CDG CDMA DEVELOPMENT GROUP

CDMA Technology

Welcome to the World of CDMA

The CDMA Revolution

The great attraction of CDMA technology from the beginning has been the promise of extraordinary capacity increase over narrowband multiple access wireless technologies. Simple models suggest that the capacity improvement may be more than 20 times that of the existing narrowband cellular standards, such as AMPS in North America, NMT in Scandinavia, TACS in the United Kingdom. Historically, the capacity was calculated using simple arguments. Reality, of course, is much more complicated than the idealized models. Real cell coverage areas are highly irregular, not the neat hexagons found in textbook models. Offered load is not spatially uniform, changes dramatically with time-of-day, and is often subject to other uncontrollable influences.

Background

An idealized multiple access mobile radio system consists of a family of base stations, or "cells," geographically distributed over the service area, and mobile stations. We use the term "mobile" generically to mean any subscriber station, whether it moves or not. The majority of new cellular sales are now in fact hand held portable units, and the market outlook is for that trend to continue for the foreseeable future. Non-traditional uses, such as wireless data modems in laptops, are also expected to grow dramatically in the near future.

Spectrum for mobile wireless is normally allocated in frequency division duplex (FDD) paired bands. Cellular systems are separated by 45 MHz, PCS bands by 80 MHz. Although there have been some proposals for the use of time division duplex (TDD), such operation inherently limits the coverage area, and have not achieved widespread acceptance.

Communication between base stations and mobile stations is established by a negotiation upon call origination. Once communication is established between base and mobile, movement of the mobile is detected and the service is handed over from one base station to another. One cell at a time services each mobile in the narrowband services. The concept of handoff is extended to a multi-way simultaneous "soft" handoff in the CDMA standards.

When cellular radio was first proposed by Bell Telephone Laboratories, cell subdivision was cited as the method by which the system would grow in traffic capacity. However the marketplace success of cellular has outstripped predictions by such a large margin that subdivision of cells to increase capacity, as a practical matter , is losing effectiveness. AMPS mobiles have a very limited ability to reduce transmitter power. High density, small cells also increase the handover rate, thereby increasing the probability that any particular call will be dropped. The overhead costs associated with large numbers of cells, such as real estate, backhauls, and maintenance, become an ever larger fraction of capital and operating costs. Zoning boards increasingly resist applications for new antenna towers that make their jurisdictions resemble oil drilling fields. An alternative method of capacity enhancement was needed.

Traditional Multiple Access Communication

Traditionally radio communication systems have separated users by either frequency channels, time slots, or both. These concepts date from the earliest days of radio. Even spark transmitters used resonant circuits to narrow the spectrum of their radiation. Scheduled net operation was probably the first manifestation of time slotting. Modern cellular systems began with the use of channelized analog FM. More recently several hybrid FDM-TDM digital systems have been developed, ostensibly to enhance service quality and capacity. In all these systems, each user is assigned a particular time-frequency slot.

In large systems the assignments to the time-frequency slots cannot be unique. Slots must be reused in multiple cells in order to cover large service areas. Satisfactory performance in these systems depends critically on control of the mutual interference arising from the reuse. The reuse concept is familiar even in television broadcasting, where channels are not reused in adjacent cities.

North American cellular allocates approximately AMPS 416 channels to each operator (30 kHz spacing, with a total allocation of 12.5 MHz in each direction - see the Frequency Plan). An idealized system geometry shown in the figure. The same frequency obviously cannot be reused in any adjacent pair of cells because a user on the boundary between those cells would receive both signals with equal amplitude, leading to an unacceptably high interference level. A plane can be tiled with hexagonal cells, labeled in accordance with the seven-way pattern shown in the figure. Thus, if a unique set of channels is assigned to each of the seven cells, then the pattern can be repeated without violating the adjacency requirement. Although this idealized pattern is not strictly applicable in all real systems, the seven-way reuse pattern is approximately correct. The capacity of systems built in this way is determined by the bandwidth per channel and the seven-way reuse pattern. In an AMPS system, therefore, the maximum capacity per cell is approximately 416/7 = 59. For three-way sectored cells, the same K=7 reuse applies over all three sectors, that is, only about 19 channels are available in each sector. In an ideal geometry the reuse pattern looks like this, representing channel sets by distinct colors ...

