit, receive spacing of 1008 MHz (for CEPT) or 700 MHz (for 38 GHz). Applying these radio solutions to new allocations requires a redesign of the system. A TDD system does not have any transmit or receive considerations in its design. To implement an FDD approach for each frequency allocation would require a substantial amount of R&D changes for each allocation (and possibly multiple systems for a given license such as in the case of the LMDS A licensee's three allocations). As a result, this greatly increases the per-unit costs of each FDD system over the more flexible TDD approach that can serve all of the markets with a single design. Furthermore, because of TDD's inherent flexibility, it is easily applied to work in split bands that had been previously allocated with FDD in mind.

Synchronized TDD

As PMP networks are built out and carriers obtain a larger mass of customers, channels within each base station will need to be reused for maximum capacity to serve a growing customer base. Once the same frequency begins to be reused in a given base station (or FR greater than 1), additional complexities for RF planning must be considered. Although both TDD and FDD suffer greater interference issues in this more built-out network, the unique pattern for TDD of base station-base station interference becomes much more acute. The solution is to implement intrahub and interhub synchronization.

Synchronizing communications so that all base stations are transmitting at the same time and all CPE are listening at the same time, and vice versa eliminates the patterns of interference that are unique to TDD deployments (CPE to CPE and base station to base station). By synchronizing transmissions, nearby base stations can use the same frequency at the same time, and TDD deployments are as robust as FDD deployments (see Figure 3). TDD deployments are as robust as FDD deployments (See Figure 4).

And with TDD's software-based design, it is easy to implement synchronization at the bit level with the time division in milliseconds. By the simple installation of a GPS receiver at each base station—as is used in every CDMA base station deployment—the synchronization capability can be implemented by adding only about $200 to the cost of each base station. No change to the CPEs is required to implement synchronization.

As a base station serves a limited region in an area, targeting identical asymmetry profile for all transceivers has only a small impact on its ability to offer variable asymmetry in real time, and in any case still holds a major performance advantage over FDD systems that use a fixed channel profile. As mentioned previously, one key benefit of TDD is its ability to offer the additional advantage of variable asymmetry in the channel based on demand from users. In an unsynchronized scenario, the intervals can vary by destination and channel. Once synchronization is implemented for FR>1, the upstream and downstream asymmetry remains flexible throughout the service area.

Frame synchronization

To make BWA deployments more robust, synchronization can be used to avoid almost all interruptions due to interference. Frames can be synchronized on an odd-even basis channel by channel, burst by burst. For example, a given CPE may be instructed to only transmit on even frames in a given channel, while a nearby CPE could be instructed to transmit only on odd frames in that same channel. The result is that any potential for CCI is eliminated as the receiver in the base station only "listens" for even frames, so any signal from the base station transmitting only odd frames will not be recognized

CPE-to-base station interference is the most difficult problem because of the wide antenna pattern that allows base stations to pick up signals from many CPE, which also makes it susceptible to picking up signals from the wrong CPE. Because CPE-to-base station is the worst-case interference scenario, odd-even frame synchronization can be employed only where it is needed most, in the upstream path.

Two examples where interference can be so great that frame synchronization should be employed are combinatorial and uncorrelated rain interference. Combinatorial interference occurs when the same channel is being re-used in multiple sectors. As re-use increases, signals from each sector can combine to provide a signal strong enough to interrupt service. If the situation reaches the point where any single signal has a greater than 50 % chance of a line-of-sight connection with a base station, it can prevent the proper signal from being received by the base station.

Rain can create upstream interference as well. Because the LMDS wavelength is disrupted by rain, the CPE terminal increases its power to boost the signal to its base station, but that power boost can also allow the signal to reach other base stations, especially if there is no rain in the path between the CPE and the other base station (see Figures 5 and 6).

Due to the bit-level synchronization used to reduce TDD-specific interference scenarios, the implementation of odd-even frame synchronization for synchronized TDD systems comes without additional cost. And TDD's variable asymmetry capabilities can be used to keep this type of synchronization from affecting upstream capacity because more capacity can be created in the upstream signal to compensate.

FDD, on the other hand, does not maintain frame synchronization, and therefore cannot take advantage of synchronization to reduce or eliminate interference issues. In the field, interference problems between CPE and base stations have been demonstrated with FDD. Operators using FDD must face complex deployment scenarios that often must sacrifice capacity in order to prevent interference.

Final thoughts

Obviously, the larger the amount of frequency an operator has available, the more flexibility that operator has in deploying a network. With a smaller amount of frequency, it is more tempting to attempt to re-use frequency aggressively, but this may result in more problems than benefits. With aggressive frequency re-use, there is a greater chance for interference, resulting in more errors. Bandwidth must be used to correct those errors, so the end result is that the amount of available bandwidth remains the same. This is the frequency reuse dilemma I (See Figure 6).

Another approach for maximizing frequency re-use is to create more sectors per cell. But with a smaller channel available, there's more sensitivity to accuracy, requiring careful aim of each antenna to avoid interference issues. Dividing the available "pipe" into smaller sectors also effectively decreases the amount of bandwidth available per sector, making it difficult for operators to reap the benefits of having multiple customers share the same bandwidth. This is the frequency reuse dilemma II (See Figure 6).

In summary, real-world deployments will provide a true test of BWA PMP systems that will fully illustrate the challenges facing RF planners. While all approaches suffer similar interference problems, TDD has additional types of interference because of its method of using a single channel for operation. Synchronization of base stations is a technique that can be used to eliminate these additional types of interference and put TDD on par with FDD in terms of interference profile. So the carrier can experience the benefits of the TDD approach while still being able to plan the network as if it were any FDD network. For ultimate robustness in the network, frame synchronization can be implemented to deal with combinatorial and uncorrelated rain interference. The good news is that this capability comes for free in TDD because synchronization for up and down transmissions has already been implemented. In FDD, implementation of frame synchronization is much more difficult. Overall, synchronized TDD best allows service providers to maximize the use of the available frequency to serve more customers and generate more revenue, while still offering high-quality broadband services.

RF

About the author

Jay Klein is co-founder and chief technology officer for San Diego-based Ensemble Communications. Prior to founding Ensemble, he served as principal and vice president of research and development for CTP Systems, a wireless systems developer and manufacturer. Before CTP, Klein served eight years in the electronic R&D department of the government of Israel, leading a range of wireless communications initiatives. For his accomplishments and contributions, he received an award from the President of Israel. Klein holds a BSEE degree with honors from Tel Aviv University. He may be contacted on the Web at: **www.ensemblecom.com.**