

A Case for Adaptive Sub-carrier Level Power Allocation in OFDMA Networks

Shailendra Singh
University of California, Riverside
singhs@cs.ucr.edu

Karthikeyan Sundaresan
NEC Labs America Inc., NJ
karthiks@nec-labs.com

Moloud Shahbazi
University of California, Riverside
mshah008@cs.ucr.edu

Srikanth V. Krishnamurthy
University of California, Riverside
krish@cs.ucr.edu

Konstantinos Pelechrinis
University of Pittsburgh
kpele@pitt.edu

Sateesh Addepalli
Cisco Systems, Inc.
sateeshk@cisco.com

ABSTRACT

In today's OFDMA networks, the transmission power is typically fixed and the same for all the sub-carriers that compose a channel. The sub-carriers though, experience different degrees of fading and thus, the received power is different for different sub-carriers; while some frequencies experience deep fades, others are relatively unaffected. In this paper, we make a case of redistributing the power across the sub-carriers (subject to a fixed power budget constraint) to better cope with this frequency selectivity. Specifically, we design a joint power and rate adaptation scheme (called JPRA for short) wherein power redistribution is combined with sub-carrier level rate adaptation to yield significant throughput benefits. We further consider two variants of JPRA: (a) JPRA-CR where, the power is redistributed across sub-carriers so as to support a maximum common rate (CR) across sub-carriers and (b) JPRA-MT where, the goal is to redistribute power such that the transmission time of a packet is minimized. While the first variant decreases transceiver complexity and is simpler, the second is geared towards achieving the maximum throughput possible. We implement both variants of JPRA on our WARP radio testbed. Our extensive experiments demonstrate that our scheme provides a 35% improvement in total network throughput in testbed experiments compared to FARA, a scheme where only sub-carrier level rate adaptation is used. We also perform simulations to demonstrate the efficacy of JPRA in larger scale networks.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

General Terms

Algorithms, Experimentation, Measurement, Performance

Keywords

Rate Adaptation, Power Allocation, Orthogonal Frequency Division Multiplexing (OFDM), Frequency Selective Fading

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiHoc'12, June 11–14, 2012, Hilton Head Island, SC, USA.
Copyright 2012 ACM 978-1-4503-1281-3/12/06 ...\$10.00.

1. INTRODUCTION

Due to its inherent ability to cope with fading, OFDMA is employed today in many commercial wireless systems. The current implementations as specified by the IEEE standards [2] for such systems do not include schemes for adaptively assigning powers to subcarriers based on CSI (Channel State Information); instead, these systems evenly distribute the total transmission power budget across all subcarriers (waterfilling [21] with equal power without CSI knowledge).

However, at a receiver, since each sub-carrier typically experiences a different fade (as seen both by our work and in prior related efforts such as [16]) the transmission rates that can be supported by the different sub-carriers can differ. Spatial and temporal diversity can further complicate the communications between a single sender and multiple receivers at different times. Each receiver will likely have a different multipath fading profile given the differences in receiver locations. Even for the same receiver, the fading effects experienced by the different sub-carriers are expected to differ in time.

Traditional rate-adaptation algorithms try to cope with fading by varying the transmission rate in response to either packet losses (e.g., [23]) or signal-to-noise ratio (SNR) variations (e.g., [1]). However, these are not done at a per sub-carrier level. Thus, poor received quality on a few sub-carriers can affect the supported transmission rate. To cope with this, Rahul *et al.* [16], propose FARA, a scheme that supports per sub-carrier rate adaptation. While FARA provides gains over traditional rate-adaptation schemes, it does not address the problem of frequency-selectivity directly and does not utilize the available power budget efficiently.

In this paper, we propose redistributing the power across the sub-carriers to better cope with frequency selectivity. Specifically, we argue for an adaptive, scenario specific, uneven distribution of transmission powers across sub-carriers. By appropriately assigning the transmission powers to the different sub-carriers (while adhering to the total power budget for the transmission), and by combining this with the choice of appropriate per sub-carrier transmission rates, we envision achieving significant gains in throughput. As our main contribution, we design a joint power and rate adaptation (JPRA) scheme towards realizing these envisioned benefits.

