Wireless Multiple Access Adaptive Communications Techniques

Gregory J. Pottie
Electrical Engineering Department
University of California, Los Angeles

Abstract: The extreme variability of wireless channels demands that adaptivity be the one constant feature of high-performance systems, whether for dedicated channels, cellular systems, or multi-hop networks. We outline the fundamentals of the physical channel, and means of achieving reliable communications over highly variable links. Methods for managing multiple access interference in cellular systems are described, as well as the additional challenges posed by multimedia traffic and non-standard network topologies. At each step, further adaptive capability is demanded of the system, and care must be taken to ensure that the adaptive algorithms interact in a complementary and stable fashion. Using such techniques, very large performance improvements over static designs may be obtained.

I. Introduction

The variability of wireless channels presents both challenges and opportunities in designing multiple access communications systems. To maximize throughput for a given power budget, the link must adapt to the actual channel conditions, changing the transmitter power level, antenna beam pattern, equalizer settings, and possibly the symbol rate and constellation size. On the other hand, the attenuation and directionality of signals makes possible re-use of the time/frequency resources in space, permitting a large number of users to access the shared medium. These users are coupled by the mutual interference they cause one another. For an acceptable quality of service, each user will typically need a signal to interference ratio (SIR) above some target. The techniques which enhance reliability of a link also affect the interference seen by other users, as do multiple access techniques. Thus, it is counterproductive to separate consideration of the physical and medium access layers in designing a wireless transmission system.

A fundamental limit on the capacity of a multiple access wireless system is imposed by the inability of the various users to perfectly estimate and predict the time-varying channel and interference. In the limit of large complexity and power levels, all other impairments can in principle be overcome if the channel dynamics (including interference) are slow enough. In this article we will describe practical techniques for ameliorating the impairments in point to point communications (and their limitations), and then discuss the extension of these techniques to the multiple access setting. Throughout we pay particular attention to how the channel and interference couple the adaptation of the different techniques, in both desired and undesired fashions.

The remainder of the article is organized as follows. In section II we discuss radio links, providing an introduction to radio propagation, and then focussing on means for dealing with multipath interference. In section III we discuss multiple access techniques for cellular systems, assuming connection-oriented traffic, including TDMA, FDMA, and CDMA, along with means to enhance the capacity of traditional approaches. These include dynamic resource management, smart antennas, and multi-user detection, as well as the interaction among the various methods. In section IV we discuss alternative access scenarios, including network topologies such as peer-to-peer and multi-hop networks, and the impact of multimedia traffic. In section V we provide our conclusions and suggestions for further research.

II. Radio Links

Wireless channels are distinguished from wireline channels by their large variability and inherent multiple access nature. Here we focus on the causes of the variability and the means for dealing with it, deferring discussion of multiple access until section III.

II.1 Fundamentals of Radio Propagation Modeling

Radio propagation can be understood at many levels: Maxwell’s equations, Huygen’s principle and ray-tracing, and statistical models at various levels of abstraction. Which method to use depends on the task at hand. Solution of Maxwell’s equations is mandated in design of high performance antennas, and to fully characterize the interaction of antenna array elements with each other and their immediate environment [1,2]. However, the computations required to characterize propagation within even a small structure are immense, and to be accurate must include precise measurements of the electromagnetic properties of the composite materials in the structure. The effort involved is comparable to directly taking measurements of the point to point link. A more fruitful approach in designing wireless links is to create a simplified model whose statistics match what is found in some combination of measurements and sophisticated mathematical modeling. For these purposes, a model which starts from the theory typically found in first year physics textbooks is reasonable [e.g. 3].

Huygen’s principle for wave propagation states that we may regard each point on a wavefront as the source for a new circular wavefront. The waves add by superposition. This principle accounts for the circular expansion (and thus square law free space attenuation) of a wavefront from a point source, the approximation of the wavefront as a plane wave at large distance from the source, and the phenomenon of diffraction. The theory of optics abstracts plane waves using the notion of ray tracing. In this point of view (which neglects diffraction), rays are emitted omnidirectionally by a point source. Reflection and refraction angles are computed through consideration of the differences in the index of refraction. Adding the attenuation of waves in various media to the theory of ray tracing provides a reasonable first cut at understanding how electromagnetic radiation intensity levels will vary. Multiple rays arriving at the same destination are added by superposition, with the phases of the rays determined by the lengths of the propagation paths and the phase changes induced at reflective boundaries. This is the level of abstraction at which propagation models suitable for radio engineering are usually presented, since it captures most of the observed dynamics in a model that is easy to visualize. It is also the basis for ray tracing and ray launching simulation models. Moreover, these simple considerations provide guidance in designing large scale propagation studies to gather data which may later be abstracted into statistical models.

The propagation losses for a radio link are usually divided into three categories: multipath, shadowing, and distance loss [4,5]. These are convenient in constructing models, and in developing techniques to combat propagation losses, which tend to be effective for only one or two of the categories. Multipath refers to the many different propagation paths between the transmitter and receiver, each characterized by its own phase, delay, and attenuation. It results in channel variations over distances on the order of a wavelength. Shadowing refers to local variations in the received signal strength caused by structures, hills, canyons, vehicles, etc. When we average over the multipath, it is recognized as a deviation from the empirically observed distance propagation law, which is the best fit for the received power to d^{-n}, where n is typically between 1.5 and 6 and d is the distance from the base station. The exponent is the result of some combination of absorption, scattering, ground reflections, and free space loss. Taking all of these factors together, the impulse response of the channel may be modeled by [6]

$$c(t,\tau) = \sum \alpha_i(t)\, e^{\theta_i(t)}\, \delta(t-\tau_i) \qquad (1)$$

where the i^{th} ray has amplitude α_i, phase θ_i, and propagation delay τ_i, at time t.

The expected amplitudes of the multipath components decline as the delay τ increases, since these rays experience larger free space and absorption losses. The length of the impulse response or “delay spread” is dependent on physical factors such as the orientation, reflectivity, and distance between strongly reflective objects (e.g., buildings, mountains, or walls of a room). Delay spreads can range from tens of nanoseconds in indoor settings to tens of microseconds in high-power outdoor applications. Figure 1 illustrates one multipath propagation scenario in which there are both strongly reflective objects (A,B,C,D) and foliage which scatters some of the energy. Thus, the direct path has the shortest length and no intervening foliage, and will be the strongest component. The multipath components arrive later due to their increased path lengths, and are attenuated with respect to the direct path because of their longer lengths and in some cases the passage through the foliage.

Depending on the bandwidth of the signal relative to the delay spread of the multipath, as the communicators move with respect to one another either fading (variations in the signal envelope) or time-varying dispersion is observed. In the absence of a strong line of sight (the worst case), the sum of the multipath components at any given time may be modeled as a zero-mean complex Gaussian random variable, and consequently the envelope has Rayleigh statistics. Small movements result in large changes in the phases of the individual rays, leading to large variations in the sum. Thus, the components of the impulse response can change dramatically for movements of less than a wavelength, even though the attenuation for the individual rays changes very little. Narrowband systems are those for which the symbol period is much longer than the multipath dispersion. For these systems, the envelope of the received signal will vary with the fading statistics. Wideband systems on the other hand will experience both fading and intersymbol interference.

As noted earlier, exact calculation of the channel impulse response at all possible transmitter and receiver locations is very complicated. In practice, a variety of statistical models are used to predict the reliability of radio communication systems. In general, the sophistication of the model must increase with the complexity of the radio system, so that all the possible limiting impairments are accounted for. For example, for point to point line of sight digital microwave transmission the three-ray Rummler model has proven to be a reasonable predictor of channel outage [7]. It consists of a direct ray, a ground bounce, and a refracted ray, with the statistics for the variations in the phases and attenuations of these rays derived from empirical observations. However, when the system also includes multiple antennas, we need to include spatial correlation among the different ray components to provide an accurate model.

A frequently used model for urban radio propagation that accounts for distance and shadowing losses is as follows [8]. Mobiles are assumed to be uniformly distributed on an infinite plane, with base stations arranged in a uniform hexagonal pattern. Power drops with the fourth power of distance, reflecting free space losses and the influence of the ground-reflected ray. In addition, there is independent log-normal shadowing for each pair of transmitters and receivers. Roughly speaking, the physical assumption is that the shadowing is due to reflections off a large number of objects, each of which produces a multiplicative loss (or gain through confinement of the wavefront); by the central limit theorem, the result is a Gaussian intensity distribution, if intensity is measured in dB. The received power level for a mobile m at distance d from a base station is then

$$P = \frac{10^{0.1\xi} P_m}{d^4} \qquad (2)$$

where P_m is the mobile's transmit power, and ξ is a Gaussian random variable with zero mean and standard deviation $\sigma=8$. The mobiles are assumed assigned to the base station with the minimum propagation loss.

The above basic model can be modified in many ways to take account of specific environmental conditions and the multipath. The shadowing standard deviation is more severe in rough terrain, and in general when the height of the base stations is low compared to the local relief (whether due to buildings or terrain). The distance power loss also depends to some extent on terrain, since dense vegetation or heavy urban build-up can lead to additional attenuation, while propagation in tunnels is enhanced by wave-guiding effects. Moreover, the ground reflected ray leads to less loss in smaller cells, so that a third power law is often assumed.

Multipath may be accounted for in several ways. For a static simulation where both transmitter and receiver use a single antenna element, we may multiply the received power by an independent variable which follows any of the popular statistical models: Rayleigh, Rice, or Nakagami. The Rayleigh distribution is in some senses a worst case assumption, as it effectively assumes no line of sight and a large number of reflective objects. The Rice distribution more closely matches observations for satellite systems and terrestrial systems in which there is relatively little ground clutter (e.g. suburban). However, in general simply generating one random variable for a link is inadequate, and we must include spatial correlation in both the shadowing and multipath to deal with motion of the communicators and to model the behavior of antenna arrays.

Spatial correlation of shadowing is by far the simpler quantity to model. Shadowing will vary roughly with the dimensions of the objects most responsible for it (e.g., trucks, buildings, mountains). To simulate this, a grid of shadow values can be created, where each value in dB is randomly generated according to the appropriate distribution. Then intermediate points can be interpolated linearly in dB. A lower storage option which captures the dynamics at the expense of giving inconsistent shadowing values to users who later occupy the same location is to create grid points only around each active communicator.

The interpolation method above is but one instance of spatial filtering, which is the usual perspective in modelling the spatial correlation of multipath [9]. At regular spatial intervals, random multipath variables are input to the spatial filter, so that the multipath value at any given location is a function of all the random variables within the filter span. The filter may be linear or multidimensional depending on whether antenna arrays are used. The filter parameters can be set to conform to the empirical correlations for a particular environment, or statistics collected using ray launching, ray tracing, or solving Maxwell's equations. Since multipath varies widely over a single wavelength, this entails many more computations than the shadowing loss. To reduce this at the expense of some storage, a feasible alternative is to store a set of multipath profiles generated according to this method, and randomly choose among them. Yet another approach is to work out error probabilities over a wide set of multipath conditions for different signal to interference ratios and relative communicator speeds, and then abstract this into the model so that in the large scale simulation multipath is not directly considered. Indeed, abstraction of the results of physical layer modeling is essential in performing simulations for large networks.

The motion of the communicators and the type of traffic they generate also have a large impact on the performance of radio communication systems. Traffic engineers have long studied the motions of vehicles to improve the flow on city streets and freeways. Taking these factors into account reveals that traffic neither moves at uniform speeds nor in uniform densities [see e.g. 10]. These issues are important in designing methods to share resources among cells and to manage handoffs of calls between cells; they also provide predictions on the range of vehicle speeds that may be anticipated, and thus the fraction of time particular multipath and resource management techniques will work. A newer issue is the nature of the data traffic. While cellular communications have been traditionally designed with only voice calls in mind, an increasing fraction of traffic in personal communications systems will be multimedia: a combination of data, voice, and video. The statistical variations of this kind of traffic are far different from voice, with implications for channel dynamics (via mutual interference) and thus the appropriate access methods and signaling schemes. We will discuss this issue in more detail in section IV, in the interim making the assumption that we are dealing mainly with traffic with a long hold time.

To decide on how detailed a model to use, the most important question is whether the model captures the essential channel dynamics. If it does not, the results will typically be wildly optimistic. Attempting to predict error rates for a given physical situation is very difficult in general, but prediction of reliability over an ensemble of conditions is fortunately easier, and what is more often needed. Simpler still is comparison of alternative algorithms, in which one need test only some representative channels which include the most stressful conditions likely to be found in practice. This is the basic idea behind the production of test suites of channels that are widely used in standards processes. In this selection of models, consideration of the underlying physical conditions is the essential guide.

II.2 Mitigation of Channel Variability

We now turn to the problem of dealing with the variable channel in point to point connections. We begin with the channel estimation problem, and then describe mitigation techniques which require adaptation only in the receiver, finishing with techniques requiring cooperation between the transmitter and the receiver. The latter are more powerful, but require a more slowly changing channel.

II.2.1 Estimation of the Channel State

The most effective methods for mitigating channel variability require that the receiver has some knowledge of the channel impulse response and the mean and variance of the combined noise and interference. We refer to this kind of information as the channel state. There are well-defined limits on how well the channel state can be learned, even given infinite receiver complexity, and much lower limits if low-complexity iterative techniques are to be used. It is the channel estimation problem that fundamentally limits the performance of digital radios in conditions of fast channel change or low SNR.

The Cramer-Rao bound relates the mean square error of an optimal sequence estimator to the SNR and number of observations [11]. For example, when estimating a deterministic parameter in Gaussian noise, the variance of the estimate must be greater than or equal to the noise variance divided by the number of independent observations. Clearly for a given level of fidelity in the estimate, as the SNR worsens we need a longer observation period (in effect, for more noise averaging). The difficulty comes when the channel state changes quickly with respect to the rate at which independent observations can be made. Using a longer observation window will eventually degrade the quality of the estimates, since it will include observations that do not reflect the current channel state.