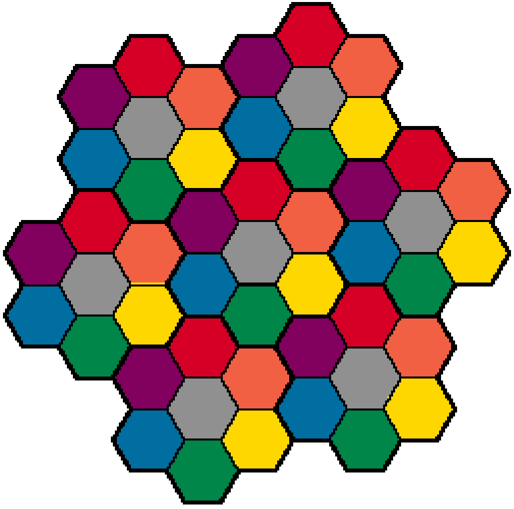

In this connection, it should be noted that achievement of the K=7 reuse, rather than an even larger number, depends on the fact that the effective propagation decay law is faster than free space. That is, in a vacuum electromagnetic radiation decays in intensity like R^{-2}. However measurements have consistently shown that the effective propagation law exponent is typically between -3.5 and -5 in the ground mobile environment. Interestingly, it is easy to show that if the propagation law were that of free space, a large cellular system would not be viable at all. The larger-than-free-space propagation exponent means that only the first tier of neighbor cells is significant in the idealized model.

The "Magic" of CDMA

CDMA offers an answer to the capacity problem. The key to its high capacity is the use of noise-like carrier waves, as was first suggested decades ago by Claude Shannon. Instead of partitioning either spectrum or time into disjoint "slots" each user is assigned a different instance of the noise carrier. While those waveforms are not rigorously orthogonal, they are nearly so. Practical application of this principle has always used digitally generated pseudo-noise, rather than true thermal noise. The basic benefits are preserved, and the transmitters and receivers are simplified because large portions can be implemented using high density digital devices.

The major benefit of noise-like carriers is that the system sensitivity to interference is fundamentally altered. Traditional time or frequency slotted systems must be designed with a reuse ratio that satisfies the worst-case interference scenario, but only a small fraction of the users actually experience that worst-case. Use of noise-like carriers, with all users occupying the same spectrum, makes the effective noise the sum of all other-user signals. The receiver correlates its input with the desired noise carrier, enhancing the signal to noise ratio at the detector. The enhancement overcomes the summed noise enough to provide an adequate SNR at the detector. Because the interference is summed, the system is no longer sensitive to worst-case interference, but rather to average interference. Frequency reuse is universal, that is, multiple users utilize each CDMA carrier frequency... The reuse pattern is now

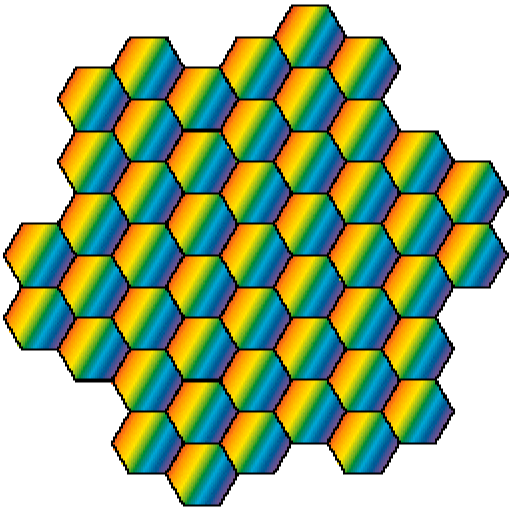

The rainbow cells indicate that the entire 1.25 MHz passband is used by each user, and that same passband is reused in each cell.