Calibration phase in JPRA: In order to determine the right power levels for each sub-carrier, a calibration phase is necessary for JPRA. The goal of calibration is twofold; (i) to identify the *practical* decoding thresholds for each transmission rate, and (ii) to correlate the required change in the received signal of a sub-carrier (in order to support a rate) with an appropriate change in the corresponding transmit power at the sender. We perform several measurements to understand if such a calibration phase is viable. We use

the error vector magnitude (EVM) [13] to determine the minimum power level needed for sustaining a specific transmission rate. We find that these minimum powers are consistent over a large number of scenarios and node locations. This suggests that infrequent benchmark measurements are sufficient to calibrate JPRA and utilize it for adaptive sub-carrier level power allocation during data-transfers.

Variants of JPRA: We propose two variants of JPRA. The first variant, named JPRA-CR (where CR stands for common rate), selects the power-assignments such that a single *best* rate can be used on an appropriately chosen set of sub-carriers. This best rate corresponds to the highest rate for which the EVM threshold is satisfied for the set of sub-carriers at the receiver. Since we operate on a fixed power budget, the scheme entails the transfer of power from relatively good sub-carriers (unaffected by fading) to those that are in deep fades. The advantage of this approach is that it is simple and reduces the transceiver complexity (since they only need to decode at a single rate). However, this may not yield the maximum achievable throughput with power re-allocation. As shown later however, it does offer an additional benefit in ensuring that transmission rate transitions are less likely than in traditional approaches (stabilizes the rate in use). In addition, it can potentially provide significant power savings for a target PDR (packet delivery ratio), compared to traditional systems.

Towards maximizing throughput via sub-carrier level rate allocation, we propose a second variant JPRA-MT. Our goal with this scheme is to minimize the total transmission time of a packet by appropriate power re-distribution. Essentially, with JPRA-MT sub-carrier powers are assigned to sub-carriers such that the total number of bits (as mapped to symbols on a constellation) transmitted on the sub-carriers in each symbol-duration is maximized. To illustrate with an example, consider a case where we have just two sub-carriers. The power assignments can be such that either both sub-carriers can be modulated using QPSK (2 bits per symbol) or, one sub-carrier can be modulated with BPSK (1 bit per symbol) while the other can be modulated with 16-QAM (4 bits per symbol). In this case, the packet airtime is minimized with the latter choice since more bits are transmitted per symbol duration (5 bits instead of 4 bits) and is consequently chosen by JPRA-MT. We wish to point out here that in order to support a higher modulation, a higher received power is required. Thus, in the latter case the power on the second sub-carrier (supporting 16-QAM) is increased compared to the power used in the former case (when it supported QPSK). Similarly, the power on the first sub-carrier has to be decreased in the latter case compared to the former one (and thus, it can support only BPSK as opposed to QPSK).

We implement both the versions of JPRA on our six node WARP radio testbed. We also implement FARA [16] for comparison. We perform extensive experiments which show that JPRA-MT outperforms FARA by as much as 35 % in terms of throughput. The gains are much more significant compared to traditional OFDMA systems (75 %). We also show the efficacy of JPRA in larger scale settings via simulations.

The main properties of JPRA are summarized below:

- *Throughput Efficiency:* JPRA-MT achieves 75% more network throughput than standard SNR based rate adaptation. It also outperforms the state of the art OFDMA rate adaptation scheme, FARA[16] by 35% in terms of total network throughput.
- *Potential Power savings:* JPRA-CR is more power efficient compared to systems with no power re-distribution. Specifi-

cally, for a fixed BER (bit error rate), we observe upto 4.5 dBm in power savings.

- *Stability of Transmission Rates :* JPRA-CR reduces the number of rate changes by 27% in comparison to standard SNR based rate adaptation.
- *Impact on Carrier Sensing:* JPRA does not affect the RSS (received signal strength) of packets since the total power budget remains fixed.

