Fair Resource Allocation in One-Dimensional Cellular Systems

Gregory J. Pottie, *Fellow, IEEE*

***Abstract*—We consider the problem of fair resource allocation in interference channels, for the case of one-dimensional cells. We propose simple interference channel models based upon the deterministic channel model, assuming zero queuing delay. This enables study of how the interference induces coupling among the rates available in the channels and evaluation of resource allocation. A sequence of models of increasing complexity provides a simple means of dissecting the different effects, and rapid simulation. We show that there is a strong benefit in explicitly pairing channels to minimize interference to users with disadvantaged locations, as compared to an averaging approach such as CDMA. We also show the benefits of power and bandwidth allocation when fair allocations are desired. A distributed algorithm is proposed that embodies the principle of interference channel pairing. We discuss the consequences for cognitive and multi-band radios.**

Index Terms—**Wireless networks, models**

I. Introduction

While there are many results on allocation of resources in multiple access wireless systems, the best schemes remain elusive in the general cellular setting. There are numerous difficulties. First, there is the question of the traffic model and the appropriate quality of service requirements. There are many application dependent choices. Second, there are many choices of propagation models, both in terms of the loss and the dynamics. Third, there are also many choices in how interference is to be modeled, including whether it is to be treated as noise or something that can under certain conditions be cancelled. The model parameter space is thus vast, with the choices dictating very different algorithmic solutions and complexities for simulation. Our objective is to design a set of models that are simple enough to yield analysis, approximations, and low complexity simulations, but yet which have enough detail to force consideration of the main design tradeoffs that are encountered in practice. Practical algorithms for performing the channel, power and bandwidth allocations follow the construction of these network abstractions. We now review some of the prior art for resource allocation before describing our general approach.

For single cell non-fading multiple access channels, fair orthogonal frequency assignments are capacity achieving [1]. However, in multi-cellular systems, information theory has historically been less successful. Recently, there has been considerable information theoretic work in channel allocations to exploit the principle of interference alignment [2,3]. Roughly, once a channel has experienced severe interference, it is better to pile up with more interference to enable other channels to have better SIR. In the more sophisticated techniques, codebooks are shared. We will show how an uncoded version of alignment, which we term coordination, can be very effective in managing interference in multi-cell systems with low complexity. Additionally, with the advent of the deterministic channel model new capacity results have been proven [4]. In this model, the effect of interference is to reduce the number of bits that can be transmitted on a link, with noise and unaccounted for interference (e.g., from distant cells) playing the role of a limit on the maximum number of bits that can be sent. Together with the assumption of transmission at Shannon capacity, it also leads to vastly simplified simulations, since probability of error calculations get abstracted away.

The resources that are allocated are power, bandwidth, and channels. The optimal allocation of the last is NP hard, while the others are much lower complexity, since algorithms that are effectively variants on gradient searches can be constructed. For example, dynamic power allocation is often accomplished by comparing error rate or SIR to some threshold, and then toggling power up or down e.g. [5,6]. There are many algorithms that perform some combination of channel and/or power allocation e.g. [7-13]. The two basic categories of practical channel allocation algorithms are interference averaging schemes (e.g., CDMA and OFDMA) and interference avoidance (e.g., the least interference algorithm [12]). Our focus will be on avoidance, although we will compare to averaging. In particular, we are interested in how the interference alignment principle can be used to improve the channel allocations, given a power and bandwidth allocation strategy.

The traffic arrival model and the quality of service requirements also significantly impact capacity. With a Poisson arrival model, statistical multiplexing is beneficial, whereas with self-similar traffic it has limited benefits. Utility may increase with rate, or fall off a cliff with too much delay. A considerable literature exists on game theoretic and other approaches for dealing with both homogeneous and mixed traffic types [e.g. 13], including issues such as prediction of future traffic state. We will focus mainly on interactive media traffic (e.g., voice or video calls) since it can be simply modeled and is a difficult class

This work was supported in part by Broadcom Corporation under a UC MICRO grant, the Australian-American Fulbright Commission, and by an Erskine Fellowship from the University of Canterbury, New Zealand.

G.J. Pottie is with the UCLA EE Department, Box 951594 Los Angeles CA 90095-1594 USA (email: pottie@ee.ucla.edu)

of traffic to manage due to the stringent delay constraints.

The propagation model has a large impact on the possible capacity. With a large number of independently fading channels, MIMO can improve the capacity of individual links [14,15]. The techniques we propose are compatible with MIMO systems. Under certain conditions, interference cancellation can also improve the capacity; however, consideration of such techniques is beyond the scope of this paper.