Capacity is determined by the balance between the required SNR for each user, and the spread spectrum *processing gain.* The figure of merit of a well-designed digital receiver is the dimensionless signal-to-noise ratio (SNR)

$$E_b/N_0 \equiv \frac{\text{Energy per bit}}{\text{Power spectral density of noise + interference}} \quad (1)$$

The "noise" part of the SNR, in a spread spectrum system is actually the sum of thermal noise and the other-user interference. The SNR needed to achieve a particular error rate depends on several factors, such as the forward error correction coding used, and the multipath and fading environment. For the receivers typically used in commercial CDMA it ranges typically from about 3 dB to 9 dB.

Energy per bit is related to signal power and data rate:

$$E_b = P_s/R \quad (2)$$

The noise + interference term is power spectral density. If the spectrum of the signals is roughly rectangular, with a bandwidth of W, then the noise + interference power spectral density is

$$N_0 = F_N k_B T_0 + W^{-1} \sum_{\text{other users}} P_i \quad (3)$$

where the first term represents the thermal noise level of the receiver (F_N = receiver noise figure). Rewriting the SNR equation in terms of the data rate and the spread-spectrum bandwidth shows where the magic lies:

$$\left(E_b/(N_0 + I_0) \right)_j = \frac{P_j / R}{N_0 + W^{-1} \sum_i P_i} \quad (4)$$

The interference in this equation is the sum of the signals from all users other than the one of interest.

This equation is the key to understanding why CDMA was not explored for use in terrestrial multiple access systems. It is also the key to the innovation that led to commercial CDMA.

Near-Far Problem

CDMA (and spread spectrum in general) was always dismissed as unworkable in the mobile radio environment because of what was called the "near-far problem." It was always assumed that all the stations *transmitted* constant power. In the mobile radio environment some users may be located near the base station, others may be located far away. The propagation path loss difference between those extreme users can be many tens of dB. Suppose, for example that only two users are present, and that both are transmitting with enough power that the thermal noise is negligible. Then the SNR, in dB, is

$$\left(E_b/N_0 \right)_{dB} = \frac{W}{R} + P_j - P_i \quad (5)$$

If there is, say, 30 dB difference between the largest and smallest path losses, then there is a 60 dB difference between the SNR of the closest user and the farthest user, because these are the *received* powers. To accommodate the farthest users, the spreading bandwidth would have to be perhaps 40 dB, or 10,000 times the data rate. If the data rate were 10,000 b/s, then W=100MHz. The spectral efficiency is abysmal, far worse than even the most inefficient FDMA or TDMA system. Conversely, if a more reasonable bandwidth is chosen, then remote users receive no service.

This observation was, for years, the rationale for not even attempting any sort of spread spectrum in any but geosynchronous satellite environments, where the path loss spread was relatively small.

Power Control

The key to the high capacity of commercial CDMA is extremely simple: If, rather than using constant power, the transmitters can be *controlled* in such a way that the *received* powers from all users are roughly equal, then the benefits of spreading *are* realized. If the received power is controlled, then the subscribers can occupy the same spectrum, and the hoped-for benefits of interference averaging accrue.

Assuming perfect power control, the noise plus interference is now

$$N_0 + I_0 = N_0 + (N-1)P_s$$

$$N_0 = F_N k_B T_0 \qquad (6)$$

where N is the total number of users. The SNR becomes

$$E_b / (N_0 + I_0) = \frac{P_s / R}{N_0 + (N-1)P_s / W} = \frac{W/R}{(N_0 W / P_s) + N - 1} \qquad (7)$$

Maximum capacity is achieved if we adjust the power control so that the SNR is exactly what it needs to be for an acceptable error rate. If we set the left hand side of (7) to that target SNR and solve for N, we find the basic capacity equation for CDMA:

$$N - 1 = \frac{W/R}{\left(E_b / (N_0 + I_0) \right)_{\text{target}}} - \frac{N_0}{P_s} \xrightarrow[P_s \to \infty]{} \frac{W/R}{\left(E_b / (N_0 + I_0) \right)_{\text{target}}} \qquad (8)$$