We acknowledge that our system is currently applicable only on (quasi) static topologies. The calibration phase required for the deployment of our scheme cannot be as effective in scenarios that include mobility. However, we seek to examine mobility scenarios as part of our future work.

Organization: The paper is organized as follows. In Section 2, we provide relevant background and overview related work. The calibration phase of JPRA is described in Section 3. Details of the two JPRA variants for power/rate allocation are in Section 4. Section 5 describes both our experimental and simulation results. Section 6 concludes the paper.

2. BACKGROUND AND RELATED WORK

In this section, we first describe studies related to our work and then we briefly provide background on the EVM calculation required by our system.

Frequency Selective Fading: There are studies that employ rate adaptation to cope with frequency selective fading. For example, in [16] the authors propose Frequency Aware Rate Adaptation (FARA) to improve system performance. They assign subsets of sub-bands to each sender-receiver pair and based on the SNR reported by the receiver on these sub-bands, the sender performs rate adaptation. However, the authors do not propose a solution to improve the performance of sub-carriers experiencing frequency selective fading or low SNR. Barthia *et al* [5] propose a smart mapping of symbols to sub-carriers. This supports partial recovery of symbols if they are lost due to frequency selective fading. They also propose an extra layer of FEC codes on top of Physical layer FEC. One of the main limitations of this work is that the proposed solution is only compatible with block FEC schemes and its not clear how it will work with convolutional or turbo code FECs [2]. A large amount of feedback information is also required for the partial symbol recovery. In contrast our scheme is not limited by the choice of FECs and the amount of feedback information is low (we only need to send sub-carrier power values and rates).

Rate Adaptation : There is a large volume of studies on rate adaptation (e.g., [6, 1]). SampleRate [6], proposed by Bicket *et al.*, probes the performance at a random rate every 10 frames, and selects the rate that minimizes the expected transmission time (including retransmissions). Wong *et al.* [23] develop Robust Rate Adaptation Algorithm (RRAA), which uses short term loss ratios to opportunistically change rate. It further incorporates an adaptive RTS filter to prevent collision losses from lowering data rates. All these (and many similar) existing schemes adapt rate according to frame loss rates. RBAR [8] uses the RTS/CTS exchange to estimate the SNR at a receiver, and picks the transmission bit rate accordingly. OAR [17] further builds on RBAR, by opportunistically transmitting back-to-back frames when the channel quality is good. CHARM [10] leverages the reciprocity of the wireless channel to estimate the average SNR at the receiver using packets overheard from the latter. The overhead of RTS/CTS (present with

RBAR and OAR) is thus avoided and implementation on commodity cards is enabled. Sen *et al* [18] propose the use of EVM (Error Vector Magnitude) to perform rate adaptation. Since the above rate adaptation schemes use information such as loss ratio, SNR and EVM averaged over a packet, they fail to capture the effects of frequency selectivity. On the contrary, we try to directly address issues related to frequency selective fading by using per subcarrier EVM measurements.

Multi-User OFDM : There exist a few studies on subcarrier, power and rate assignments in multi-user scenario. In this case a single channel is shared among multiple users. In [12] Javidi *et al* propose a scheme to outperform “water filling based multi-user subcarrier assignment” by introducing a subcarrier allocation scheme which takes the queue lengths of different users into account. Adaptive power allocation for a multiuser OFDM environment has been proposed in [22, 9]. An OFDM channel is divided into multiple subbands and these subbands are assigned to different users. To alleviate frequency selective fading experienced by different users, redistribution of power and modulation on these subbands is done according to the SNR experienced by the users. However, the majority of these studies are evaluated only via simulations by making assumptions about channel conditions. Furthermore, power adaptation is done for each user and not on a per OFDM sub-carrier basis.