Cognitive radio is a term with many interpretations [16-19]. Given that smartphones now have multi-band and multi-protocol capabilities (and multiple radios), we will take these features as baseline functionality. We will make varying assumptions regarding the analog front end of the radio. Current technology limits low-cost radios to support digital modulation formats of 6 bits per symbol or less; based on technological trends, we will suppose that up to 10 bits is plausible in the medium term. Many radios currently deployed only transmit at the maximum power (particularly in WiFi). We will accordingly compare use of power control to having a fixed level, since there is always a cost in demanding a large dynamic range.

Our approach is to exploit both recent and classical information theory abstractions to simplify modeling of the physical layer, moving from very simple to more complex models to develop insights into particular resource allocation design choices. We begin with the deterministic channel model. Transmissions are error-free and at the Shannon capacity. It is the simplest model that captures the damaging effects of interference. In the same spirit, we choose as the first propagation model fourth power loss with distance. This is the simplest model that captures the location disadvantages of some users with respect to others, and the variable interference effects caused to other cells. (Fading will be considered later). We also initially limit the users to two or four symmetrical locations in each cell. This enables separating out different effects that are important for multiple access algorithm design, such as power allocation and how to coordinate channel allocations in response to loads in nearby cells. For the traffic model, we compare two choices. The Erlang model [20] is representative of the quality of service constraints of interactive media traffic: as high a rate as possible with low blocking probability and small delays relative to message length. It implicitly demands strict fairness in rates. While other utility functions are possible, in most instances there is some minimum acceptable rate. At high loading, with sublinear utility in rate above the minimum, solutions will tend towards fair rate allocation. Moreover, using the sum rate can result in extreme geographic unfairness in allocation that would typically be unacceptable. We will also consider infinite queuing as a limiting case (corresponding to a large number of channels) that is more amenable to quick analysis. This is useful in establishing limits that might be attained with distributed algorithms, and reasonably models transmission of data. Finally, we begin with one-dimensional cellular networks, as might be encountered along interstate highways.

With these simple models, we have been able to demonstrate a number of important results. First, it is easily shown that power control is a suboptimal strategy for achieving strict fairness in CDMA; bandwidth allocation is considerably superior in the uplink, and a combined bandwidth and power allocation scheme is better still. Moreover, allocation strategies that pair to the interference (including channel vacancies) from surrounding cells and then match low interference conditions against high loss users do better than CDMA. In such coordination schemes, generally a combination of power and bandwidth allocation is required to achieve fairness. The cost of not doing power control is in some important conditions quite small, if bandwidth allocation is performed. In others, the price is very steep, and fair allocations very difficult to achieve without some power control. We have produced a distributed algorithm that performs the pairing and bandwidth allocation, along with power control based upon propagation loss, that outperforms OFDMA. Current wide-band systems (e.g., WiFi) access a limited number of channels. Our results indicate that there is a large potential gain if more channels can be accessed across multiple bands (e.g., as in cognitive radios), as compared to allocating these bands separately. At small numbers of channels, this effect dwarfs the differences among well-designed multiple access schemes.

The remainder of the paper is organized as follows. In section II, we discuss the nature of the interference channel, and present simple examples that illustrate the tradeoffs and provide some limiting results. In section III we present our simulation results for two and four point cells, and the proposed distributed algorithm. We apply the algorithm to continuous location distributions, illustrating the benefits of coordination and power control. We compare the alignment scheme and OFDMA. In section IV we explore the implications of adding static fading, and present comparisons for the coordination and OFDMA. In section V we present our conclusions and suggestions for future research.

II. TRAFFIC AND INTERFERENCE CHANNEL MODELS

We begin with standard queuing theory [20]. The probability of n active messages for c queues, normalized offered load ρ and zero delay is given by the Erlang loss formula,

$$p_n = \frac{\rho^n / n!}{\sum_{j=0}^{c} \rho^j / j!} \quad (1)$$

where $n=c$ for blocking. As the number of servers (channels) increases, we can get closer to full utilization of the server capacity with low blocking. Thus as bandwidth increases, for a given offered load, blocking probability decreases since

the number of channels available goes up. In this model, the delay is expressed in multiples of the average message duration. For multimedia streams, this means for purposes of queuing analysis that we have essentially zero delay—the Erlang loss model is thus appropriate. Similarly, since data can be queued many times the message duration, an infinite queuing model is appropriate for simplifying analysis, reflecting a utility function in which only rate must be optimized. As either the queue depth or the number of servers goes to infinity, the models converge. We will thus consider infinite *n* as the limiting case, and finite *n* for actual networks.