Consider for example a situation where one communicator moves at velocity v, while transmitting at a wavelength of λ. The phases of the multipath components can change completely in the time it takes to move a wavelength, and thus the impulse response of the channel undergoes large changes in this time interval of v/λ. (Actually appreciable change can take place over shorter intervals, but this will serve for rough calculations). The appropriate estimation strategy depends on the length of this interval compared to the observation period. The minimum observation period is the delay spread of the channel, since we must form accurate estimates of all components within this span to take full advantage of the SNR available. We may make estimates of the frequency response, as might be done for multi-carrier modulation, or in the time domain, as would be required for single-carrier modulations, but in any case the same number of parameters per unit time must be estimated. For each parameter a certain minimum time is required to get an accurate estimate. If the total time required to estimate the parameters to the desired fidelity is a small fraction of the time between large channel changes (e.g. less than 1%), then decision directed techniques such as the recursive least squares (RLS) algorithm can easily track the channel state. If for example we consider the equalization problem for single carrier modulations, at high SNR and for the channel changing slowly during the observation span, the RLS algorithm will converge in roughly 2M iterations, where M is the span of the delay spread in channel symbols [12].

However, for faster changing channels or channels with lower SNR, the observation window needed to form accurate channel estimates becomes a larger fraction of v/λ. In this case, one must devote some fraction of the energy to explicit training, or coherent communications becomes difficult. Many cellular standards send pre-ambles, midambles or even pilot frequencies containing such training information. Another feasible strategy is to send periodic training sequences, estimate the channel during these intervals, and then interpolate to estimate the channel state in between [13]. The training symbol intervals must be long enough to form accurate estimates according to the Cramer-Rao bound, and frequent enough so that by the sampling theorem the intervening channel values can be reproduced. Here however we must take explicit account of channel changes during the training interval or error floors result [14]. Substantially the same diversity techniques and constellation sizes may be used as if decision directed techniques were employed; the penalty lies in the fraction of symbols that must be devoted to training, as well as the elevated signal processing cost for forming reliable estimates. When the conclusion is that substantially all symbols must be used for training, there is no option but to use non-coherent modulations and combining techniques, in effect giving up on estimation of the multipath. This has a very large performance penalty, as we will now detail.

II.2.2 Diversity Techniques

Fading leads to serious degradation in the quality of digital communications. For a Gaussian channel without fading, the probability of error declines exponentially with signal to noise ratio γ. However, for a Rayleigh fading channel perturbed by additive Gaussian noise the error probability varies only as $1/\gamma^*$, where γ^* denotes the average SNR [6]. The difference in SNR to achieve a particular error probability on the Rayleigh versus the Gaussian channel can be extremely large--it is already 40 dB at an error probability of 10^{-5}. Simply increasing the transmitted power to provide an adequate fading margin is not practical given the high additional cost in power and circuit complexity, and would be totally pointless in interference-limited multiple access systems.

This power penalty can in principle be completely overcome using diversity techniques, provided that the receiver has time to make proper estimates of the channel state. Throughout this article we consider the channel to have a time-varying impulse response of the form of equation (1), with the additional impairment of a combination of additive noise and time-varying interference. At different times, locations, and frequencies the instantaneous SNR's will be independent for large enough spacing. Diversity of order L is obtained by receiving L independent copies of the signal, and then using optimal combining techniques [6]. In essence, the copies are weighted according to their relative reliability, and then coherently combined so that the minimum mean squared error (MMSE) estimate of the received signal can be formed. Any number of adaptive techniques may be employed to achieve the MMSE solution [12]. With diversity of order L, the error probability varies as $1/\gamma^{*L}$; as L goes to infinity, performance approaches that of the Gaussian channel. For example, suppose that a very large L is achieved by spreading the signal in time using an orthogonal transformation and interleaving. Then the information is spread over many symbols which experience independent fading and additive noise. Following weighting using estimates of the fading depth, application of the inverse transform results in averaging over all the SNR's. The aggregate channel which includes the transformation, fading, weighting, and inverse transform will appear to be a channel with a nearly constant SNR of γ^*.

Diversity may be realized in the time domain through the use of error correcting codes. If it is possible to make good estimates of γ and use them in the decoder, then for binary signaling the diversity order is approximately equal to the minimum distance of the code, M, and in addition we realize the usual coding gain against additive noise. However, diversity is obtained only if fading is nearly uncorrelated between successive code symbols. Interleaving may be used to provide this independence at the cost of delay. Channel codes devised for the Gaussian channel are in general not the best choice for providing diversity protection, and additionally the best metric to use in decoding is not necessarily the squared Euclidean distance. For example, we would want to weight the codeword components according to the variable SNR in each segment. This becomes possible for moderate fade rates when interleaving is used. The type of weighting used depends upon the quality of our channel estimates [15]. Additionally, for multi-level constellations the Euclidean distance is not the most important design criterion for the code in the presence of fading. Thus, trellis coded modulation schemes which leave some bits uncoded are not suitable as techniques to provide diversity. Use of increased constellation expansion (redundancy) is one potentially useful approach [16]; additionally the codes can be designed using different criteria of optimality [17].

Space diversity is obtained by using multiple antennas, either spaced several wavelengths apart or with different polarizations or beam patterns. With L elements receiving independent signals and using optimum combining, the diversity order is L. We are generally limited in the number of antenna elements that may be used by the size of the platform or by the cost of using multiple receivers. However, this form of diversity is highly attractive because it is achieved at no cost in latency or bandwidth, and adaptive antenna arrays provide gain even in the absence of multipath.

There are three popular methods for combining signals from multiple antenna branches: selection, equal gain, and maximal ratio. In selection diversity, only one downconverter is required for multiple antennas. The various antennas are periodically and sequentially monitored to see which one produces the largest SIR at the decision point of the receiver. Then this antenna is selected for reception and/or transmission. The performance is reasonable for a two-antenna system, but is highly suboptimal for multiple antennas, and rapid channel variations result in severe losses. Equal gain combining is used for noncoherent systems when it is difficult to obtain accurate channel estimates (e.g., a fast frequency-hopped system). Its performance is not much better than selection diversity, but it may be used even for quickly changing channels. A separate receiver is required for each antenna, since the combining takes place after demodulation. In maximal ratio combining a separate receiver for each branch is also required. However, now the outputs of each branch are weighted by the estimate of the signal to noise amplitude for that branch. Decisions are then made using a weighted sum of the demodulator outputs. This procedure amounts to maximum likelihood detection.

Frequency diversity may be realized directly in the frequency domain, or in the time domain. An equalizer for a wideband system will not only deal with the intersymbol interference but it will also provide a diversity benefit, since the time dispersion is just the dual of frequency selective fading. The diversity order will be essentially the same as that obtained by attempting to directly obtain diversity by frequency division multiplexing of identical signals. In either case, if W is the bandwidth and C is the coherence bandwidth of the fading process (roughly the inverse of the delay spread), then frequency diversity of order L = W/C is possible.

In a direct sequence spread spectrum signal, the information sequence is multiplied by a pseudo-noise (PN) sequence whose components have a “chip” duration much less than the symbol duration. The modulated signal then has a much larger bandwidth than direct modulation of the original information-bearing sequence. In the absence of multipath, a receiver which correlates the received sequence to the PN code suffices for maximum likelihood detection. For dispersive channels, this is replaced by a RAKE receiver [6], which consists of a tapped delay line with PN correlators and an adaptive complex tap for each delay. The receiver may then be used to directly estimate the multipath components at delays that are multiples of the chip period. To reduce complexity, it is possible to search for the largest multipath components over the delay span of the receiver so that only a few correlators and adaptive taps are active at any time.

In a frequency hopped spread spectrum system, the instantaneous frequency is governed by a pseudo random sequence. In contrast to DS systems, the bandwidth is unrelated to the hopping rate. The frequency can be changed one or more times per data symbol (fast hopping) or only after many data symbols (slow hopping). These systems obtain a frequency diversity benefit only if coding is used. The simplest but least efficient technique is to use the repetition code, as is often done in fast-hopped systems. A better approach is to use some orthogonal transformation (e.g., the Hadamard transform) on the data in the transmitter to impress memory, and use the inverse transform in the receiver [18]. The delay can however be quite large, and no coding gain is obtained with respect to additive noise. Channel codes spread information more compactly. The combination of frequency hopping, coding, and interleaving over hops provides a very efficient way to achieve frequency diversity in slow-fading channels without excessive delay.

Diversity obtained in independent domains is multiplicative, and it is generally easier to achieve moderate diversity order in several domains than it is to consistently achieve high order in one. However, while frequency, time and space are often modeled as independent domains, this is not quite the case in practice. Techniques which include multiple antennas and an equalizer or RAKE receiver fall under the rubric of space-time methods. Generally the antenna array is adapted to minimize mean squared error, which results in a combination of interference nulling, antenna gain, and diversity combining. This usually makes more benign the impulse response seen by any equalizer that follows (as certain multipath directions are discriminated against), and thus also reduces the potential diversity benefit of the equalizer on its own. Structures which include for example linear equalizers for each antenna branch can yield better performance [13,14,19,20]. A multi-channel diversity combiner of this type is illustrated in Figure 2, where there is a tapped delay line on each antenna branch (actually assuming down-conversion has already been accomplished). This structure clearly has considerably higher signal processing cost than a simple combiner, and also has tighter requirements for channel estimation to provide performance gain. For a given number of adaptive elements, antenna arrays usually provide better performance than equalizers when the multipath has a large angular spread. However, each antenna element implies a separate down conversion path, which is far more costly than an equalizer tap, and in any case equalizers may be mandated with narrow multipath angular spreads but large delay spreads (e.g., cellular radio base stations).

There are also means to link diversity methods for transmitters and receivers without the requirements for a feedback communications system. One method uses a combination of multiple transmit antennas and either an equalizer or RAKE receiver. The motivation is to use the same single antenna handset in both mobile and indoor applications. In mobile systems, the fading is relatively narrowband, with coherence bandwidths on the order of 100 kHz, while for indoor channels flat fading over 10 MHz is not uncommon. A RAKE receiver may be used to resolve individual multipath components in the mobile setting for spreading over a 1.25 MHz bandwidth, but this hardware would be of no use in providing diversity in the indoor channel. The solution is to put multiple antennas in the base station, with transmission delays between them of the chip duration [21,22]. If the antennas are separated by several wavelengths, then the signals from each will be received by the handset with independent fading levels. Because of the delays introduced in the distributed antenna, the RAKE receiver will be able to resolve them and perform maximal ratio combining. The antennas may of course also be used for diversity reception in the base station. A similar principle can be applied to wideband TDMA systems making use of adaptive equalizers.

Multiple antennas may also be used to take advantage of the time diversity built into some systems. For example, suppose a narrowband system employs coding and interleaving to mitigate multipath. With the delay constraints inherent in transmission of speech, this diversity approach will fail for low vehicle speeds. The solution is to induce artificial fast fading. On one antenna the signal is sent with a fixed phase. On another, separated by several wavelengths, the phase of the signal is advanced at a rate just sufficient for the coding/interleaving system to perceive independent signal levels in successive decoder inputs [23]. Thus, two transmission elements may activate the entire diversity protection of the error control code, with low delay.

Channel coding and antenna arrays may be more directly coupled via space-time coding [24,25]. The standard approach of sending the same information from each transmit antenna amounts to repetition coding over the antenna branches, which is clearly not the most efficient means to provide diversity. With channel codes, the information can be spread more compactly while also obtaining a noise resistance benefit. The codes can be designed with specific numbers of transmit and receive antennas in mind to achieve particular coding gain and diversity targets, with complexity comparable to codes designed for the Gaussian channel.

II.2.3 Cooperation between Transmitter and Receiver

Up to now we have assumed the absence of feedback between receiver and transmitter. Feedback does not increase the Shannon capacity of stationary Gaussian channels [26], but it does improve the capacity for any dispersive channel and reduces the complexity of high-performance systems. Moreover, for slowly time-varying channels, the use of feedback can greatly increase the average throughput if we can tolerate variable rate transmission. Consider the following motivating example. Suppose we must transmit binary PSK at a constant symbol rate over a slowly-fading Rayleigh channel, subject to a peak power constraint so that the average SNR is 10 dB. To meet a probability of error requirement of 10^{-3}, an SNR per bit of 24 dB is required [6], which could be accomplished by repeating symbols 25 times. (With slow fading, no diversity results). Alternatively, for a Gaussian channel an SNR per bit of 7 dB is sufficient. With feedback, we could conservatively instruct the transmitter to send only when the SNR exceeds 7 dB, as happens 60% of the time. We could do even better by instructing SNR-dependent repetition. Clearly throughput is increased by a large factor with feedback, even without power control, at the cost of variable transmission rate. We now consider more powerful techniques. These may be used alone or in combination with diversity techniques, for example dealing with shadowing while the diversity techniques mitigate the multipath.

We first focus on single carrier modulation and suppose the channel to be flat-fading. Without feedback, the data rate would be restricted by the x% worst case SNR condition. However, we can do much better with variable rate transmission. The optimal system includes a channel state estimator in the receiver. Having formed an estimate of the SNR, a message is communicated back to the transmitter to indicate the size of the digital signal constellation to be used, the power level, and the code rate to be employed.

The Shannon capacity-achieving power distribution for this situation is the classic “waterfilling” solution derived originally for parallel Gaussian channels [26]. The time-slices are regarded as the parallel channels for the purposes of bit and power allocation [27]. A close approximation of this power distribution is to allocate the same power whenever the SNR exceeds some threshold γ_0, and none at all when the SNR is below γ_0. The best threshold may be determined by means of numerical integration. As for the suboptimal protocol outlined above, we allocate more power and bits when the channel conditions are good, and send nothing when they are highly unfavorable. Capacities for Rayleigh fading and log-normal shadowing are computed in [27]. With the proper power allocation and a moderate diversity order, capacity is again very close to that of a constant SNR channel.