Error Vector Magnitude(EVM): In this work we use the EVM per subcarrier as the CSI (carrier state information) feedback from the receiver, for the sender to perform power redistribution and rate selection. EVM is a vector measurement taken in terms of peak (or rms) percentages between the ideal symbol position and the actual measured position in the constellation space for a particular modulation. The error vector is a vector in the I-Q plane between the ideal constellation point and the data interpretation by the receiver[13]. In other words, it is the difference between actual received symbols and ideal symbols. The average power of the error vector, normalized to the signal power, is the EVM. It can be expressed as a percentage:

$$EVM(\%) = \sqrt{\frac{P_{error}}{P_{reference}}} * 100$$

where, P_{error} is the RMS power of the error vector and $P_{reference}$ is defined as the reference constellation average power. In contrast to SNR, higher EVM values correspond to bad channel conditions, while lower EVM values represent good channel conditions.

3. CALIBRATING PHASE OF JPRA

JPRA is a measurement driven system and requires a set of calibrating measurements. These measurements are:

1. *Per sub-carrier EVM for each received packet.* This provides the system, knowledge with respect to both the sub-carriers that are experiencing deep fade and those that are relatively unaffected.
2. *EVM threshold for supporting each available rate.* This identifies the modulation that a sub-carrier can support.
3. *A mapping between a power adjustment and the corresponding EVM change.* This is critical in determining the effect of a power change on the rate that can be supported by a sub-carrier. In other words, it seeks to answer: By increasing (or decreasing) the transmission power on a sub-carrier by a certain magnitude, what is the higher (or lower) rate that that can be achieved on that sub-carrier ?

Note that the the per sub-carrier EVM should be obtained for every received packet (both during the calibration phase and when JPRA is being employed). The latter two measurements are to be obtained prior to applying power redistribution in JPRA, for the purposes of calibrating the system.




Figure 1: Calibration phase of JPRA

In what follows we provide more details on the way we obtain these measurements. Figure 1 outlines the main steps involved in calibrating JPRA. We defer a detailed description of our testbed to later; simply put, we have a six node WARP testbed that we use for all of these calibrations.

Calculating per sub-carrier EVM: EVM is an indicator of modulator or demodulator performance in the presence of impairments. For a receiver to calculate the EVM, both received and transmitted packets are required. The WARP boards provide the per sub-carrier EVM for every correctly received packet. If the packet passes the CRC check, the receiver can reconstruct the exact signal that was transmitted. The bits are re-mapped on to a constellation space to obtain ideal symbol positions based on the modulation used for the transmission. With the ideal symbol positions, it is easy to calculate the EVM as described in Section 2. With the above process it is not possible to calculate the EVM for packets that fail the CRC check; the receiver cannot deduce what was transmitted and thus, cannot obtain the ideal symbol positions. Thus, the EVM estimates are updated only based on packets that are correctly received.

Mapping rate to the EVM threshold: In order for a symbol on a sub-carrier to be demodulated with high probability, its EVM at the receiver needs to be below a specific threshold. This threshold is different for different modulation schemes. Similar thresholds has been determined in terms of SNR [16], but cannot be directly used with JPRA. The mapping between the EVM threshold and rate are calculated empirically via measurements, by varying the transmission power and the locations of the sender-receiver pair. The maximum EVM value that yields a target PDR (90% in our case) provides us with the required information. We observe that this mapping does not change from one setting to another. However, we recognize that it may vary for different hardware platforms. Thus, although we use one time pre-deployment measurements for calibration, in practice, they may have to be periodically repeated. The mappings that we obtain are presented in Table 1 and used to select the transmission bit rates with JPRA.

Estimating power adjustments for specific EVM changes: To perform joint rate and power adaptation we need to estimate the magnitude of increase/decrease in the EVMs of the received sub-carriers, for given changes in the transmit power. This mapping will facilitate the power redistribution across the available sub-carriers. In the current implementation of OFDM on WARP Boards, the total transmission power (called the power budget) is distributed equally among the available sub-carriers. The radio board applies power to the analog baseband/RF waveforms. In other words, the time domain signal after the IFFT stage is processed for power alloca-




Figure 2: Average change in EVM(%) with change in sub-carrier Power from 0 to 1 for QPSK.




Figure 3: Average change in EVM(%) with change in sub-carrier Power from 1 to 2 for QPSK.