Interference greatly complicates analysis since server rates depend on occupancy in neighboring cells. In the deterministic channel model we must account only for the strongest such interferers. Noise limits the maximum number of bits that can be achieved per transmission, while interference lowers the number of bits to below that level. With this abstraction we can choose to include fading, distance loss, shadowing and the like, all of which get folded finally into a number indicating how many bits can be sent. A propagation model with distance loss alone is the simplest that still yields important results. First, including distance loss enables frequency re-use, which is universal in wireless systems. Second, it leads to issues of fairness: some transmissions within a cell are more advantaged than others with respect to resource requirements to achieve a given bit rate. This forces the use of very different allocation algorithms than if all channels were equal, and in particular, consideration of variable bandwidth and power allocation. We will consider two basic classes of symmetrical one-dimensional cells before applying the lessons learned to more complicated situations. In the first, there are only two transmission locations per cell, halfway to the cell boundary from the center. In the second, there are four equally spaced points. For two points per cell, there are no geographic fairness issues, and we can directly study interference coupling among cells. With four points per cell, both coupling and fairness issues must be considered jointly. We also consider two different arrival models. In the first, arrivals are spatially periodic with period 2. Thus both neighbors behave in the same way, resulting in a highly symmetrical problem that is amenable to analysis. In particular, it is straightforward to identify optimal fair allocation policies for the situations of either sparse or full loading. In the second, we assume there are four independent arrival processes. This already considerably complicates the problem, and due to distance loss considerations, will yield results comparable to consideration of an infinite cell network. Yet simulations can be limited to a small network, increasing speed. Having developed algorithms based on these simplified models, we later apply them to more realistic situations, including continuous location distributions.

The interference channel is a system of coupled queues, in which we have found that even approximate solutions tend to be numerically unstable. Thus, we must consider special cases and simulations. In the limit of few channels, most neighboring channels are vacant. A straightforward Erlang model assuming a single rate corresponding to all neighbors vacant (i.e., the noise limit) yields good results for symmetrical situations such as two points per cell and two repeating arrival processes. In the limit of a very large number of channels, essentially all channels will be occupied. There are however still multiple rates possible, since interference depends on which neighbor occupies the same channel, and loss in general depends upon position. This case is similar to infinite queuing. First consider the two location, two channel model. There are only two possible channel allocations for the two cells. Let the symbol | denote a cell boundary. We can have the channel allocations:

| 1 2 | 1 2 | 1 2 | … or | 1 2 | 2 1 | 1 2| 2 1|…

If the user positions within the cell are each unit distance from the base station, then in both cases for either channel there is one neighbor at distance 3 and one at distance 5 in the neighboring cells. The SIR is thus 71.7, and the number of bits at Shannon capacity is 6.2. The obvious allocation policy is for each position to have the same power and bandwidth, and to occupy every position in every cell. Leaving vacancies only decreases the total throughput. This produces the same result as CDMA (assuming orthogonal assignments within cells). This channel arrangement is highly symmetrical. For each channel it reduces to the classic binary interference channel.

In going to four independent cells, periodically repeating, in addition to the above cases we have situations where the right and left interfering cells can have reversed orders of channels, for example:

| 2 1 | 2 1 | 1 2 | 1 2 |…

If we consider the third cell, channel 1 has two close neighbors, achieving 5.4 bits, while channel 2 has two far neighbors, achieving 8.3 bits. Notice that every cell has two different neighbors, so this set of rates is achieved in every cell, albeit for different channels. On average this is more than in the two-cell case, but it is unfair. However, we could time share among a number of realizations so that every location gets the average rate of 6.8 bits. If the ordering of channels were random in each cell, then an average over all these realizations yields a rate of 6.5 bits. Achieving the higher rate requires coordination across the network, which is impractical. The random ordering is a more plausible target rate for distributed algorithms to aim for. The rate is still higher than for two cells, as there are both near and far neighbors, as opposed to only (effectively) near ones.