Bit and power allocation may also be used with wideband signaling in the form of discrete multitone transmission (DMT), also known as orthogonal frequency division multiplexing (OFDM) [28]. For frequency-selective fading, this approach has the merit of achieving variable bit rate on the subchannels while having a much lower variation in the aggregate bit rate at any given time. The same result may also be achieved with wideband variable rate QAM signalling, provided an adaptive equalizer and the waterfilling power distribution are employed. Both wideband approaches amount to a combination of frequency diversity and adaptive bit allocation.

In addition to controlling parameters directly related to the modulation, feedback may be used to adjust the transmission rate either directly or through control of the redundancy of the source and channel coders. For example, in a data packet system, when the channel conditions are poor the receiver could request that either more error protection be used or no packets be sent until the SNR improves (say indicated by test packets). An example of variable rate-channel coding over time-varying channels is a scheme which uses a family of punctured convolutional codes [29]. The encoder produces all the parity bits corresponding to the lowest rate code, but only sends the data on the first attempt. If a negative acknowledgment is received, one set of parity bits is sent, corresponding to highest rate code. This process is repeated until the data is successfully decoded or we reach the rate limit. With punctured codes, the same decoder hardware may be used on each decoding pass. This scheme uses the minimum redundancy to reliably decode the data. The redundancy may be adjusted up or down based on the error rate of a set of packets.

II.2.3 Discussion

The severe effects of multipath and shadowing must be effectively dealt with to ensure robust, reliable, low-power design. A combination of diversity techniques and feedback control of variable bit and power allocation can in principle be used to completely overcome the variable propagation conditions. In practice, constraints on the delay, dynamic power range, and complexity limit the improvements that can actually be obtained. In addition, the rate of change of the channel imposes fundamental limits on the quality of the channel state estimator. For example, in mobile communications it is difficult to form accurate estimates of the instantaneous SNR. This degrades the effectiveness of all techniques, but is especially damaging to approaches that require cooperation between the transmitter and receiver since the rate of updates to the transmitter will be at least an order of magnitude slower than the rate at which updates are available in the receiver. Compensation for multipath in the receiver allows the transmitter to perform reliable allocation using an average channel state, i.e., the residual channel after diversity combining, which has smaller variations than the original channel. A mix of feedback and diversity techniques might lead to a lower overall complexity or consistently higher quality of service than reliance on only one approach. The payoffs for adaptive techniques are very high in radio systems, and as these techniques are implemented at baseband rather than RF, result in a favorable chip area tradeoff compared to simply using higher power RF front ends.

III. Cellular Multiple Access Techniques

For multiple access systems we shift the focus from coping with variable propagation effects to ways of dealing with the mutual interference caused by common access to a shared band of frequencies. We begin with conventional multiple access schemes, and then discuss how to enhance performance through interference averaging, avoidance, and cancellation techniques. We also discuss how these interact with the multipath and shadowing mitigation techniques of section II, which implicitly assume that the noise and interference do not change as a result of actions taken by the transmitter. This is in general not true in multiple access systems. Throughout this section we focus on cellular systems with connection-oriented traffic, i.e., the message duration is very long compared to the frame size. Alternative network topologies and traffic types will be considered in section IV.

The propagation losses outlined in section II with reference to equation (2) make cellular systems at once possible, and difficult to optimize. The severe attenuation with distance isolates the signals generated in one cell from cells that are far away. The large variations in signal strength within a cell force the use of power control to keep the average signal to interference ratio (SIR) at an acceptable level. At the same time, because the cells are not perfectly isolated from each other, anything done to the power level of the desired signal changes the interference experienced by users in other cells. The multiple access strategy also has an impact on interference levels and robustness with respect to interference generated in other cells. In large measure, the design of a multiple access scheme is concerned with management of the interference coupling. We now outline the main multiple access cellular strategies.

III.1 Conventional TDMA, FDMA, and CDMA

In TDMA systems, each call is assigned one time slot within a frame which it keeps until it is handed off to another cell. No other calls within the same cell are assigned the same slot, and thus users within a cell do not interfere. Frequency division multiple access (FDMA) systems assign unique frequency slots instead; as there is no fundamental capacity difference between these systems in the absence of multipath we shall not further distinguish among them. Of course, as a practical matter achieving orthogonality through good synchronization in TDMA is usually easier than achieving orthogonality with tight filtering in FDMA, and TDMA is more flexibly adapted to handle multiple bit rates. TDMA is highly spectrally efficient for single cell systems. The problem arises in cellular systems, where users in nearby cells may be assigned the same slot. If both users are on the cell boundary and experience similar shadowing, then the SNR can be persistently poor regardless of the power control strategy, as illustrated in Figure 3. The traditional solution to this problem is to assign disjoint frequency bands to neighboring cells. Regular re-use patterns of 4 and 7 frequencies are in common use to deal with this problem, although this immediately reduces the potential capacity by the same factor. Even this does not fully eliminate the problem, since large SNR margins must be built in to deal with the variability caused by shadowing. Since call blocking probabilities must be on the order of 1% over 90% to 99% of the coverage area, performance is essentially limited by the tails of the shadowing distribution, just as communication over Rayleigh fading channels is limited by the probability of experiencing deep fades. Direct attack on the shadowing margin with frequency re-use patterns, channel coding, and antenna sectorization yields relatively modest improvement in coverage because the tail of the shadowing distribution decays slowly.

A more fruitful approach is to design a multiple access system which deals with the average interference rather than the interference from individual mobiles, in effect providing interferer diversity. Spread spectrum systems may be used for this purpose. In a direct sequence system, the correlation receiver spreads the interference over the whole band. With large spreading factors, the residual interference is well approximated as white noise with a level equal to the average interference power. For multiple access, each user is assigned its own spreading code for accessing the time-bandwidth resources, and thus these systems are known as code division multiple access (CDMA). In frequency hopped systems, the same effect could be achieved by using very large spreading factors in combination with orthogonal transformations to spread information over all frequency hops. Direct sequence systems are generally operated in an asynchronous mode for the upstream (mobile to base) transmission, since maintaining synchronism to the level of a chip is usually impractical. The downstream transmission is synchronous since all have a common clock at the base station. Frequency hopped systems can be synchronous in both directions, and thus provide an orthogonal decomposition of the resources in a cell, exactly as for TDMA systems. The price of asynchronous access in the direct sequence uplink is that interference from other users appears as noise to conventional receivers, and power needs to be controlled so that the received levels at the base station are approximately the same for all mobiles. As the in-cell interference is typically in aggregate much larger than the out of cell interference, this represents a substantial performance penalty. We will discuss means to mitigate this loss and improve robustness with respect to power control variations in section III.3

CDMA systems are said to be interference-limited. Any technique which reduces the Eb/No required for reliable operation or which reduces the average interference directly increases the capacity. Thus, techniques which mitigate multipath lead to large improvements in capacity, since the Eb/No required for reliable operation is greatly reduced. Likewise, a coding gain of 3 or 6 dB will improve capacity by factors of 2 or 4 respectively, irrespective of the shadowing distribution. Exploitation of voice activity is also relatively easily implemented; if no power is transmitted during silent periods, the interference is diminished. Sectorization reduces the number of interferers in view, and thus capacity linearly increases with the number of sectors (neglecting the antenna sidelobes). The techniques which raise the capacity of DS CDMA may also be employed in frequency- and time-hopped systems, provided the basic approach is to average over the interference. They are also available in some degree to hybrid systems, such as GSM, which have a basic TDMA structure together with slow frequency hopping.

III.2 Enhanced TDMA Systems

We now discuss how better performance may be obtained in TDMA systems in the multi-user environment using interference averaging and avoidance. In this section, we consider techniques such as channel probing, dynamic power, bit, and channel allocation, and adaptive antenna arrays as means to improve the capacity of TDMA or slow-hopped TDMA systems, including discussion of how these techniques interact with each other and the diversity methods of section II.

III.2.1 Dynamic Channel Allocation

Shadowing and average loading can be taken into account in the frequency planning of TDMA systems using quasi-static channel assignment algorithms [e.g., 30]. Significantly greater capacity improvements can however be obtained using dynamic channel assignment (DCA) algorithms, even with the suboptimal power control strategy of equal received levels [31]. These algorithms slot users so that the mutual interference is below some fixed threshold for every user, and re-adjust assignments as channel conditions change. With a very large number of time slots, DCA will cause the minimum SIR for each slot to be approximately equal, and this SIR will be at least that obtained with a synchronous CDMA system with the same power control strategy. When most slots are occupied, DCA will not perform much better, since DCA does little to reduce interference. However, for smaller occupancy, users have more freedom to *avoid* mutual interference, and performance improves. Hybrid strategies where some channels have fixed re-use and some are subject to dynamic allocation are possible, including also traditional borrowing of the "fixed" re-use channels to accommodate non-uniform loading.

The general problem of finding the optimum assignment of all users in the sense of maximizing the minimum SIR is very computationally intensive. Instead various ad hoc techniques can be used, see e.g. [32]. In a simple algorithm, slots with low total interference are checked to see if the target SIR of γ can be met when admitting a new user, for all users in the slot. If not, more slots are tried until the user is admitted or all slots are exhausted. In a variation, first invoke the above procedure. When the slots are exhausted, try again, this time with the option to nudge out one user already in the slot, if this user could be admitted in another slot using the simple algorithm. If not, nudge the next user. This continues until admission, or all users have been nudged in all slots. This procedure is far more computationally intensive, but leads to only modest performance gains.

The main benefit of DCA is that it transforms a TDMA system into one that is interference-limited in the sense of CDMA, and thus techniques which reduce the Eb/No required for reliable operation directly improve the overall capacity. Shifting the problem to try to find assignments such that each user satisfies its own SIR requirements, it is possible to derive distributed forms of DCA [e.g. 33]. These achieve performance similar to centralized omniscient algorithms, with only imperfect local measurements of the channel conditions, and no cooperation among base stations. This comes at the expense of additional convergence time, and relies upon the local nature of the interference coupling. Distributed solutions, when they exist, are desirable from many points of view, not least because they require minimal revision of existing network software.

A number of distributed DCA algorithms are compared in [34], under the assumption that interference levels are easy to measure but SIR is not. The one leading to the largest capacity is the well-known least interference algorithm (LIA), which is implemented in the cordless phone standards DECT and CT-2. It chooses the channel with lowest measured interference. Thus, in effect it rewards interference avoidance. This outperforms more complicated algorithms with adjustable thresholds based on heuristics such as trying to pack users compactly into channels. As we shall see, interference avoidance is also an effective heuristic for power and bit allocation.

III.2.2 Dynamic Power and Channel Allocation (DPCA)

We now consider the joint optimization of power and channel assignments, beginning with the problem of optimal power allocation among users competing for one channel [35-43], here following [38]. Consider the situation depicted in Figure 4. Let G_{ij} be the propagation gain between the j^{th} base station and the mobile to which the i^{th} base station is communicating directly for some time slot. Let P_i be the power transmitted by the user belonging to base station i. If $P_i=0$, there is no mobile communicating with base station i in the slot. The SIR for basestation i is then given by

$$R_i = \frac{G_{ii}P_i}{\sum_{j \neq i} G_{ij}P_j} \qquad (3)$$

For reliable operation, we require $R_i > \gamma$, for all basestations. This set of conditions is equivalent to the following set of linear inequalities

$$G_{ii}P_i - \gamma \sum_{j \neq i} G_{ij}P_j > 0 \qquad (4)$$

where we must find whether there exists a set of positive powers $\{P_i\}$ for which all the inequalities are satisfied. It may be shown that this problem is equivalent to solving the matrix equation $\mathbf{G}\Pi=\mathbf{1}$, where Π is the vector of power assignments, $\mathbf{G}_{ii} = G_{ii}$, and $\mathbf{G}_{ij} = -\gamma G_{ij}$, with appropriate renumbering of entries if not all basestations are active. Then if every component of the vector Π is positive, there is a feasible assignment. Moreover, the vector of powers is the best one to use apart from a scaling factor determined by the background noise, in the sense that it lies in the center of the region of feasible power assignments, and thus is least sensitive to perturbations.

To apply this single-channel algorithm to a multi-channel environment, we must in the worst case check among all possible combinations of users and find the assignment that leads to the largest number of users being admitted. We may instead use a DCA algorithm. As for DCA, fully distributed forms of DPCA exist [39], in which the only information exchange required is that between a mobile and its own base station, revealing their measurements of the interference power and the propagation loss in that single link. If the SIR is below the target, the receiver signals the transmitter to increase the power by a small fixed amount in dB; if above the target, it signals for a decrease. This amounts to a gradient search on SIR, and has a geometric convergence rate. Specifically, all users adjust their power according to the following rule:

$$p_i(t+1) = \frac{p_i(t)\,\delta\gamma_i}{r_i(t)} \qquad (5)$$

where $p_i(t)$ is the power at time t, $r_i(t)$ is the measured SIR at time t, δ is a small safety margin to protect active users (those already admitted) and γ_i is the target SIR for user i. Numerous variations upon this basic strategy are possible, e.g. [43].

In the absence of dynamic range limitations on the power control, the algorithm not only converges to the optimum power vector if a feasible solution exists, but it also rejects new users in the absence of a feasible solution. Dynamic power range limitations can be accommodated by either limited centralized network control operations, or through voluntary termination of admission attempts when the resistance of other users is apparent from rapidly increasing interference levels [40,41]. Finally, while the above implicitly assumes that the mobile should be assigned to the base station with least propagation loss, this is not necessarily the best assignment from the point of view of network capacity; the optimal hand-off conditions are presented in [42].

III.2.3 Channel Probing

One way to perform DPCA is to run the distributed algorithm for a set of channels, choosing one that appears to yield a feasible solution. This is undesirable for two reasons: it disrupts the power levels for other users, and takes a long time, reducing robustness to mobility. Channel probing is a means to speed up the call admission process while causing minimal disturbance to other users.

When the only additive impairment is noise, probing is simply a matter of learning the noise level and the channel impulse response. From these measurements, the best achievable SNR can be computed, and the appropriate combination of symbol rate, power level, and constellation size can be chosen. However, in systems with distributed adaptation of power, an increase in the power level of a new user will cause a reaction of increased power by the other users sharing the channel. The aim of probing is therefore not to measure the channel as it is, but to predict the SIR after the power control algorithm would have reached the steady state.