Using the numbers for IS-95A CDMA with the 9.6 kbps rate set, we find

$$N \approx \left(W/R \right)_{dB} - \left(E_b / (N_0 + I_0) \right)_{\text{target, dB}} \approx 21.1 - 6\text{dB} = 15.1\text{dB} \qquad (9)$$

or about N=32. The target SNR of 6 dB is a nominal estimate. Once power control is available, the system designer and operator have the freedom to trade quality of service for capacity by adjusting the SNR target. Note that capacity and SNR are reciprocal: a three dB improvement in SNR incurs a factor of two loss in capacity, and vice-versa.

We've neglected the difference between N and N-1 in (9). This is convenient in the capacity math, and is usually reasonable because the capacity is so large.

There are factors we haven't taken into account yet. Some of the things we have not yet considered actually help; others hurt. But on balance, there is a major improvement over the narrowband technologies.

The sustainable capacity is proportional to the processing gain, reduced by the required SNR. While there are several considerations we have yet to look at, there is already a suggestion of the capacity enhancement possible. With E_b/N_0 in the 3-9 dB range, equation (9) gives a capacity in the neighborhood of 16-64 users. In the same bandwidth, a single sector of a single AMPS cell has only 2 channels available.

Embedded Cell Capacity

The discussion leading to equation (9) assumes only a single cell, with no interference from neighboring cells. One might ask what has been gained here. The capacity of an isolated AMPS cell likewise is very high. In fact, there is nothing to stop you from using *all* the channels if there are no neighbors; reuse is not needed. The capacity of that fully populated AMPS cell would be about 42 channels (1.25 MHz/ 30 kHz channel spacing). This is not greatly different than the number that we just calculated for CDMA.

So what is the point of using CDMA?

This is where the big difference is. To find what happens *with* the neighbor cell interference, we have to add that interference into equation (3) above. The math of this can be found in several of the references [e.g. Gilhousen, et al.]. It turns out that the fraction of the reverse link interference that comes from the neighbor cell is about 60% of the own-cell interference. And, importantly, this answer is not terribly sensitive to the parameters of the model, provided we assume that the mobiles are power-controlled in a sensible way.

For doing the calculation, we just introduce the effective frequency reuse factor F, defined as

$$F \equiv \frac{\text{Total interference power}}{\text{Own - cell interference power}} \tag{10}$$

The effect of the additional interference can be carried through the capacity math by just replacing N everywhere it appears by F *N. The pole capacity, assuming now that all cells are equally loaded, is now

$$N_{\text{pole}} = \frac{(W/R)}{\left(E_b \big/ (N_0 + I_0)\right)_{\text{target}} \cdot F} \tag{11}$$

From equation (11), it is evident why we call F the effective frequency reuse factor. It plays the same role in the CDMA capacity equation (11) that the narrowband frequency reuse factor K does in AMPS capacity. An omnidirectional CDMA cell, even with other-cell interference, has a higher capacity than an AMPS cell by a factor of K/F = 7/1.6 or about 4.4. The improvement in a sectored cell is even more dramatic because the CDMA capacity is largely *unaffected* by the sectorization. That is, (11) applies to one *sector,* with only small modifications because of the interference leakage between sectors, while AMPS *does not gain from the sectorization.* In this case the gain is K/F = 21/1.6 or about 13 times.

Voice Coding

Because the interference is averaged, anything that can be done to reduce the average transmitted power enhances capacity. An obvious target for such power optimization is the speech coding. Human speech is an intermittent information source. Measurements at Bell Laboratories many years ago suggested that the activity factor in natural human conversation is in the range of 35-40%. If that activity factor can be translated into power gating, then a further increase in capacity of perhaps two times or more is possible. This is in fact accomplished by features in both the air interface standards and the voice coder service option standards. The pole capacity becomes

$$N_{\text{pole}} = \frac{(W/R)}{\left(E_b \big/ (N_0 + I_0)\right)_{\text{target}} \cdot F \cdot v} \tag{12}$$

where v is the voice activity factor, approximately 0.5. The reduction achieved in practice is less than the measured activity factor because the transmission rate in the air interface is not reduced to zero during idle periods. The gain over AMPS capacity, by the same assumptions as above, is now about 26 times - perhaps optimistic given the crude nature of the model, but suggestive of the substantial improvements possible by converting to CDMA.