For four points per channel, we now have fairness issues immediately in the two-cell case. The best fixed arrangement is

|1 2 3 4 | 2 1 4 3 |…

In this case we are pairing outer points with inner points in the neighboring cell to partly equalize SIRs. To achieve fairness, the outer points can have transmission power 5.3 times that of the inner points, so that all points achieve 6.8 bits per transmission. Now consider four independent cells. Assume that the maximum number of bits per transmission is 10, and that the pairings of inner to outer points include the ensemble of possible combinations of near and far neighbors of the same type (i.e., inner or outer). With these

assumptions, channel coordination with power control leads to the highest throughput at 7.2 bits. We considered OFDMA with orthogonal assignments within cells and overlaps with other users in at most one subchannel in neighboring cells. This form of OFDMA with coordinated subchannels is the next best scheme at 6.9 bits, but requires a much larger spread in powers. If the subchannels in the two neighbors are not coordinated, a substantial penalty results. CDMA with power control is the worst of all, with only 4 bits. Among pure bandwidth allocation schemes, a time-sharing arrangement between inner and outer points does best at 5.7 bits. In this scheme, paired inner and outer points alternate frames with outer points only. Similar schemes for OFDMA and CDMA are less effective. For CDMA the best scheme is a hybrid of bandwidth allocation and power control, only slightly outperforming bandwidth allocation alone (4.9 vs. 4.7 bits).

When we now consider zero queuing, vacancies in channels break the optimal pairings, since the number of arrivals being processed at any given location is stochastic. For four points per cell, inner points become much more advantaged since they gain high bit rates whether they are paired with vacancies or other inner points, while outer points get very low bits rates when paired with other outer points, regardless of the power control strategy. Thus, in addition to power control and bandwidth allocation, we must also consider the priority with which the various classes of locations are coordinated with vacancies in nearby channels. The other effect of vacancies is that the number of possible patterns for channel assignments becomes very large. In practice, the allocation policy will thus need to be based on a local heuristic.

III. SIMULATIONS FOR NON-FADING CHANNELS

Simulations were completed for five cases involving two points within each cell, as shown in Table 1. Bit rates per channel per unit time are quoted for 1% blocking probability. We compare channel coordination where channels are randomly occupied or channels are coordinated such that channels are first occupied to align with vacancies in the neighboring cells. This was accomplished through an indexing scheme, where odd numbered cells fill from the bottom and even numbered cells fill from the top. While random assignment is beaten by coordination at lower numbers of channels, in the limit of large numbers it performs as well, since the vacancies disappear. The same procedures were repeated for four independent cells, with similar results but somewhat higher throughputs due to the mix of near and far neighbors. Finally, CDMA is examined. It performs competitively to random channel alignments up to moderate numbers of channels, with slightly lower throughput at larger numbers.

Channels	2 cell random	2 cell coord	4 cell random	4 cell coord	4 cell CDMA
4	1.9	2.1	1.7	2.0	1.8
8	3.0	3.5	2.9	3.3	3.0
16	4.0	4.6	3.9	4.4	3.9
32	4.8	5.2	4.7	5.1	4.6
64	5.3	5.6	5.4	5.6	5.2
128	5.6	5.9	5.8	5.9	5.5
256	5.9	6.0	6.0	6.1	5.8
512	6.0	6.1	6.3	6.3	5.9
1024	6.1	6.1	6.4	6.4	6.1
infinity	(6.2)	(6.2)	(6.5)	(6.5)	(6.2)

Table 1. Bits per channel per unit time for 2-Point Scenario

In extending the simulations to four points per cell, the main complication is that the service rates are different for the inner vs. outer points. At low loading, the most important mechanism for providing fairness is to preferentially align the outer points with vacancies in the neighboring cells (the focus with the two-point model). At high loading, we aim to align them with inner points in neighboring cells. We found that both power and bandwidth allocation were required for achieving fairness. A distributed algorithm was developed whereby low index channels are assigned in cells alternately among the low and high loss users. Further, also using the indexing, vacant channels in neighboring cells will tend to align, and be paired with high loss users. Rates were adjusted by comparing the average rate obtained for an arrival against the average for the system. If below the average rate by some threshold, more channels were assigned; if above by some threshold, no resources were assigned. Power was taken to be approximately the square root of the propagation loss. The benefits are lower dynamic range of power adjustment compared to traditional power control algorithms, but yet higher throughput than leaving all powers fixed.