Figure 4: Average change in EVM(%) with change in sub-carrier Power from 0 to 1 for 16QAM.

tion. Since, in the time domain the individual sub-carriers are not explicitly visible, the power is uniformly distributed across all frequencies used. To achieve a non-uniform power allocation across the sub-carriers, we process each sub-carrier prior to applying the IFFT. The coefficient used to establish the power level of a sub-carrier does not correspond to actual transmission value but is a scaling factor. It takes real values between 0 and 2, with 1 representing the default power (uniform distribution of the power budget across sub-carriers), 0 mapping to zero transmission power and 2 to the maximum transmit power allowed (double the default transmission power for that sub-carrier).

In order to determine the scaling factor described above, we perform another set of calibration measurements. We use a pair of WARP board transceivers and we vary the transmission power of a sub-carrier by scaling its transmission power as explained above, with a step size of 0.1. At each step the transmitter sends back to back packets for 3 minutes; we perform 10 trials for each step. Figures 2 and 3, show the average EVM(%) change when QPSK is used when decreasing and increasing the transmission power of a sub-carrier, respectively. Analyzing the obtained data we find that a third-degree polynomial can fit the data fairly accurately (the corresponding R^2 value¹ is 96%). Consequently, this polynomial can be used to predict the change in EVM(%) with respect to sub-carrier power. We also looked at different modulations and we observe that qualitatively the results do not change with the change in modulation. To illustrate, Figure 4, depicts the results obtained when 16-QAM is used. A polynomial of degree 3 can again be used to fit the data (R^2 value is 95%). As one might expect though, the coefficients of the polynomial are different as compared to the QPSK case since the EVM(%) range is different. Finally, we repeated the above experiments using different sub-carriers obtaining similar results. In particular, the fitting polynomial coefficients for different sub-carriers for the same modulation are very similar. The above measurements are obtained for static links and the results do not significantly change for small variations in the transceivers' locations. However, for large deviations from the initial measurements positions the results will not hold; in particular, the polynomial coefficients will be different. Therefore, we need to perform these calibrating measurements when a new link appears or when the topology changes significantly. This limits the applicability of the current version of JPRA to static or slowly changing topologies. Extending JPRA to more dynamic settings is deferred to the future.

4. POWER/RATE ALLOCATION WITH JPRA

In this section we describe the two versions of JPRA in detail. In a nutshell, for both schemes, each receiver upon packet reception calculates the per sub-carrier EVM. Using this information JPRA

Maximum EVM(%)	Modulation	Coding
18.0	BPSK	1/2
10.2	BPSK	3/4
6.6	QPSK	1/2
4.0	QPSK	3/4
1.67	16-QAM	1/2
1.26	16-QAM	3/4
1.1	64-QAM	1/2

Table 1: EVM to bit-rate(modulation + coding) mapping

redistributes the power among the sub-carriers depending on the objective (i.e., sub-carriers use the same maximum possible common rate with JPRA-CR *or* the the packet air time is minimized while allowing sub-carriers to be modulated at different rates with JPRA-MT). Figure 5 depicts the high level functionalities of JPRA.

Upon executing the appropriate JPRA version, the receiver obtains a tuple for each sub-carrier specifying (i) the magnitude of increase or decrease in transmission power and (ii) the corresponding bit rate to use; this is represented as $(carrier\#, TX\ power, Rate)$. This is transmitted to the sender in the ACK (or NACK). The sender can then adjust its sub-carrier's transmission powers and bit-rates for the next packet. Note that in the case of JPRA-CR, all tuples will have the same rate information. If the channel conditions remain constant for a train of packets, the receiver may omit sending this information with every packet and the sender will use the latest sub-carrier settings for the next packet.