A very small number of measurements suffice to predict the final power level. Having measured the propagation gain (loss) on one’s own link, a few power control steps establish the starting interference level and the change in interference as a result of increasing the power [41]. To probe the channel, let $q_i(0)$ be the measured received interference value before transmission begins. Let $q_i(1)$ be the measured interference value and $s_i(1)$ the measured received signal power value after a fixed power level has been transmitted, and the other users have adjusted their power in response. Then the maximum feasible SIR is estimated as

$$\hat{\gamma}_{max} = \frac{s_i(1)}{q_i(1) - q_i(0)} \qquad (6)$$

As there will be measurement errors, several iterations are needed in practice to form reasonable estimates of the final SIR, but probing is orders of magnitude faster than going through the complete adaptation. Since it may be accomplished at relatively low power levels, it results in minimal disruption of active users. Probing may also be used to predict the power level required to sustain any particular target SIR [41].

Another approach to probing and power allocation is to partition channels amongst dynamic and fixed re-use assignments. The motivation is as follows. Users in certain locations, such as near the edges of cells, will always require large re-use distances; indeed fixed re-use patterns are designed based on the needs of these users. The interference they generate is felt over many cells. This will tend to strongly couple distributed power adaptation, slowing convergence in many cells and leading to infeasible solutions in immediate neighbors. Thus, if these users could be easily identified, removing them from contention for dynamic allocation channels would not reduce system capacity and could lead to faster convergence. Having a set of fixed re-use channels also permits rapid call admission and gives better guarantees of continued service in hand-offs.

A channel partitioning scheme which embodies these ideas is as follows [44]. Each slot is divided into a short probing phase and a longer data phase. Probing and power control are as described above, with all power control based upon the probing segment of the slot. There are two classes of channels: fixed and regular. Fixed channels have fixed frequency re-use, while regular channels (which are much more numerous) are subject to dynamic allocation. All users are initially admitted in fixed channels. The probing segment for the fixed channel however is set to a free regular channel. If probing indicates that the SIR will be above threshold, the user is moved into the regular channel. Fixed slots thus provide immediate admission, with probing providing a means to rapidly transition a call to a regular channel. If probing is unsuccessful, the call remains in a fixed slot. This procedure shows increased robustness to high mobility compared to a pure dynamic allocation system, at the cost of slightly reduced capacity in static systems. These fixed slots may also be subject to a slower time scale reassignment among cells, to accommodate long-term non-uniformities in the cell loading.

III.2.4 Dynamic Bit Allocation

In a frequency hopped system, there are additional degrees of freedom beyond choosing power levels and channel assignments. If hopping patterns are aligned across cells and we assign the same power to every slot, then the DPCA problem is the same as for a simple TDMA system, except that hopping affords the possibility of coding over slots to provide diversity protection. We may alternatively randomize the hopping patterns over cells [45], in which case coding also provides some diversity protection with respect to the varying interference levels. Since in this case the choice of channel is less critical, time spent probing can be reduced. The former approach leads to a greater possibility of interference avoidance, but is less robust in mobility than the latter.

The above schemes do not fully exploit the interference avoidance possibilities inherent in frequency hopped systems. It is not necessary to transmit the same number of bits in every hop, or indeed to use every hop. Consider for example a system in which hopping patterns have been randomized among cells. The digital constellation size is fixed, but we can choose in which M out of N slots in a hopping pattern to transmit bits. Coding is performed across the hops so that we are most concerned with the average power, although with a fixed constellation size minimum SIR requirements remain unless large redundancy is contemplated. Then the allocation of slots can be accomplished by using channel probing and choosing the M slots leading to the lowest power consumption. This choice causes users to arrange themselves for near-minimum mutual interference, i.e., it promotes interference avoidance [46]. Lower M permits more avoidance, but increases the bandwidth requirements for each user. Th optimal ratio of M/N can be determined by simulations.

Larger capacity is obtained if we can allocate a variable number of bits to the slots. As noted in section II, the optimal allocation of bits and power when dealing with signal-independent impairments is the waterfilling distribution with respect to the noise floor. However, in multiple access systems the variable power allocation also changes the interference. In [47] it is shown that the allocation problem is NP hard (as is DCA), and so several ad hoc methods are attempted. The most successful are those which aim to achieve maximum interference avoidance, rather than a minimization of average transmitted power--that is, users tend to occupy as small a number of slots as possible consistent with the maximum constellation size permitted, with some smaller number using medium size constellations, and many with no bit allocation at all. Two algorithms which give good performance are reverse waterfilling based on maximum estimated SIR, and minimization of the total number of slots in the hopping pattern that are occupied, subject to having an acceptable SIR. In both cases, similar error rates are found in all hops for which bits are allocated, through the joint adjustment of constellation size and power level. Here accurate channel probing is essential, although channel coding over slots can permit reliable communications even if there remains some variability in uncoded error rates due to small errors in the SIR prediction.

III.2.5 Adaptive Antenna Systems

Antenna arrays may be used for many different purposes. Because the use of antenna arrays need not be specified in a standard, the performance of any system can be upgraded through the effective use of multiple antennas in the transmitter and receiver. Moreover, the capacity improvements are limited in principle only by such factors as physical size constraints on the platform and the ability to track channel variations, although of course complexity and power consumption are also issues in practice. In general, antenna arrays give some combination of interference reduction due to beamforming (e.g., sectorization), interference cancellation, and multipath rejection.

In beamforming, directional elements or a phased array of (approximately) omnidirectional elements are employed to provide antenna gain over a set of preferred directions, with suppression of signals in all other directions. For M elements, a main beam size of roughly 360/M degrees may be formed. If used in a base station, the pattern might be fixed to provide sectorization. For example, supposing there to be no antenna sidelobes in a 3-sector system, only those mobiles in a particular 120 degree sector could contribute interference. In an interference-limited system such as CDMA this immediately results in roughly a three-fold improvement in capacity [18]: the density of interferers could be increased to balance the interference reduction.

Something less than a linear increase in capacity with the number of sectors will be obtained in practice. Conventional TDMA also sees a roughly linear reduction in the average interference seen in each slot, but the probability of damaging interference is not reduced enough to allow for large improvements in frequency re-use based on sectorization alone. For interference-limited systems, the capacity increase is less than linear because of antenna sidelobes.

Multiple antennas may also be adaptively controlled in a base station to track individual mobiles, e.g. [48-50]. Adaptation may be accomplished by any classic least squares algorithm.

For cellular radio, most of the significant multipath components of the signal received from the mobile arrive with a small angular spread. Consequently, a feasible alternative is to pre-compute a set of beam patterns and store the antenna weights, switching between patterns as the mobile moves [51]. This technique can also be applied to analog standards; one experiment with 24 antenna beams indicates an average SIR improvement of 5 dB over a three sector dual diversity system [52]. For interference-limited systems, arrays of this type can potentially greatly increase the capacity by re-using frequencies within the same cell. This approach is however not viable for many indoor systems, since the multipath can have very large angular spread, and the base station antennas are likely to be electromagnetically coupled to their surroundings. In this case, the beam pattern must be adaptively constructed.

We may alternatively view narrow sectorization as interference cancellation, in which the receiver adapts the weighting of the elements to completely cancel an interferer. Physically, this corresponds to placing a null in the beam pattern in the direction of the interferer, and a maximum in the direction of the desired signal. With M elements, M-1 interferers may be cancelled in this way. Thus, M users may share a cell, exactly as from the point of view of sectorization. However, in practice the performance improvement is less than it would first appear. There may be many multipath components which arrive from different directions, for both the desired signal and the interferers. To actually cancel interferers, all the large multipath components must be tracked, and these may arrive from a wide range of angles, particularly for indoor systems. The adaptive algorithms will not be able to distinguish between the two tasks of diversity protection and interference cancellation; they simply dictate the weighting that results in the best SIR.

III.2.6 Interference Cancellation

Error control codes may be used in the fashion of antenna arrays for interference suppression. For example, suppose M users in a single-cell TDMA system employ the same rate-1/M repetition code and share the same slot. Then as for an antenna array with M elements, one can solve a set of M linear equations and cancel the M-1 interferers. Unfortunately, no capacity has been gained due to the redundancy of the repetition code. However, the method extends to linear channel codes with a lower redundancy penalty [53]. Consequently, channel codes of this type may be used alone or in combination with antenna arrays to reduce interference. As for antenna arrays, there is a trade-off between interference rejection and effective diversity order. The potential implementation advantages of channel codes are the absence of multiple down-conversion paths, and the ability to potentially achieve large rejection for platforms that are too small to support multiple antennas.

Another approach to interference mitigation is joint detection of signals, which has received most attention in the context of asynchronous CDMA [53-61]. Multi-user detection can also be of use in TDMA, as illustrated by the following example. A base station in a sectorized system may treat the signals received from each sector separately, with the signals picked up on the antenna sidelobes being regarded as noise. But in fact more information is available, since much of the "noise" is actually demodulated in the other sectors as the desired signals. A receiver which jointly estimates the three desired signals could achieve better performance, and largely mitigate the effects of antenna sidelobes. Note that for improved performance we do not necessarily need to perform perfect joint estimation, since we could also perform independent interference cancellation in the three sectors. This principle applies equally well to systems which use adaptive antenna arrays in the base station. In the limit of a large number of elements, some of the simplified algorithms proposed for CDMA systems may become appropriate.

Indeed, the term "interference limited" is somewhat misleading in light of the possibility of interference cancellation. In [62] it is shown that if a receiver has perfect channel state information for the desired signals and all the interferers, then by the use of diversity and channel coding the capacity will be limited only by the noise. In reality, the degree to which interference limits capacity is a function of the receiver complexity and our ability to estimate the signal and interference channels, the latter being connected both to complexity and the channel dynamics. Fundamentally, we are limited only by noise, bandwidth, and channel variability when messages are long.

III.2.7 Interaction of Adaptive Methods

We have so far described methods to adapt the channel assignment, power level, bit allocation, antenna beam pattern (and in section II) the equalizer, with receiver-based or distributed algorithms, and probing to speed up the adaptation. Since there are conditions in which each technique will fail to provide the desired benefit, to design a robust system it is desirable to include a variety of these techniques. The question naturally arises to what extent they will work together or against each other, in a multiple access system coupled by mutual interference.

We consider first a system in which the transmitter and receiver both have adaptive antenna arrays, adaptive equalization, and a wide dynamic range of power adjustment. Adjustment of either the power or the transmit beam pattern changes the interference seen by everyone else, and also modifies the impulse response seen by the intended receiver through selective weighting of multipath components. Adjustment of the receiver beam pattern changes the interference levels and the impulse response, changing the power level required for reliable operation and the residual channel seen by the equalizer. Thus, the algorithms are coupled at two levels: internally, with respect to the desired link, and externally, via the interference coupling to other users. There are many possible ways to combine power control, antenna arrays, and adaptive equalizers, not all of which are useful. We will outline some insights arising from a study for high-speed indoor radio [20,63].

Closed form solutions can be derived for the optimal weights of the receiver array given any particular transmitter pattern, but the joint optimization is already difficult for an adaptive transmitter array and receiver for a single link. The problem is even more difficult when the interferers also have adaptive arrays, since for strongly coupled links (large gains) an action by one produces a large reaction in the others. Figure 5 illustrates one simple scenario, where users (1,2) and (3,4) are the communications partners, and the barriers A-F are highly reflective. Some of the principal propagation paths are shown, with interference being caused only by multipath in this example. Suppose the link (1,2) is silent, and the users (3,4) have adaptive arrays but no equalizers. Then users 3 and 4 should adapt their arrays to suppress the multipath off barriers C and E and put high gain in the direct connection. However, suppose now that users 1 and 3 transmit in odd intervals and users 2 and 4 in even intervals. Then during reception user 3 must additionally suppress the interference from user 2, while user 1 need only be concerned by multipath from user 2 off barriers A and B. User 2 must suppress both multipath from its partner and multipath interference from user 3, while user 4 will be concerned only with multipath from user 3. If user 2 succeeds in nulling out the interference from user 3 and re-uses the same antenna pattern during transmission, then user 3 will experience very little interference from user 2 and can concentrate on multipath suppression. As there is no closed form solution to the general form of this problem, the approach taken in [20] was adapting the receiver using the LMS algorithm, with receiver weights re-used for transmission. The heuristic behind this choice is that in a time division duplex system the nulls placed in the direction of interferers during reception will also reduce interference transmitted towards the other users who caused most of the interference. Thus, receiver adaptation will also serve to reduce the interference generated. Additionally, gain in the direction of the communications partner is preserved, so that SIR is improved for the link. These two effects also serve to decouple interactions, promoting more rapid convergence.

One interesting result is that when users in the same cell are prohibited from using the same channel the system always converges. By restricting channel access in this way we insure that the dominant factor in adaptation is the desired link, rather than any particular interferer. Thus, such restrictions decouple the system, permitting rapid adaptation by least squares (gradient descent) methods. When this restriction is released, there are many occasions in which adaptation fails to converge. This is unsurprising, since there is no prior expectation that the error surface should be quadratic. One can also ensure stability by using an approach similar to that used in distributed power control algorithms, namely, voluntary drop-out when little progress is being made in SIR levels. This mechanism also decouples the system, leading to stable convergence by the remaining users. Indeed, every distributed iterative method based on gradient descent that affects transmission will fail in some circumstances in multiple access systems, and so there must be some non-linear rule that either limits coupling in the first place, or makes a decision to halt adaptation by some subset of the users based on their priorities. That is to say, while in most circumstances gradient descent methods appear to work, we must either engineer the system so that the exceptions are extremely rare, and/or develop a procedure which determines when one of these exceptional circumstances is about to occur, so that a back-up procedure can be used (e.g., try another channel, go to a non-adaptive mode).