With a larger number of locations per cell, it is difficult to jointly optimize pairing, power, and bandwidth in a situation where there is variable channel occupancy. However, some insight may be gained from the case of infinite numbers of channels. For power control laws of the form $|x|^q$, where $|x|$ is the distance from the basestation, we computed the best bandwidth allocation for a given value of q, and then stepped through values of q. For coordination we obtain 6.9 bits with q=2.5, and the range in bandwidth allocation, W, being a factor of 1.5. For OFDMA, the figures are 6.4 bits, q=2.5, W range of 2.1, while for CDMA, the results are 4.2 bits, q=2.0, and W range 5.0. Using these results, we then compared the three assignment strategies in simulation. We performed simulations for cells of size 32, with users randomly located within them. A 6-cell network was assumed, with interference included from the nearest and next-nearest neighbors on the right and left. This effectively simulates an infinite network. CDMA results approached the bound, but given the inferiority of the method are not reported below. Coordination and OFDMA

with random channel assignments are compared in Table 2 for the best value of q and for flat power.

n	Coord q=2.5	Coord q=0	OFDMA q=2.5	OFDMA q=0
4	1.9	1.9	1.7	1.5
8	3.2	3.0	2.8	2.5
16	4.3	3.9	3.8	3.3
32	5.2	4.5	4.4	3.8
64	5.6	4.8	4.9	4.1
128	6.0	4.8	5.2	4.3
256	6.3		5.4	4.5
512	6.4			
1024	6.5			
infinity	(6.9)		(6.4)	

Table 2: Bits per channel per unit time for distributed algorithm in 6-cell model

Increasingly inferior results are obtained when flat power is assigned, as the value of n increases. For OFDMA, performance for q=2.5 is only slightly superior to q=4—in simulations only about 0.1 bits. The main benefit is a large reduction in the dynamic power range required. Even more striking than the performance differences among the various multiple access schemes is the large statistical multiplexing gain achieved as more channels become available. In every instance, going from 4 to 16 channels more than doubles the bits per symbol per channel. This illustrates an easily obtained gain that might be achieved in using multi-band or cognitive radios: treating all the available bands as a block gives far superior performance than performing individual management with non-agile radios. Since many current and proposed systems have very few channels at the higher bit rates, the implications for reliable transport of interactive media traffic are obvious. It is a simple consequence of classical Erlang models. Of course, different conclusions would result from traffic models that yield less statistical multiplexing gain.

Similar conclusions apply for the downlink, although the tradeoffs are slightly different. Users near cell boundaries experience increased interference and propagation loss compared to those near base stations, whereas in the uplink case they only suffer increased loss. (In that case, they generate more interference). Compared to the uplink, maximum throughputs are slightly lower in the downlink (0.1 to 0.2 bits). The major difference is that for power control the most effective rule changes from q=2.5 to q=3.

IV. STATIC FADING

Distance loss is fundamental to all cellular systems, allowing frequency re-use and in the end dominating over other losses. However, it is also interesting to examine the effect of large fading. We will not consider dynamical fading, which if it takes place on a rapid time scale can be largely mitigated through some averaging means (e.g., buffering and diversity in time). Rather we will consider (quasi-) static fading, as might occur in indoor environments. We begin with the simplest possible situation: the two-cell model with two multiplicative power fade levels of either $1/f$ or f, here with f=9 (8 dB in power). It is assumed that the fade levels are independent for the gain to a user's own base station and to its neighbor. Even in this case, there are 16 fading combinations for losses to a user's own base station (high or low for each of 4 locations), and 16 interference combinations for a total of 256 in all. So a user might have gain to its base station and extra loss to a neighbor, or loss to its own base station and gain to its neighbor. Pairing of users will be of little help in such situations due to the asymmetry of losses and interference; an arrangement that is good for one user is unlikely to be helpful for the neighbor, and even less so in more general situations with multiple neighboring cells each with independent fading. In general, each different fading realization will have different channel assignment, power and bandwidth optimizations.

For example, let g_{ij} denote the fading gain between user i and base station j. Suppose in cell 1, the two users have gains g_{11} ={1/9,9}, and g_{12}={9,1/9} while in cell 2 we have g_{22} ={1/9,9} and g_{21} ={1/9,9}. The interference gain also includes the distance propagation loss factor of either 81 or 625 depending upon whether the neighbor is near or far. We will further suppose for simplicity of illustration that all cells have identical neighbors with identical fade parameters. Thus each neighbor appears once as being close and once as being far. Suppose first we try to equalize the bit rates for users sharing the same channel, using power control. There are two possible channel assignments:

| 1 2 | 1 2 |… or | 1 2 | 2 1 | …

We illustrate the first one. For channel 1, we have for cell 1,

$$SIR_1 = \frac{P_1 g_{11}}{P_2 g_{21}(1/81 + 1/625)} = \frac{P_1/9}{P_2/640}.$$

Similarly, for cell 2 we have

$$SIR_2 = \frac{P_2 g_{22}}{P_1 g_{12}(1/81 + 1/625)} = \frac{9P_2}{P_2/7.97}.$$

Equating the SIR's, we obtain $P_2 = 8.96P_1$, or a capacity of 3.16 bits. For channel 2, a similar calculation obtains $P_1 = 8.96P_2$ for a capacity of 9.33 bits. In this example, exactly the same rates are obtained in the second assignment. Thus, time-sharing will be required to achieve fairness. In this case, the left-most users in both cells will almost half the time have both channels, resulting in a fair rate of 4.72 bits.

Now suppose power is fixed, and we can only use time-sharing for fairness. We obtain the following results for the two channel assignments:

(1): Ch. 1—SIR={80.3,0.886}, Ch. 2—SIR={71.7,5760}

(2): Ch. 1—SIR={0.886,80.3}, Ch. 2—SIR={5760,71.7}

Limiting the bits per symbol to 10, and time averaging among the assignments to equalize rates in the two cells, location 1 in each cell gets 3.63 bits and location 2 in each cell gets 8.09 bits. Time sharing so that the "channel 1" user pairs get two channels 38% of the time, a fair rate of 5.01 bits is obtained.

The situation in CDMA is interesting. In the limit of a very large spreading gain and a large number of channels, each fading independently, the interference from neighboring cells under certain special conditions is constant. Specifically, we require that power be adjusted only in response to losses to one's own base station (and not to meet a target SIR) and that bandwidth adjustments be compensated through decreased power spectral density. In this particular case, power control is sufficient for fairness if it fully compensates for fading loss. However, because it increases the interference, the fair bit rate is only 2.2 bits. We do better *on average* with power and bandwidth allocation, where $P=(\text{loss})^{q/4}$. In this case, q=1 gave the best results at 3.2 bits, with the ratio of max to min bandwidth being roughly 6. The main price is that the range of bit rates depending on the fading realization ranged from 1.9 to 6.2, whereas with full compensation for loss the rates were flat. We also evaluated CDMA in the condition of only two channels per cell, assuming full occupancy. There are now 16 different interference cases for each of the four fading loss conditions for a cell, each demanding a separate power and bandwidth allocation. We tallied the average rates obtained for each value of q, and found that the best was q=3, for which a fair rate of 5.03 bits was achieved. However, now the minimum rate achieved was 0.91 bits and the maximum hit the ceiling of 10 bits, depending upon the fading realization. In practice, this makes it much more difficult to consistently achieve a desired bit rate.

Turning to the situation of interest, namely, a finite number of channels and a fixed blocking probability, the increased variability in achievable rates due to fading will clearly decrease capacity. Since the best bandwidth allocation depends upon the fading realization, it is apparent that some type of adaptive mechanism will be required. It is also clear that power control will be important, just as for distance loss alone. Therefore, we simulated the effect of adding lognormal fading with variance of 8 dB to the distance loss for our six-cell model. In Table 3, we compare coordination with OFDMA, using the same allocation algorithms as in the previous two sections. Coordination was now conducted using the fourth root of the loss as the measure for the hash sort. With large fading the resulting alignments are less effective, and consequently we found that a larger power exponent was required for best results. OFDMA also benefited from a higher exponent, with best performance for full compensation for the channel losses. Compared to Table 2, for the best q values, the loss for coordination is roughly 1.2 bits for large n, but only 0.6 bits for OFDMA. The larger figure for coordination perhaps reflects the decreased effective alignment caused by the shadow fading. Consequently, OFDMA is competitive with coordination in this setting, at the cost of a greater dynamic range of power control. We again observe that the choice of power control law is very important, with large losses for q=0. In particular, for both coordination and OFDMA the algorithm was unable to generate fair allocations for 1% blocking probabilities for n>16. Average service times displayed a spread of greater than 10% between low and high loss locations in this instance. We also ran some simulations for CDMA in fading, but the algorithm failed to provide fair allocations at any reasonable offered load, with and without power control. The likely reason is that large interferers affect the channels of all others, rather than simply one subchannel, as in OFDMA.

n	Coord q=3	Coord q=0	OFDMA q=4	OFDMA q=0
4	1.7	1.5	1.5	1.4
8	2.6	2.3	2.5	2.1
16	3.5	3.2	3.3	2.6
32	4.1		3.8	
64	4.4		4.1	
128	4.7		4.3	
256	4.8		4.5	
512	5.0			
1024	5.1			