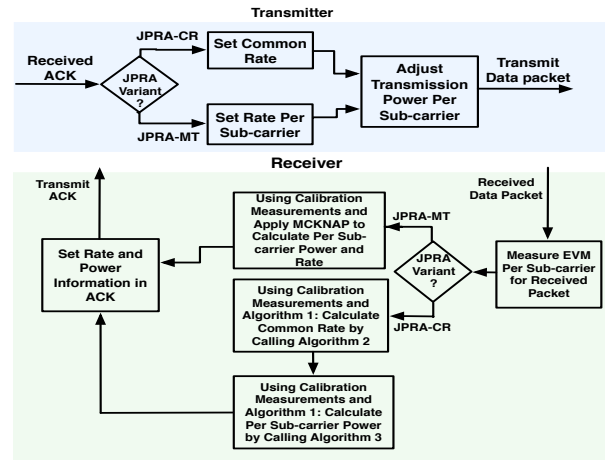


Figure 5: Power and Rate allocation with JPRA

¹ R^2 value[19] shows the goodness of fit of a model.

4.1 Achieving a common rate with JPRA-CR

Due to frequency selectivity, for a specified bit-rate, some of the sub-carrier's EVMs will be higher than the maximum tolerable EVM for that rate (class 1 sub-carriers); on the other hand, some others will have EVMs lower than what can be tolerated (class 2 sub-carriers). In conventional OFDMA systems, the modulation used is conservative. In particular, all sub-carriers are modulated with the minimum supported rate from among them. One can make a case for redistributing the *extra* power from class 2 sub-carriers to class 1 sub-carriers. This in turn, will reduce the EVM of the latter and can potentially allow the use a higher bit-rate modulation on all the sub-carriers. However, selecting a common rate as above for all the sub-carriers may not be optimal since there may exist carriers (in deep fades) that may be unable to carry data at acceptable rates. Thus, with JPRA-CR we propose to find a common rate on a *subset* of sub-carriers such that the aggregate rate is maximized. This constitutes removing sub-carriers which require high powers per bit delivered, and redistributing their powers to other sub-carriers; at the same time, it is possible that powers from some of the best sub-carriers are also carried over to other sub-carriers towards achieving the above objective. Note that there is an inherent trade-off here; while eliminating poor sub-carriers can increase the common rate that can be supported on the other sub-carriers, it also decreases the number of available sub-carriers and thus, contributes to a lowering of the capacity. JPRA-CR tries to find the best point of operation while accounting for this trade-off.

Let us assume, for simplicity, that all sub-carriers are considered for setting the common rate. Upon receiving a packet, the receiver iterates over the list of available rates that are greater than the currently supported lowest rate, to identify the maximum bit-rate that all the sub-carriers can satisfy after power redistribution. In each iteration, it has to solve the problem of per sub-carrier power redistribution for the specific rate considered. One can map this problem to a form of the well-known Knapsack problem (which we describe below) and based on this, we design JPRA-CR.

The solution to the considered problem is as follows. In every iteration, the receiver first subtracts the EVM of the sub-carriers under consideration from the EVM threshold of the rate under consideration (obtained from the calibration phase). For some sub-carriers this value will be negative, while for others it will be positive. If there are no negative values (the EVM of every carrier is lower than the rate's EVM threshold), all sub-carriers can be modulated using this rate. If not, it maps the EVM differences to the corresponding sub-carrier transmission power increases required in order to meet the threshold (negative weights), or the excess transmission powers present in the sub-carriers whose EVM adheres to the required threshold (positive weights). Again, this mapping is done based on the measurements from the calibration phase. The positive weights are summed to get the total extra energy available (in Knapsack problem terminology, the total weight that can be carried by the knapsack). Each of the absolute values of the negative weights is considered as the weight of the corresponding sub-carrier (item). Our objective is to select as many items as possible with the given total weight. Ideally, we want all items to be included in our Knapsack. If this can be achieved, all sub-carriers can be modulated with the rate considered and we further examine the next available higher rate. If this cannot happen, we terminate our iterations, since it is guaranteed that no other higher rate can be supported (the proof is trivial and we omit this for space reasons).