Variable rate data is accommodated in the air interface by providing a basic traffic data rate that can be reduced by binary ratios (1, 1/2, 1/4, and 1/8). Transmission is never reduced to zero because this would present problems related to channel supervision. Two rate families are currently supported: the first based on 9600 bps, the second on 14,400 bps. These, of course, have different capacity characteristics.

The voice coding used in CDMA is standardized separately from the air interface as a *service option*. The first coder standardized was IS-96; others are in widespread used, and will also soon be standardized. The IS-96 standard operates in the nominal 9600 bps rate set, and achieves an average effective rate of something just over 4kbps, or about 50% of the actual air interface payload data rate of 8550 bps.

Multipath Propagation

System capacity, as you might expect, is affected by propagation phenomena. Users of analog cellular phones are familiar with the fading that is so annoying, especially in handheld portables when standing nearly still. Fading in a moving vehicle is more rapid, being caused by motion of the vehicle through stationary interference patterns, where the spatial scale of the interference pattern is the wavelength, about one foot. CDMA is much more robust than the analog technologies in the presence of multipath, but it does affect capacity.

There are two questions that one must address regarding multipath fading and CDMA. First, under what circumstances will CDMA experience fading, and second, what is the effect of fading, when it occurs, on the CDMA channel?

When does multipath cause fading, and when does it not?

When the multipath components are "resolved" by the CDMA waveform, that is, when their delays are separated by at least the decorrelation time of the spreading, then they can be separated by the despreading correlator in the receiver. They do not interfere because each component correlates at a different delay. When the multipath components are separated by less than the decorrelation time, then they cannot be separated in the receiver, and they *do* interfere with one another, leading to what is sometimes called *flat fading*.

Fading is also characterized as Rayleigh or Rician. Rayleigh fading is the result of a vector sum of multiple signal components, each having a random amplitude. It can be viewed alternatively as a signal whose I and Q amplitudes are Gaussian random deviates. Rayleigh fading exhibits deep signal dropouts, as shown in the figure.

If there is a strong, constant component to the signal, in addition to the multiple random components of Rayleigh fading, then the fading is said to be Ricean. Ricean fading is typical of line-of-sight situations, where there is a direct, unobstructed path between stations, as well as reflecting or scattering surfaces.

The duration of one spreading *chip* is 1/1.2288MHz = 814 ns, or at the speed of light, 244 meters. Multipath differences less than this will lead to flat fading; greater will lead to resolved multipath, which will be diversity combined by the receiver.

To address the second question, that of the effects of fading, the answer is complex and is different in the forward and reverse links. It also depends on the fading rate, which in turn depends on the velocity of the mobile station. Generally fading increases the average SNR needed for a particular error rate. The increase can be as much as perhaps 6 dB. In the reverse link, the power control will mitigate the effects of fading at low speed; at high speed it has little effect. At high speed, and in both links, the FEC coding and interleaving becomes more effective as the characteristic fade time becomes less than the interleaver span.

Coverage versus Capacity

There is some bad news arising from the CDMA capacity equation. Namely, the fact that the power that the mobiles are required to transmit goes to infinity as the capacity pole is approached. As the required power increases, mobiles at the fringe of coverage will begin to run out of transmitter power. That is, they will be asked to transmit more than their capability allows. The practical consequence of this is that the system load should really be controlled so that the planned service area never experiences coverage failures because of this phenomenon.

There are some interesting mathematical models of this, that we talk about in our Coverage-Capacity pages. It is not really so much a problem as it is a system design consideration. You cannot simultaneously achieve maximum capacity and maximum coverage. It is a tradeoff.