Adding equalizers to the system greatly expands the number of configurations possible. In multichannel combiners (MC) there are linear equalizers for each antenna branch, and then after combining, either MLSE or a DFE. Alternatively, all equalization can follow adaptive antenna combining. The MC deals with the distinct impulse responses seen in each antenna branch, and (in slowly changing channels) has better performance than structures where the equalizer follows simple combining. However, the complexity is higher and it is difficult to produce appropriate beam patterns for the transmitter array. For the indoor environment simulated in [20], there was little performance loss from using a simple combiner followed by equalization, so that the receiver antenna weights could be re-used for transmission.

Adding power control to the mix adds further interactions. A change in the beam pattern even with constant radiated power changes the power measured at all receivers. An interesting result is that if the antenna weights are renormalized to maintain constant gain in the direction of the desired users, then including small power control steps actually improves convergence of the distributed beamforming algorithm [63]. The intuition is that power control directly determines the minimum gain required for the desired link, quickly reducing unnecessary interference, and thus the coupling in adaptation among cells.

Probing to estimate the maximum SIR is made much more difficult by the presence of adaptive transmitter and receiver antenna arrays. When probing at low power levels, receivers can easily slightly shift their beam patterns to almost completely remove the effect of the new interference. Later, at some higher power level, shifting the beam pattern is insufficient and the transmitted power levels increase. Nevertheless, a simple model which assumes that most interference is due to one interference path predicts the final SIR with remarkable accuracy [64], provided the network has some explicit re-use restrictions (e.g., users sharing the same channel cannot be in the same cell).

An intriguing question is whether there are distributed non-linear optimization methods that can efficiently solve the problems, without resort to ad hoc procedures. There are several problems in the path to such a prospective procedure. First, many of the subproblems are known to NP hard (e.g. bit allocation, DCA), and so even centralized algorithms will in some sense be ad hoc. Nevertheless, there appears to be room for a more rigorous mathematical analysis from the point of view of coupled dynamical systems, the results of which may suggest new algorithmic approaches to stable and rapid convergence. Second, from the implementation point of view, nonlinear methods are more computationally intensive than gradient descent approaches, and have not found significant application even in the simpler situation of single channel communications. The menu of low-complexity, rapidly converging algorithms is limited. Third, it is impractical in most circumstances to gather together information about the state of the system to attempt a globally optimum solution, even if it were computationally possible. In a cellular system, this would for example mandate re-writing network software, coordinating probing among numerous cells to determine actual mutual interference coupling, and (with antennas) attempting to characterize the local multipath. Each of these on their own are unlikely prospects, but all would be needed for a system which optimally determined channel assignments, power levels, antenna settings, and equalizer coefficients.

III.2.8 Summary of Enhanced TDMA Techniques

Despite the great variability of radio channels, many things can be done to engineer a multiple access radio system to behave in a stable fashion. In conventional TDMA, most deleterious interference conditions are eliminated by fixed frequency re-use patterns, with the back-up procedure of dropping calls when this fails. This has the merit of simplicity of implementation, as the interference coupling has been made weak enough to enable each cell to make independent decisions on resource access. It unfortunately also greatly reduces the efficiency of spectrum utilization in the network from what is possible through such techniques as dynamic power, channel, and bit allocation, and adaptive space/time techniques. These latter techniques exploit variations in the radio channel with angle, frequency, and location, but to yield higher capacity must necessarily deal with much stronger interference coupling in the network. Much of the time distributed algorithms based on gradient descent techniques suffice to produce stable convergence, but mechanisms are required to recognize conditions which may lead to instability, and then to deal with these situations using a back-up procedure. In effect, we apply one set of rules to the main part of the probability distribution, and another set to the (pathological) tails, with some effort spent in recognizing the situation (e.g., through probing). While at this time many of the solutions are ad hoc, in aggregate they promise orders of magnitude capacity improvement using distributed iterative algorithms.

III.3 Enhanced Direct Sequence CDMA

The uplinks and downlinks in direct sequence CDMA systems have different characteristics. On the downlink from base station to mobiles, all signals are synchronous apart from the effects of multipath. Consequently there is little mutual interference among users sharing the same cell, and there is considerable flexibility in the power control strategy. In conventional DS-CDMA, on the uplink the signals are asynchronous, and thus the other signals appear as noise, suppressed by the processing gain. A standard result is that the capacity maximizing solution is to control the power so that the received signal levels at the base station are the same for each mobile [65]. Deviations from this lead to considerably reduced capacity.

As noted previously, CDMA is an interference-limited system, and thus any measure which reduces the interference or the SIR required for reliable operation directly improves the capacity. Essentially all the methods discussed in the context of TDMA systems have a CDMA counterpart, although except in the areas of antenna arrays, joint estimation, and power control relatively little research has been pursued in these directions, in part because only the former techniques are compatible with existing standards. We will outline some of the possibilities and difficulties with these methods in the following subsections.

III.3.1 Multiuser Detection (Joint Estimation)

In asynchronous CDMA systems, joint estimation of the signals has two main advantages: (1) increased capacity from suppression of in-cell interference and (2) reduced dynamic range in power control. In principle, power control is not required at all for a single-cell system using perfect joint estimation. In practice, thermal noise and out-of-cell interference will bring some power control back into the picture, as well as limiting the capacity improvement [61,66]. The optimal joint estimator has complexity that is exponential in the number of users, and requires very tight estimation of the channels for each user [54]. However, with the prospect of increasing the capacity of CDMA systems without changing standards, a number of algorithms have been developed which achieve some performance gain over conventional receivers with more modest complexity costs [54-61]. For example [59], if all transmissions are synchronized to within one chip period, the matched filter receivers in the base station may be replaced with relatively simple decorrelating receivers [56], achieving most of the performance gains available with optimal detection. Of course, against this must be balanced the complexity cost of more accurate synchronization.

A useful principle which avoids the need for tight synchronization or estimation of the timing is a subspace decomposition of the set of interfering signals. Using the minimum mean square error criterion, for each desired signal a projection operation is performed onto a subspace determined only by the spreading waveform for that particular signal [58,60]. This operation requires only second order statistics and can be performed independently for each user, without knowledge of the interferering waveforms, with and without multipath. The receivers are very resistant to variations in the power levels of the interferers with respect to the desired signal. An important issue is the speed with which the channel variations (including multipath) can be tracked, as the number of coefficients to be estimated for each user can be as large as the spreading factor. Reduced complexity forms with a smaller number of adaptive coefficients provide reduced performance in slowly changing channels but increased ability to track channel variations. The use of orthogonalization preprocessing to decorrelate the interference can speed convergence of standard gradient descent techniques such as the LMS algorithm, and enhance tracking ability [67].

The other main approach to multi-user detection is explicit interference cancellation. For example, in [57] a partial decorrelation of the channel is performed which results in one user being free of interference. The second bit has interference only from the first user, and thus can be made interference free by subtracting the decision for the first one, and so forth until for the last user all previous decisions have been used. In [61] it is pointed out that the various detectors have similarities to equalizers. The decorrelating receiver is similar to a zero forcing linear equalizer, the MMSE solution to a linear equalizer, and the interference cancellation approach to a zero forcing DFE, with the appropriate differences in complexity and performance according to the accuracy of the channel estimates. The analogy arises because the multi-user interference is in many ways similar in its undesired effects to intersymbol interference. The large number of equalizer structures and adaptive algorithms are thus available to be adapted to this new setting.

III.3.2 Antenna Arrays and Space-Time Methods

Sectorization has been proposed as an integral part of the IS-95 standard from the beginning, as (apart from imperfections in the beam patterns) it leads to an almost linear increase in capacity with the number of sectors. Each sector views fewer potential interferers, and radiates power to a smaller area. Three sectors per cell have been proposed, with adaptive techniques being required for a variety of physical reasons if this linear increase is to be maintained with higher directionality of beams [68,69].

Sectorization does not provide multipath immunity. As in TDMA systems, this can be achieved with space time methods, which in the context of CDMA amounts to a combination of adaptive antenna arrays with RAKE receivers. Variations upon this theme are sometimes termed smart antennas. A single RAKE receiver may follow the antenna array combiner, or there may be a separate RAKE receiver in each antenna branch, for potentially better performance but considerably higher complexity. All operations for the RAKE receivers must run at the chip rate, which is orders of magnitude higher than for the corresponding TDMA systems if large spreading gain is to be achieved. Consequently, there is a large incentive to reduce the number of adaptive elements.

As noted previously, one approach to this is to create a large number of narrow beams using antenna arrays in the base station, and switch among these precomputed patterns according to which one yields the better SIR. As the angular spread of the multipath is typically not that large at the base station and there is little electromagnetic coupling to the immediate environment, this can be quite successful. We may then follow the beamforming with a single (standard) RAKE receiver, or use multi-user detection techniques. While not providing additional diversity protection in either direction of transmission, it provides the benefits of very narrow sectorization.

III.3.3 Dynamic Resource Management

As in TDMA systems, users require vastly different power levels to achieve an acceptable SIR. The high power users generate a large amount of interference to neighboring cells. Since with conventional CDMA systems it is the in-cell interference that limits capacity, there has hitherto been little incentive to deal with what would be problem interferers in the TDMA context. However, should effective multi-user detectors be used, these interferers will be the limiting impairment. In this case, partitioning the channels into those which are used in every cell and those which have some frequency re-use restrictions may become attractive. In this case probing would serve to classify to which category of channel the users should be assigned.

III.3.4 Similarities to Enhanced TDMA Techniques

Conventional TDMA manages the shared resource through separation of users within a cell onto orthogonal channels, reducing out of cell interference through large frequency re-use distances. This amounts to interference avoidance through channelization restrictions. Conventional CDMA improves upon this by using interference averaging for both in-cell and out of cell interference; since the system is now interference limited, techniques such as sectorization, exploitation of voice activity, and channel coding can be used to either reduce the interference or the SIR required for reliable operation. These lead directly to improved capacity. The enhancements to TDMA described in III.2 relied upon a combination of interference averaging and interference avoidance, with the effect that TDMA also becomes interference limited. Dynamic interference avoidance is a more powerful means than averaging to increase the capacity of a cellular system, but demands more active control of the channel resources than does averaging, and is consequently less robust with respect to rapid channel variations. Explicit interference cancellation can be applied to either TDMA or CDMA systems, representing one more escalation in receiver complexity, through its requirements for some combination of antenna arrays and joint detection of signals. Interference avoidance may also be a viable strategy for CDMA if multi-user detection is used. Thus, as one climbs the complexity/performance ladder, TDMA and CDMA systems will resemble each other in many ways, and one might consider a hybrid approach as a means to leverage the large technical efforts that have gone into both approaches.

IV. Alternative Access Scenarios

So far we have assumed a cellular topology, with long holding times for the calls. In this section we consider alternative network topologies and traffic types, and their influence on design requirements for high capacity communication networks.

IV.1 Network Topologies

The network topology has a profound influence on the capacity of the communications system. Even with a fixed infrastructure as in cellular systems, there are a variety of means to organize access to the backbone network. Other access scenarios include peer to peer networks, and multi-hop networks. In each case, the challenge is to deal effectively with the very sharp drop-off of power with distance and the large variability of the radio channel.

IV.1.1 Macro Diversity in Cellular Systems

A powerful means for increasing the capacity of a wireless network is to provide the users with many access points, i.e., macro diversity. This is at the heart of the cellular concept; calls are assigned to base stations for which the propagation losses are small and for which channels are available. Cells are planned by placing base stations in locations that minimize the effects of shadowing. Calls are handed off to neighboring stations when the signal quality is degraded or when a cell site becomes too busy. To first order, capacity is increased linearly with the number of base stations in the network, but of course costs also increase. In addition, handoffs become more frequent, using more network resources and decreasing the benefit of further cell splitting. Thus, the cost of increasing the number of cells must be balanced against the cost of higher capacity access methods. This tradeoff also depends to some degree on the scale of the cells. Microcells with base-stations potentially below the level of buildings behave very differently than macro-cells in terms of propagation characteristics [70], and pico-cells within buildings have vastly different characteristics [71].

Shadowing causes considerable problems in cell planning, since hills, valleys, urban canyons and the like can cause holes in coverage or anomalously high interference across many cells. One method for both boosting capacity and mitigating shadowing is to distribute antenna sites around the cell [72]. Then mobiles may receive the same signal from different directions, providing diversity protection against shadowing and distance propagation losses. The base station also achieves diversity by selecting the received signal with the highest SIR (or by using more complicated combining methods). No handoff is necessary so long as the mobile stays within the cell. This system decreases the dynamic power control range required, and also reduces the interference to nearby cells. Thus, the capacity of interference-limited systems will be improved, and the cost of the mobile units may be reduced since lower power electronics become feasible. Minimal signal processing is performed at the antenna sites, which essentially function as relays. Additionally, the frequency of handoffs is reduced since signal quality stays high until near the cell boundary, resulting in less switching back and forth due to local shadowing conditions [73]. This concept has also been applied to buildings, with substantially improved capacity being reported for a CDMA system [22].

One successful form of macro diversity is soft-handoff in DS-CDMA [74]. During handoff, two base stations listen to the mobile, and both transmit to it, with the mobile able to combine the signals using its RAKE receiver. This significantly expands the coverage area for base stations under a peak transmitted power constraint. The ultimate form of macro diversity would have joint processing of signals from *all* cell sites [75]; on the uplink, everyone is effectively in soft handoff with every basestation, all the time. With this information, interference from other cells could be cancelled with techniques similar to the multi-user detection schemes noted earlier. The large attenuation of signals with distance suggests that local collaboration should be sufficient to achieve most of the available gain, but this remains a topic for future research. The problem of joint transmission to maximize the capacity of the downlink is even more complicated. For a single cell, the problem is closely related to coding for broadcast channels, with the general form of the solution being signaling via superposition of messages, and a power control strategy quite different from the ones outlined in earlier sections [76]. The multi-cell problem is so far unsolved.