Table 3: Bits per channel per unit time for distributed algorithm in fading

Potentially, better results could be obtained with power control and channel selection mechanisms that were more responsive to the particular fading conditions. For distance loss only, the channel indexing scheme is sufficient to slot high loss users in low interference channels with high probability. With fading, an explicit SIR-based slotting might do better, but considerable modifications to the algorithm would be required. Harder still, but potentially with more gain, would be algorithms that also consider the interference effects on nearby cells. This was also accomplished in the case of distance-only losses, but due to the asymmetry in fading would require information exchange between cells in the more general setting.

V. Conclusion

In this paper we have considered two contrasting delay constraints: infinite and (near-) zero. The former corresponds to data, while the latter is characteristic of multimedia. In the limit of large numbers of channels, the two cases are essentially the same in that almost fully loading all channels is possible with low blocking probabilities. Moreover, this leads to the highest throughput. We have shown that while the infinite delay case leads to some insights in the design of resource allocation algorithms, (e.g., superiority of combined power/bandwidth allocation over power control for CDMA, superiority of power control with interferer alignment over bandwidth allocation), it misses very important effects when the number of channels is small to moderate. Since this is the usual situation in current high speed access networks such as WiFi, it is important to understand the design tradeoffs. Our results indicate large statistical multiplexing gains can be obtained if channels can be aggregated, as for example in multi-band and cognitive radios.

Computing the optimal throughput in such cases is

difficult because effectively we are dealing with coupled queues. We explored the limit where the number of channels goes to infinity and loading is high, so that there is full occupancy. However, further progress in algorithmic design required simulations. Our results for different models point to different facets of algorithm design. Designing for the condition of infinite queuing (full loading) points to the principle of aligning high loss channels with low interference channels in neighboring cells. This significantly reduces the degree of power disparity required for fairness. Less effectively, it also somewhat reduces the relative bandwidth required for fairness. The two-point case illustrated the benefits for capacity in aligning the vacant channels in both nearby cells together where possible, to create more low interference channels. This alignment of vacancies in neighbors is especially helpful when the number of channels is small, whereas the interferer alignment is more important (and effective) when the number of channels is large. The general principle is that we should try to arrange for low interference channels through alignment of vacancies and low interference users, and slot the high loss users in such channels. This is balanced against increased resources to disadvantaged users; as they get more power and bandwidth their bit rate improves but so does the interference they generate. Using the channel coordination, we can reduce the resources and thus the interference, and therefore produce higher throughput. This accounts for the superiority of our allocation methods over CDMA at high loading. Our results for the four-point case illustrated that power control alone could not achieve fairness in all loading conditions. The best scheme combined alignment of channel vacancies and interference with power control and variable bandwidth. The uplink and downlink cases differed in some details, most notably the power control rule to use, but had similar tradeoffs.

In the most realistic scenario of distance loss together with log normal fading, the distributed algorithm with alignment continued to outperform OFDMA. However, the margin was much reduced due to the reduced effectiveness of our simple channel indexing scheme in pairing high loss users with low loss neighbors. The asymmetry of fading among different links considerably complicates the design of optimal algorithms.

A large number of interesting problems remain. The results reported in this paper all apply to one-dimensional cells with static channels. We believe that systematically building from simple to more complicated models can provide similar insights in other settings, albeit with more effort. Coordination of low interference channels is much easier to achieve in one dimension than in two, since there are two rather than six neighbors in a regular packing. We conjecture that the gap between optimal power and bandwidth allocation with coordination and a randomized OFDMA scheme should decline in two dimensions. However, it is significantly more difficult to assess what the optimal resource allocation scheme would be in two dimensions. In this paper we have considered a hierarchy of responses to the inherent unfairness of access in wireless systems: alignment of vacancies and low interference users in nearby cells, prioritized pairing of disadvantaged users with such channels, increased bandwidth, and increased power. The next logical step is to consider cancellation of the strongest remaining interferers, and how cancellation interacts with these simpler measures. That is, to what extent they can help or hinder interference cancellation, especially given practical limitations in synchronization, knowledge of channel parameters, SNR, etc. that influence the degree of cancellation achieved.

Finally, there is the question of performance in dynamical settings, where channel conditions change quickly, rendering the more elaborate optimizations less effective. Our distributed algorithm has essentially zero convergence time, in contrast to iterative power control schemes where convergence time depends upon the degree of interference coupling. However there may be stability questions in rapidly changing channels.