IV.1.2 Mobile Multi-hop Networks

An even more challenging set of problems arise when there is no fixed infrastructure available for communications. Examples include disaster relief and military operations, in which team members must coordinate activities over large geographic areas. Others include low power sensor networks in which the objective is to minimize the communications energy cost. In such circumstances, the optimization criteria are quite different from ordinary wireline or cellular network design. For example, consider a team fighting a wildfire on mountainous terrain. It would be highly desirable if team members could communicate by voice to one another with low-weight devices for hands-free operation. One possible scenario is illustrated in Figure 6. Because of shadowing and propagation loss, not all team members may directly communicate. The possible links are indicated by a solid line between nodes. The maximum transmitter power and minimum data rate constraints together with the topography dictate which radios can establish direct links. Otherwise, team members may be connected only by multihopping. Since the team members are in motion, the network must be dynamically configurable, and for safety reasons must be fault tolerant. Thus, radios must in a distributed fashion collaborate to share base station functions such as routing, congestion control, and synchronization so that communications bottlenecks are avoided and a good compromise is reached between throughput and power consumption.

Again referring to Figure 6, users A and D wish to communicate at the same time as users B and E. There is no direct path between users B and E, and we suppose radios can relay or receive only one message at a time. Thus, the feasible connections are AD and either BCE or BCFE. On each relay the transceiver consumes power and is unable to process other messages, and so it is desirable to use as few hops as possible. On the other hand, by using BCFE less interference will be generated to the connection AD since CF is a much shorter path than CE, and less transmit power needs to be allocated. If disjoint frequency bands are used for these transmissions, then interference is of less concern; however, we must now control the frequency assignments to avoid conflict. If channelization is achieved using asynchronous DS-CDMA, then the limited processing gain must be considered when deciding on routing and the compatibility of virtual connections in the network. Thus, in either case routing and management of the multiple access scheme resources may not be separated. Moreover, since the network has variable topology and must be fault tolerant, it must be managed in as distributed a fashion as possible.

A centralized access control strategy will fail in a multi-hop network. Large overhead and latency will be incurred in conveying information about local traffic and propagation conditions over multiple hops to the controller. This also wastes energy and can exhaust nodes which serve as gateways. At the same time, fully independent decisions are not desirable in all matters. For example, suppose we would like to implement a frequency hopped system. The peers themselves can and should regulate power, bit assignments, antenna patterns, and other matters for which efficient distributed algorithms exist. However, as discussed in section III synchronous access together with enforcement of orthogonal channelization in local regions can significantly decrease the mutual interference. Thus, not all nodes should be equal at all times. It is also useful to define a hierarchy of nodes for update of routing tables, timing, channel assignments, and other matters that are of concern to larger portions of the network. The hierarchy should reflect the locality of the parameters being controlled--control data should move only over the radius where it has an appreciable influence, and otherwise the algorithms must operate in a distributed fashion. Most control traffic is exchanged between communications peers, with decreased control overhead contributions as we move up the hierarchy. This framework has the further advantage of increased scalability compared to either centralized control or flat networks.

One useful construct for implementing the middle level of control functions is the cluster [77,78]. Nodes which may communicate with each other in a single hop as peers are grouped together in a unit, designating one node as the cluster head. This designation may depend on some combination of factors such as average energy cost to communicate with it and its remaining energy reserves, or may in some fashion be arbitrary. Clusters are different from cells in two major respects. First, almost all traffic flows between the peers. Second, the designation of cluster membership is dynamic, and so is the designation of the cluster head. The cluster head is responsible for maintenance of synchronism and channel assignments within the cluster, but does not mediate most calls. It is also responsible for dissemination of networking information such as routing, and thus the designation of nodes as gateways to other clusters, as well as determination of cluster membership in cooperation with the nodes and the neighboring cluster.

Due to terrain and the motion of the network, certain links will be required to carry a disproportionate share of the long-distance traffic, leading to temporary bottlenecks. These bottlenecks may come and go more quickly than the network flow control can react. In multi-hopping networks it is thus very important for long messages (e.g. video conversations) to be prioritized in such a way that it is easy for nodes to drop the less important packets, without first consulting the source. For example, we may consider an embedded speech coding, channel coding, and packetization scheme proposed in [79] which is well suited to a multi-hop network with decentralized control. The speech coder arranges the bits in order of perceptual significance, and the channel coding on blocks of bits at each priority level can be performed using a punctured code. For less reliable reception, we can simply drop some of the redundant symbols and still be capable of decoding with the same hardware. This new stream can be packetized again in order of priority, and any node in the multi-hop network facing congestion can decide to drop the lower priority packets. While the reconstructed speech will then be of degraded quality, it will be much better than if packets were randomly dropped.

Large scale multi-hop networks typically have very low efficiencies--a ratio of control traffic to data of 3:1 or higher. To improve upon this, careful attention must be paid to improving the physical links so that they are more reliable (including interference rejection techniques where possible), assigning control functions to the appropriate level in the network, efficient routing and network maintenance protocols, and careful and possibly adaptive selection of packet size.

IV.1.3 Wireless Sensor Networks

Advances in integrated circuit fabrication technology have made possible the integration of sensing, actuation, signal processing, and wireless communications in a single small package [80]. The small size and cost can be enabling for a very large number of applications: security, environmental monitoring, manufacturing process control, condition based maintenance, medical monitoring, and remote exploration. Sensors embedded in appliances and other devices may form one level of a home or office integrated information network. In effect, this technology amounts to a means to connect the physical world to communications/computer networks, and as such may further multiply the impact information technology has already had on many spheres of life. Deployment modes for the networks could range from careful hand placement to random dispersal. These various applications mandate quite different network access strategies.

Consider first a security application, where the objective is to determine whether or not various classes of intruders are present, and then to communicate this information to the relevant authorities. Suppose further that the nodes are to be powered off small batteries, so that peak power consumption and average power consumption are both strictly limited. The various sensors can be sampled at relatively low rates, and thus be on with high duty cycle without leading to a large energy drain. However, RF communications is always costly, since the relevant circuits must operate at high frequencies. Therefore an architecture where the nodes always communicate the raw data is infeasible. Likewise, given the limited signal processing resources in each node and the fact that nodes see only a small part of the big picture, the nodes will not be able in some circumstances to make a firm determination on the presence or absence of an intruder. A compromise solution which preserves reliability while conserving power at the expense of latency is for the nodes to form tentative decisions with fairly high false alarm probabilities. A message is passed up including some reliability estimate. Messages from multiple nodes are collected, and if in fusing them there is sufficient certainty then security personnel may be alerted. Alternatively, if the outcome is still unclear, the raw data can be requested and processed by special purpose nodes.

The above procedures require two way messaging. To conserve power it is also necessary for receivers to be on with low duty cycle, which mandates a synchronous network. The arbitrary node placement suggests the need for multi-hopping; indeed, this is also required by the limits on peak power consumption, which preclude long distance communications by ordinary nodes. The data rates from individual nodes are quite low, and a few seconds of latency will not degrade the functionality of the network. Thus, such networks operate in a regime where the traditional constraints on latency and bandwidth are very weak, but energy consumption is of paramount importance. In contrast to our discussion in section III, the only purposes of dynamic resource management in terms of routing and channel assignments is to conserve energy, and preserve connectivity when nodes fail. Time division with large spatial re-use factors is adequate for the multiple access. Diversity techniques will be helpful in lowering peak power; as the nodes are too small for multiple antennas, the options are frequency hopping and channel coding. Frequency hopping permits the additional feature of avoidance of fixed external interferers if we can choose to be silent on particular hops, and also allows efficient avoidance of faded frequency slots for what will be an essentially static network.

The set of networking trade-offs is quite different in an automated manufacturing application. Here the network must make rapid decisions, and there are near ubiquitous power outlets. Therefore, a higher duty cycle of communications is permissible. Wireless communications is still desirable for a number of reasons: the expense of laying new communications lines, the difficulty of instrumenting rotating machine parts or retrofitting sensors to existing machinery, and communication with RF tags or sensors on the parts being fabricated. Parts of the network may be in motion, while other components are fixed. The electromagnetic environment may be unfavorable, with unintended emissions from the factory machines, so that longer range communications may well be by wire. For the wireless part of the network, the frequency and predictability of the communications (and interference) may be such that the techniques of section III become more attractive.

IV.2 Influence of Traffic Type on Access Strategy

The traffic type has a very large impact on the type of network that is designed. As noted with regards to sensor networks, even if all the traffic consists of data, the latency requirements and predictability of traffic bursts may vary widely. Multimedia traffic imposes even greater challenges [77,78,81,82], which we may only briefly outline here. Among the further complications are different tolerances for delay, error probability, and information rate variability on the part of the different applications. This task differs from traditional network designs in that the physical layer intrudes onto several layers of the network hierarchy. For example, if the interference conditions determine that at most some particular data rate is possible over a certain link, then the video coder and channel coding scheme must in tandem decide on new rates that give the best perceptual quality to the end user. But this decision on the maximum rate is made only after considering the different options for routing and bandwidth assignment, which in turn is the result of decisions made at many nodes in the network based on local observations. Thus, not only is there coupling between adaptive techniques at the physical layer, but also coupling through to the application layer.

Most of the interference suppression, avoidance, and averaging methods of section III rely upon the desired signals and interference having relatively long hold times compared to the frame duration, and furthermore implicitly assume all users occupy similar bandwidths. Successful multimedia networks must deal with both of these assumptions being violated. Consider a system with mixed voice and data traffic, where the radio knows what kind of traffic it will convey. Assuming we do not exploit the voice activity factor, the voice calls are characterized by a low but constant data rate over a long hold time, with moderate probability of error requirements but sharp limits on latency. The data traffic can be very bursty, with most messages typically being isolated packets (e.g. acknowledgments) but others (such as bulk file transfers) demanding high data rates. Latency requirements are loose but probability of error requirements are tight.

For short packets, significant gain can be had through statistical multiplexing. A number of random access techniques have been developed for wireless networks [e.g. 82-84]. Wireless networks differ in the spatial locality of packet collisions from wireline systems. Not all collisions are fatal, and the effects of collisions can be asymmetric, while listening to see if the channel is busy before sending is of less use since it may well be busy at the intended destination even if apparently available at the transmitter.

While interference avoidance is clearly not an option for bursty traffic, interference averaging and limited power control are. As part of the ordinary network maintenance, users can easily learn the propagation losses for their links. Then when data needs to be transmitted the power level can be set appropriately, limiting the interference caused to other users. If a packet is in error, some combination of more aggressive channel coding and usage of power can be used during retransmission as the latency limit is approached. Additionally, averaging techniques can be used to provide resistance to multipath and variable interference. This could for example be a combination of frequency hopping and channel coding, or simply wideband direct sequence CDMA.

How to mix connection and bursty traffic in multimedia wireless networks is an interesting research question. Connection traffic can be more efficiently transported if the interference mitigation techniques of section III are applied, but these will either fail or be less effective with bursty interference. This argues for an orthogonal separation of the traffic types across cells or clusters. On the other hand, this creates inefficiency in spectrum allocation since the traffic mix is dynamic, and variable across cells. This presents a complicated resource management problem, but mixing the traffic presents increased difficulty in meeting the divergent quality of service requirements. The assumptions on the mix of traffic, packet length, speech and video coding algorithms, and available interference mitigation techniques will all influence the design of the access scheme.

Matters become even more muddied when one considers applications such as an integrated wireless home or office information network, in which the devices may range from low power sensor nodes to work stations capable of multi-megabit transfer rates. It is clearly not feasible for sensor nodes to use the same protocol as the work stations, but communications must be coordinated in some fashion or the sensor node traffic will be lost in the higher power transmissions of other users, and work stations must be able to talk with the nodes. The general problem of designing networks with heterogeneous power levels, adaptive capabilities, and data rate requirements is very challenging, especially in the context of this network being connected to a larger wired network such as the Internet.

V. Conclusion

The wireless multiple access channel is characterized by variability in space, time, frequency, and interference. The design of an efficient communications system is largely an exercise in the management of this variability. It is our ability to estimate the state of the channel (including interference) which most fundamentally limits network capacity. The channel and interference can be engineered to a considerable extent through resource re-use restrictions, separation of users according to locations or traffic types, and the use of adaptive transmission and reception techniques. In this article we have moved from physical models, to point to point communications, to cellular multiple access schemes, and finally to networks with mixed traffic types, throughout emphasizing how the channel physically couples the users, and thus all algorithms and protocols that are built to tame this medium. The complicated nature of the channel and the way it impinges on all network layers up to and including applications makes wireless design very challenging and exciting. Careful modeling, simulation, and extensive testing with highly adaptive and adjustable radios are all required to better understand these interactions, and produce the designs that will realize the promise of anywhere, anytime, reliable communications.

Acknowledgements

Support for the research of the author and his students has been supplied by DARPA in contract JFBI 94-222/J4C942220, a Motorola University Partnership in Research contract, and a student fellowship from Rockwell.

References

[1] M.A. Jensen and Y. Rahmat-Samii, “EM interaction of handset antennas and a human in personal communications,” Proc. IEEE, vol. 83, no. 1, Jan. 1994, pp. 7-17.

[2] M.A. Jensen and Y. Rahmat-Samii, “Performance analysis of antennas for hand-held transceivers using FDTD,” IEEE Trans. Antennas Propagat., vol. 42, no. 8, Aug. 1994, pp. 1105-1113.

[3] D. Halliday and R. Resnick, Physics. 3rd Ed. Wiley: New York, 1978.

[4] H. Hashemi, “The indoor radio propagation channel,” Proc. IEEE, vol. 81, pp. 943-967, July, 1993.

[5] D. Parsons, The Mobile Radio Propagation Channel. Wiley: New York, 1992.

[6] J.G. Proakis. Digital Communications. McGraw Hill: New York, 1989.

[7] W.D. Rummler, “A new selective fading model: application to propagation data,” Bell Syst. Tech. J., vol. 50, May-June 1979, pp. 1037-1071.

[8] K.S. Gilhousen et al., “On the Capacity of a Cellular CDMA System,” IEEE Trans. Vehic.

Tech., vol. 40, pp. 303-312, May 1991.