ACKNOWLEDGMENT

We thank Stephen Hanly for suggesting the bandwidth and power allocation comparison for CDMA, and Yue Zhao and Suhas Diggavi for many interesting discussions.

REFERENCES.

[1] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, Inc., New York, 1991.

[2] Cadambe, V.R.; Jafar, S.A., "Interference Alignment and Degrees of Freedom of the K-User Interference Channel," *IEEE Trans. Information Theory*, vol.54, no.8, pp.3425-3441, Aug. 2008

[3] M. Maddah-Ali, A. Motahari, and A. Khandani, "Communication Over MIMO X Channels: Interference Alignment, Decomposition, and Performance Analysis," *IEEE Trans. Information Theory*, vol. 54, no. 8, pp. 3457–3470, 2008.

[4] A. Avestimehr, S. Diggavi, and D. Tse, "A deterministic approach to wireless relay networks," *Proceedings of Allerton Conference on Communication, Control, and Computing, Illinois, USA*, Sept 2007.

[5] Zander, J., "Performance of optimum transmitter power control in cellular radio systems," *IEEE Transactions on Vehicular Technology*, vol. 41, no. 1, pp. 57-62, Feb 1992.

[6] Bambos, N.; Pottie, G.J., "Power control based admission policies in cellular radio networks ," *IEEE Globecom*, pp. 863-867 vol. 2, 6-9 Dec 1992.

[7] G.J. Pottie, "Wireless multiple access adaptive communication systems," in *Encyclopedia of Telecommunications* vol. 18, Ed. Froelich and Kent; Dekker, 1999, pp. 1-41.

[8] Y. Wei and J.M. Cioffi, "FDMA capacity of Gaussian multiple-access channels with ISI," *Communications, IEEE Transactions on* , vol. 50, no. 1, pp. 102-111, Jan 2002.

[9] I. Katzela and M. Naghshineh, "Channel assignment schemes for cellular mobile tele-communication systems: a comprehensive survey," *IEEE Personal Comm*., pp. 10-31, June 1996.

[10] J. Huang, R. Cendrillon M. Chiang and M. Moonen, "Autonomous Spectrum Balancing (ASB) for Frequency Selective Interference Channels," *IEEE International Symposium on Information Theory*, pp. 610-614, July 2006.

[11] J. Bae, E. Beigman, R. Berry, M. Honig, and R. Vohra, "Sequential Bandwidth and Power Auctions for Distributed Spectrum Sharing," *IEEE Journal on Selected Areas in Communication,* vol. 26, no. 7, pp. 1193-1203, September 2008.

[12] K.N. Sivarajan, R.J. McEliece, and J.W. Ketchum, "Dynamic channel assignment in cellular radio," *IEEE 40th Vehicular Technology Conference*, 1990, pp.631-637, 6-9 May 1990.

[13] H. P. Shiang and M. van der Schaar, "Queuing-Based Dynamic Channel Selection for Heterogeneous Multimedia Applications over Cognitive Radio Networks," *IEEE Trans. Multimedia*, vol. 10, no. 5, pp. 896-909, 2008.
[14] I.E. Teletar, "Capacity of multi-antenna Gaussian channels," *Eur. Trans. Telecom,* vol. 10, pp. 585-596, Nov. 1999.
[15] H. Weingarten, Y. Steinberg, and S. Shamai (Shitz). The Capacity Region of the Gaussian Multiple-Input Multiple-Output Broadcast Channel. *IEEE Trans. Information Theory*, vol. 52, no. 9, Sept 2006.
[16] S. Haykin, "Cognitive radio: brain-empowered wireless communications," *IEEE JSAC*, vo. 23, No. 2, pp. 201-220, Feb. 2005.
[17] S.M. Mishra, A. Sahai and R.W. Brodersen, "Cooperative Sensing among Cognitive Radios," *IEEE ICC '06,* vol.4, pp. 1658-1663, June 2006.
[18] D. Cabric, I.D. O'Donnell, M.S.-W. Chen and R.W. Brodersen, "Spectrum sharing radios," *IEEE Circuits and Systems Magazine*, vol. 6, no. 2, pp. 30-45, 2006.
[19] IEEE 802.22 Working Group on Wireless Regional Area Networks, http://www.ieee802.org/22/.
[20] Gross and Harris. Fundamentals of Queuing Theory. Wiley: 1984.