[9] W.C. Jakes, Microwave Mobile Communications. Wiley: New York, 1974.

[10] J.G. Markoulidakis, G.L. Lyberopoulos, D.F. Tsirkas, and E.D. Sykas, “Mobility modeling in third-generation mobile telecommunication networks,” IEEE Personal Comm. Mag., vol. 4, Aug. 1997, pp. 41-56.

[11] H.L. Van Trees. Detection, Estimation, and Modulation Theory. Wiley: New York, 1968.

[12] S. Haykin. Adaptive Filer Theory. Prentice-Hall: Englewood Cliffs NJ, 1991.

[13] N.W.K. Lo, D.D. Falconer, and A.U.H. Shiekh, “Adaptive equalization and diversity combining for mobile radio using interpolated channel estimates,” IEEE Trans. Vehic. Tech., vol. 40, Aug. 1991, pp. 636-645.

[14] H-N. Lee and G.J. Pottie, “Fast adaptive equalization/diversity combining for time-varying dispersive channels,” to appear, IEEE Trans. Comm.

[15] V. Lin, Channel Coding and Power Control for FH/CDMA Radios. Ph.D. Dissertation, Electrical Engineering Dept., University of California, Los Angeles, 1995.

[16] R. Wesel and J.M. Cioffi, “Fundamentals of broadcast for OFDM,” Proc. 29th Asilomar Conf. on Signals, Systems, and Computers, Pacific Grove CA, 1995.

[17] E. Biglieri, D. Divsalar, P.J. McLane, and M.K. Simon, Introduction to Trellis-Coded Modulation with Applications. Macmillan: New York, 1991.

[18] G.R. Cooper and R.W. Nettleton, “A spread-spectrum technique for high-capacity mobile communications,” IEEE Trans. Vehic. Tech., vol. 27, pp. 264-275, Nov. 1978.

[19] P. Balaban and J. Salz, “Dual diversity combining and equalization in digital cellular mobile radio,” IEEE Trans. Vehic. Tech., vol. 40, pp. 342-354, May 1991.

[20] E. Perahia and G.J. Pottie, “Adaptive antenna arrays and equalization for indoor digital radio,” Proc. ICC ‘96.

[21] C. Wheatley, “Distributed antennas and microcells for CDMA,” CDMA Tech. Forum, San Diego, Feb. 28-Mar. 1, 1994.

[22] H.H. Xia, A.B. Herrera, S. Kim and F.S. Rico, “A CDMA-distributed antenna system for in-building personal communication services,” IEEE J. Select. Areas in Comm., vol. 14, May 1996, pp. 644-50.

[23] N. Seshadri, C-E.W. Sundberg, and V. Weerackody, “Advanced techniques for modulation, error correction, channel equalization, and diversity,” AT&T Tech. J., vol l72, No. 4, pp. 48-63, July/Aug 1993.

[24] A.F. Naguib, V. Tarokh, N. Seshadri and A.R. Calderbank, “Space-time coded modulation for high data rate wireless communications,” submitted IEEE J. Select. Areas in Comm, Sept. 1997.

[25] V. Tarokh, N. Seshadri, and A.R. Calderbank, “Space-time codes for high data rate wireless

communication: performance criterion and code construction" AT&T Research manuscript.

[26] T.M. Cover and J.A. Thomas. Elements of Information Theory. Wiley: New York, 1991.

[27] A. Goldsmith and P. Varaiya, "Increased spectral efficiency through power control," Int'l Conf. on Comm., Geneva, May 1993.

[28] J.A.C. Bingham, "Multicarrier modulation for data transmission: an idea whose time has come," IEEE Comm. Magazine, pp. 5-14, May 1990.

[29] J. Hagenauer, "Rate-compatible punctured codes (RCPC codes) and their applications," IEEE Trans. Comm., vol 36, pp. 389-400, April 1988.

[30] A. Afrashteh, N.R. Sollenberger, J.C.-I. Chuang, and D. Chukurov, "Performance of a TDM/TDMA portable radio link for interference, noise, and delay spread impairments," IEEE Trans. Vehic. Tech., vol 43, pp. 1-7, Feb. 1994.

[31] I. Katzela and M. Naghshineh, "Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey," IEEE Personal Comm. Mag., vol. 3, June 1996, pp. 10-31

[32] K.N. Sivarajan, R.S. McEliece, and J.W. Jackson, "Dynamic Channel Assignment in Cellular Radio," IEEE Vechic. Tech. Conf., May 6-9, 1990, Orlando FL, pp. 631-637.

[33] L.J. Cimini and G.J. Foschini, "Distributed algorithms for dynamic channel allocations in microcellular systems," 42nd IEEE Vehic. Tech. Conf., pp. 641-644, 1992.

[34] M. M-L. Cheng and J. C-I. Chuang, "Performance evaluation of distributed measurement-based dynamic channel assignment in local wireless communications," IEEE J. Select. Areas in Comm., vol 14, May 1996, pp. 698-710.

[35] J. Zander, "Distributed cochannel interference control in cellular radio systems," IEEE Trans. Vehic. Tech., vol. 41, pp. 305-311, Aug. 1992

[36] S.A. Grandhi, R. Vijayan, and D.J. Goodman, "A distributed algorithm for power control in cellular radio systems," Proc. Allerton Conf. on Comm., Control, and Computing, 1992.

[37] G.J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," IEEE Trans. Vehic. Tech., vol. 42, pp. 641-646, Nov. 1993.

[38] N. Bambos and G.J. Pottie, "Power control based admission policies in cellular radio networks", IEEE Globecom, Dec. 6-9, 1992, Orlando FL, pp. 863-867.

[39] S. Chen, N. Bambos, and G.J. Pottie, "Radio link admission algorithms for wireless networks with power control and active link protection," technical report UCLA-ENG-94-25

[40] V. Lin and G.J. Pottie, "Implementation of distributed power and access control for a frequency hopped wireless transceiver," Proc. Internat. Conf. on Comm, 1995.

[41] C.J. Hansen, C.C. Wang, and G.J. Pottie, "Distributed dynamic channel resource allocation in wireless communication systems," 28th. Asilomar Conf. on Signals, Systems, and

Computers, Oct. 1994 pp. 78-82.

[42] R.D. Yates, “A framework for uplink power control in cellular radio systems,” to appear, IEEE J. Select. Areas in Comm., 1995

[43] S.A. Grandhi, R.D. Yates, and D.J. Goodman, “Resource allocation for cellular radio systems,” IEEE Trans. Vehic. Tech., vol. 46, Aug. 1997, pp. 581-587.

[44] C.J. Hansen and G.J. Pottie, “Distributed access control in wireless and wireline systems,” Proc. IEEE Int. Symp. Info. Theory, 1995.

[45] G.J. Pottie and A.R. Calderbank, “Channel coding strategies for cellular radio,” IEEE Trans. Vehic. Tech., vol. 44, Nov. 1995, pp. 763-770.

[46] C.C. Wang and G.J. Pottie, “Interference avoidance and power control strategies for coded frequency hopped cellular systems,” Proc. Internat. Conf. on Comm., 1995.

[47] C.C. Wang and G.J. Pottie, “Bit allocation algorithms for FH-CDMA wireless communication systems,” 34th Allerton Conf. on Comm., Control, and Computing, 1996.

[48] J.H. Winters and M.J. Gans, “The range increase of adaptive vs. phased arrays in mobile radio systems,” 28th Asilomar Conf. on Signals and Systems, Oct. 1994, pp. 109-115.

[49] J.H. Winters, J. Salz, and R.D. Gitlin, “The impact of antenna diversity on the capacity of wireless communication systems,” IEEE Trans. Comm., vol. 42, pp. 1740-1751, April 1994.

[50] J.H. Winters, “Signal acquisition and tracking with adaptive arrays in the digital mobile radio system IS-54 with flat fading,” IEEE Trans. Vehic. Tech., Nov. 1993.

[51] D. Gerlach and A. Paulraj, “Adaptive transmitting antenna methods for multipath environments,” Proc. Globecom 1994.

[52] Y. Li, M.J. Feuerstein, and D.O. Reudink, “Performance evaluation of a cellular base station multibeam antenna,” IEEE Trans. Vehic. Tech., vol. 46, Feb. 1997, pp. 1-9.

[53] N. Seshadri, A.R. Calderbank, and G.J. Pottie, “Signal design for co-channel interference suppression with applications to wireless communications,” Proc. Vehic. Tech. Conf. 1995.

[54] S. Verdu, “Minimum probability of error for asynchronous Gaussian multiple access channels,” IEEE Trans. Inform. Theory, vol. IT-32, pp. 85-96, Jan. 1986.

[55] R. Lupus and S. Verdu, “Near-far resistance of multiuser detectors in asynchronous channels,” IEEE Trans. Comm., vol. 38, pp. 496-508, April 1990.

[56] S. Verdu, “Multiuser Detection,” Advances in Statistical Signal Processing, Vol 2, JAI Press, 1993.

[57] A. Duel-Hallen, “Decorrelating decision-feedback multi-user detector for synchronous code-division multiple-access channel,” IEEE Trans. Comm., vol. 41, pp. 285-290, Feb. 1993.

[58] U. Madhow and M.L. Honig, “MMSE interference suppression for direct-sequence

spread-spectrum CDMA," IEEE Trans. Comm., vol. 42, pp. 3178-3188, Dec. 1994.

[59] R.A. Iltis and L. Mailaender, "Linear multiuser detectors for quasi-synchronous CDMA systems," 28th Asilomar Conf. on Signals, Systems and Computers, Oct. 1994, pp. 104-108.

[60] S.E. Bensley and B. Aazhang, "Subspace-based channel estimation for code division multiple access communication systems," IEEE Trans. Comm., vol. 44, Aug. 1996, pp. 1009-1020.

[61] S. Moshavi, "Multi-user detection for DS-CDMA communications," IEEE Comm. Mag., Oct. 1996, pp. 124-136

[62] G. Caire, G. Taricco, J. Ventura-Traveset, and E. Biglieri, "A multi-user approach to narrowband cellular communications", IEEE Trans. Inform. Theory, vol. 43, Sept. 1997, pp. 1503-1517.

[63] Eldad Perahia. Diversity Combining, Adaptive Antennas, and Equalization for Digital Radio. Ph.D. Dissertation, Electrical Engineering Dept., University of California, Los Angeles, 1995.

[64] Christopher J. Hansen. Probing Techniques for Multiuser Channels with Power Control. Ph.D. Dissertation, Electrical Engineering Dept., University of California, Los Angeles, 1997.

[65] R.L. Peterson, R.E. Ziemer, and D.E. Borth. Introduction to Spread Spectrum Communicaitons. Prentice-Hall: Upper Saddle River, NJ, 1995.

[66] A.J. Viterbi, "The orthogonal-random waveform dichotomy for digital mobile personal communications," IEEE Personal Comm. Magazine, vol. 1 No. 1, 1994, pp. 18-24.

[67] K.B. Lee, "Orthogonalization based adaptive interference suppresion for direct-sequence code division multiple access systems," IEEE Trans. Comm., vol. 44, Sept. 1996, pp. 1082-1085

[68] A.F. Naguib, A. Paulraj, and T. Kailath, "Capacity improvement with base-station antenna arrays in cellular CDMA," IEEE Trans. Vehic. Tech., vol. 43, Aug. 1994, pp. 691-697.

[69] J.S. Thompson, P.M. Grant, and B. Mulgrew, "Smart antenna arrays for CDMA systems," IEEE Personal Comm. Mag., vol. 3, Oct. 1996, pp. 16-25.

[70] M.V. Clark, V. Erceg, and L.J. Greenstein, "Reuse efficiency in urban microcellular networks," IEEE Trans. Vehic. Tech., vol. 46, May 1997, pp. 279-288.

[71] S. Dehghan and R. Steele, "Small cell city," IEEE Comm. Mag., Aug. 1997, pp. 52-59.

[72] L.J. Greenstein et al., "Microcells in personal communications", IEEE Comm. Magazine, Dec. 1992, pp. 76-88.

[73] K.J. Kerpez, "A radio access system with distributed antennas," IEEE Trans. Vehic. Tech., vol. 45, May 1996, pp. 265-275.

[74] A.J. Viterbi, A.M. Viterbi, K.S. Gilhousen, and E. Zehavi, "Soft handoff extends CDMA cell coverage and increases reverse link capacity," IEEE J. Select. Areas in Comm., vol. 12,

Oct. 1994, pp. 1281-1288.

[75] S.V. Hanly. Information Capacity of Radio Networks. Ph.D. Dissertation, King's College, University of Cambridge, Aug. 1993.

[76] A. Goldsmith, "Multi-user capacity of time-varying channels," 28th Asilomar Conf. on Signals, Systems and Computers, Oct. 1994, pp. 83-88.

[77] M. Gerla and J.T.C. Tsai, "A distributed mobile wireless infrastructure for multimedia applications," Fifth WINLAB Workshop on Third Gerneration Wireless Networks, April 1995.

[78] A. Alwan et al., "Adaptive mobile multimedia networks," IEEE Personal Comm. Mag., April 1996, pp. 34-51.

[79] Benjamim Tang. Adaptive Wireless Voice Communications with Embedded Source and Channel Coding. Ph.D. Dissertation, Electrical Engineering Dept., University of California, Los Angeles, 1995.

[80] K.Bult et al., "Wireless integrated microsensors," Proc. 1996 Hilton Head Transducers Conference.

[81] R. Jain, J. Short, L. Kleinrock, and J. Villasenor, "PC-notebook based mobile computing: Algorithms, architectures, and implementation," to appear, ICC 95.

[82] P. Agrawal et al., "SWAN: A mobiel multimedia wireless network," IEEE Personal Comm. Mag., April 1996, pp. 18-33.

[83] F. Borgonovo, L. Fratta, M. Zorzi, and A. Acampora, "Capture division packet access: a new cellular architecture for future PCNs", IEEE Comm. Mag., Sept. 1996, pp. 154-162.

[84] S. Nanda, D.J. Goodman, and U. Timor, "Performance of PRMA: A packet voice protocol for cellular systems," IEEE Trans. Vehic. Tech., vol. 40, Aug 1991, pp. 584-598.

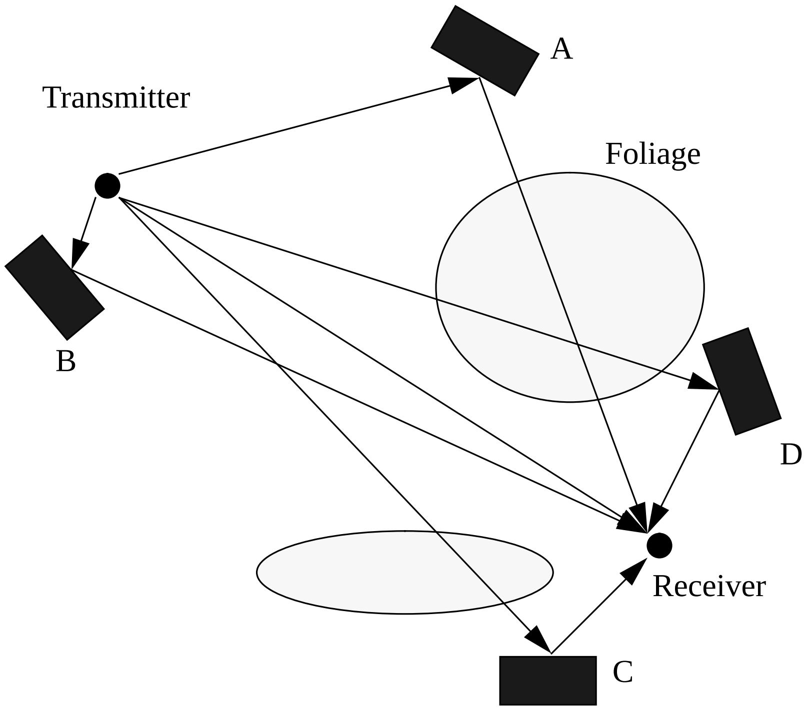

Figure 1: Multipath radio propagation

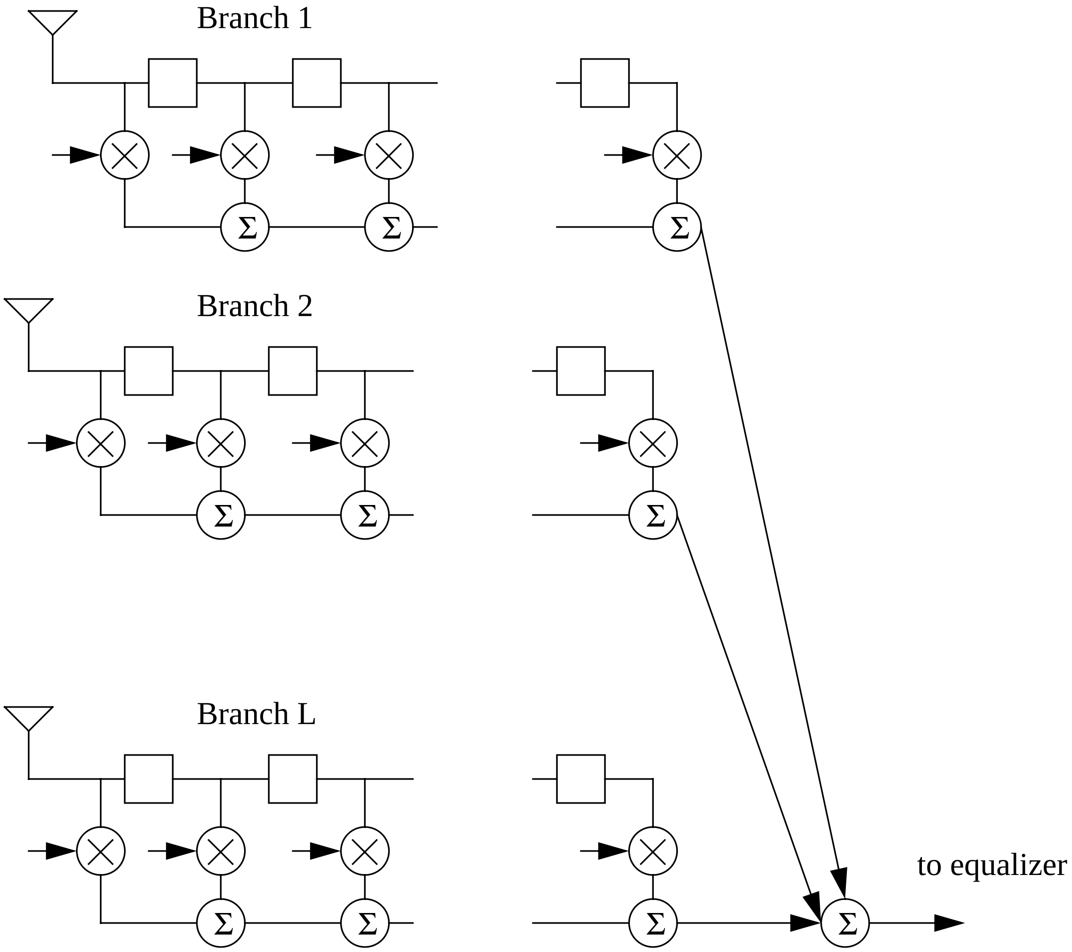

Figure 2: Multi-channel diversity combiner

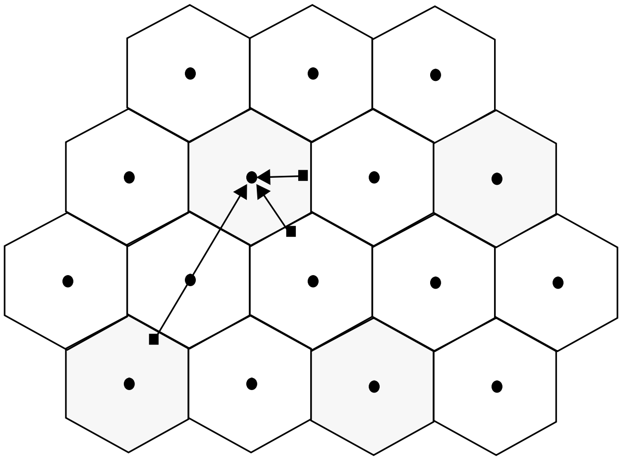

Figure 3: Reduction of worst case interference with frequency re-use K=4.

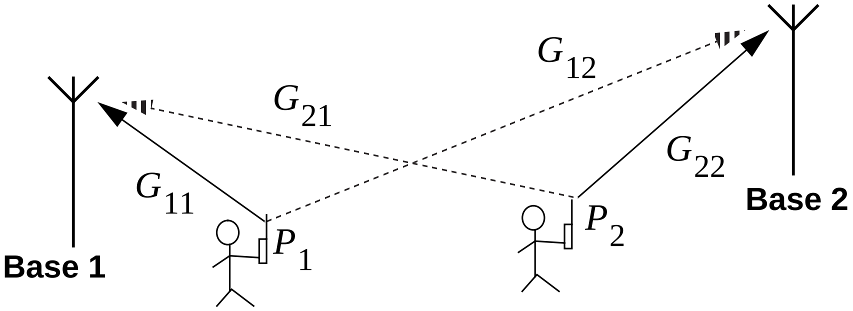

Figure 4: Propagation scenario between mobiles and base stations.

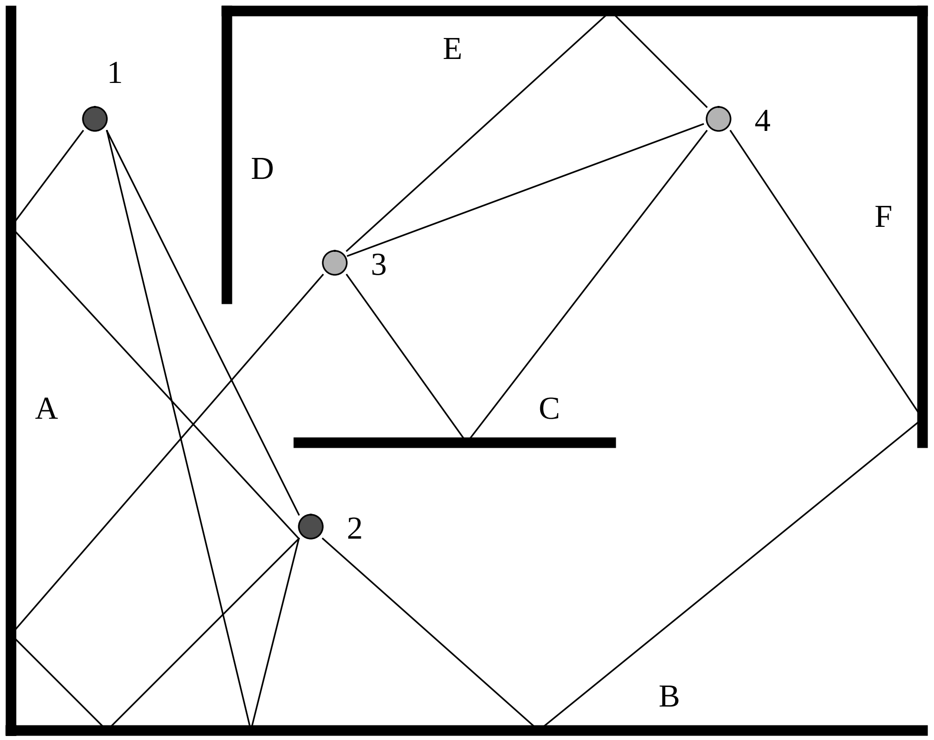

Figure 5: The interference-coupled beamforming problem

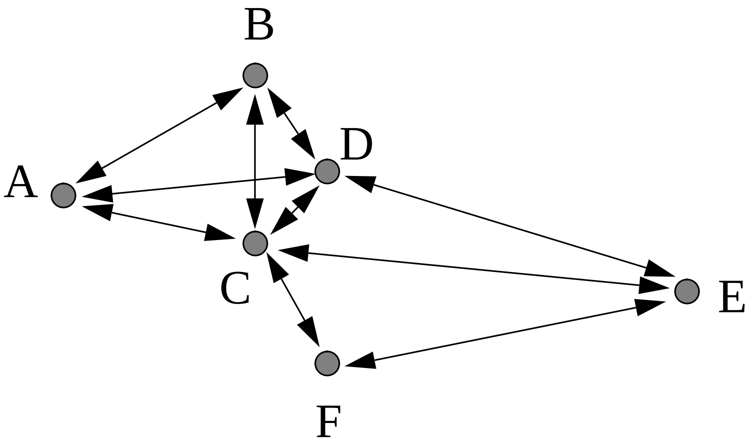

Figure 6: Feasible connections in multi-hop network

Glossary of Acronyms and Abbreviations

CDMA	Code Division Multiple Access: a spread spectrum multiple access technique in which each user has a different spreading code, with low cross-correlations among codes.
CT-2	Cordless Telecommunications 2: a UK second generation cordless telephone standard.
dB	Decibels: $10\log_{10}$(power)
DCA	Dynamic Channel Allocation: re-use of channels is accomplished according to the actual traffic, rather than according to static rules.
DECT	Digital European Cordless Telecommunications: a time division multiple access standard, employing Gaussian minimum shift keying modulation
DFE	Decision Feedback Equalizer: a device which includes a linear filter, and either or both of a noise predictive or linear filter in a feedback section whose inputs are determined by a zero delay decision device.
DMT	Discrete Multi-tone modulation: a method for sending information over multiple subcarriers, making use of the orthogonal properties of the inverse Fourier transform
DPCA	Dynamic Power and Channel Allocation: re-use of channels and the assignment of transmit powers are controlled by an adaptive algorithm, rather than being fixed a priori.
DS	Direct Sequence (spread spectrum): direct modulation of the message bearing waveform by a pseudo-random sequence (spreading code).
Eb/No	Energy per bit divided by received noise energy per bit: used in comparing modulations and codes as a measure of their efficiency in achieving particular bit error rate targets.
FDMA	Frequency Division Multiple Access: orthogonal decomposition of the band of available frequencies among the users.
GSM	Groupe Speciale Mobile: a European time division multiple access standard for cellular radio, employing minimum shift keying modulation.
IS-95	Interim Standard 95: a (US) Telecommunications Industry Assocation standard for direct sequence code division multiple access cellular radio, to operate in the same frequency band as current North American cellular systems.
LIA	Least Interference Algorithm: a dynamic channel assignment principle, based upon causing the least multiple access interference to other users.
LMS	Least Mean Squares: an iterative gradient descent algorithm which minimizes the squared error at the decision point of the receiver
MC	Multichannel Combiner: a device with adaptive linear filters on each of several antenna branches
MLSE	Maximum Likelihood Sequence Estimation: in the context of equalization or decoding, a device which finds the most likely symbol sequence to have been transmitted, typically using the Viterbi Algorithm
MMSE	Minimum Mean Square Error: a criterion of optimality for an equalizer or other adaptive device, as opposed to zero forcing (elimination of all intersymbol interference).
OFDM	Orthogonal Frequency Division Multiplexing: see DMT
PN	Pseudo Noise (sequence): a sequence designed to appear statistically similar to a random sequence, typically binary and periodic
RF	Radio Frequency: either refers to the fact that a communications system uses radio waves, or to the section of the transceiver that operates at radio frequencies (e.g., the front end).

RLS	Recursive Least Squares: a class of iterative algorithms that achieve rapid convergence to the minimum mean square error solution (e.g., for an adaptive equalizer or antenna array).
SIR	Signal to Interference Ratio: signal power divided by the total of interference and noise power, measured at the decision point of the receiver.
SNR	Signal to Noise Ratio: signal power divided by the noise power, measured at the decision point of the receiver
TDMA	Time Division Multiple Access: multiple access through decomposition of each frame into distinct time slots allocated to each user; in cellular systems, TDMA usually implies a hybrid structure where FDMA channels are divided up among multiple users.