The problem of finding the maximum number of sub-carriers that can satisfy the given EVM threshold requirement (for the considered rate j) at each iteration, can be formalized by the following optimization problem. Let us assume that n sub-carriers are above

```

Input: Total number of sub-carriers  $n$ , Per sub-carrier  $EVM_i$ , Rate
 $j \in S \equiv \{6, 12, \dots\}$ , Rate threshold  $EVMTH_j$ , EVM to
Power Mapping, Initial Per sub-carrier Power  $p_i^j$ 
Output: Rate  $R(m)$ , Per sub-carrier Power  $p_k''(m)$ 
 $R(n) = \text{Algorithm-2}(EVM_i, \text{Rate } j \in S, EVMTH_j, \text{EVM to}$ 
Power Mapping);
Aggregate Rate  $AG_n = R(n) \times n$ ;
Remove Sub-carrier which requires highest amount of power per bit
delivered and add it's power to total excess power  $W$ ;
 $m \leftarrow n - 1$ ;
while  $m > 1$  do
   $R(m) = \text{Algorithm-2}(EVM_i, \text{Rate } j \in S, EVMTH_j, \text{EVM}$ 
to Power Mapping);
  Calculate  $AG_m = R(m) \times m$ ;
  if  $AG_m < AG_{m+1}$  then
     $p_k''(m+1) = \text{Algorithm-3}(p_i^j, \text{List } L_{j,1}, \text{List } L_{j,2})$ ;
    return  $R(m+1)$  and  $p_k''(m+1)$ ;
    break;
  end
  else
    Update  $S = S - R(m)$ ;
    Remove Sub-carrier which requires highest amount of power
per bit delivered add it's power to total excess power  $W$ ;
    Update  $m \leftarrow m - 1$ ;
  end
end
 $p_k''(m) = \text{Algorithm-3}(p_i^j, \text{List } L_{j,1}, \text{List } L_{j,2})$ ;
return  $R(m)$  and  $p_k''(m)$ 

```

Algorithm 1: Common rate selection and Power redistribution using JPRA-CR

the EVM threshold for rate j and have negative weights w_i^j (these are the items to choose from). We introduce a slack variable for each such item i , x_i^j , which is 1 if the sub-carrier is chosen for placement in the knapsack (assignment of power for rate j), and 0 if not. Then we need to solve the following problem:

$$\text{maximize } \sum_{i=1}^n x_i^j \quad (1)$$

$$\text{subject to } \sum_{i=1}^n x_i^j \cdot w_i^j \leq W^j \quad (2)$$

where $x_i^j \in \{0, 1\}$ for each $i \in \{1, 2, \dots, n\}$ and, W^j is the total excess power available.

However, with JPRA-CR we are not interested in finding the maximum number of sub-carriers that can be assigned additional powers to meet the rate requirement (items that can be put in the knapsack); our goal is to only check if all such sub-carrier requirements can be accommodated. Specifically, we need to check if the value of the objective function is equal to n . If so, rate j can be supported; if not, rates greater than or equal to j cannot be supported. Thus, the receiver only needs to check if upon setting $x_i = 1, \forall i \in \{1, 2, \dots, n\}$, the constraint 2 is satisfied. If so, it moves to the next higher rate; if not, it chooses the immediately lower rate and is done with rate selection. Algorithm 2 summarizes the steps of our common rate selection algorithm.

Algorithm-1 is the main algorithm in JPRA-CR. It iteratively invokes Algorithm 2, to identify a *set* of sub-carriers that support the maximum number of bits to be transmitted per symbol, using a common rate. To begin with it considers all sub-carriers (as described above), and calls Algorithm 2, to find the maximum com-

mon rate that can be supported in this case (say $R(n)$, where n is the number of sub-carriers). The total number of transmitted bits per symbol is then $n \times R(n)$ (after performing power redistribution). It then removes the sub-carrier that requires the highest power per bit delivered (*poorest* sub-carrier) and adds its power to the excess power budget (as discussed); the common aggregate bit-rate with the remaining $(n - 1)$ sub-carriers is now calculated by invoking Algorithm 2 to be $(n - 1) \times R(n - 1)$, where $R(n - 1)$ is the maximum common rate that is now supported. If $(n - 1) \times R(n - 1) < n \times R(n)$, it is easy to see that the algorithm has converged, since no further optimization is possible. Note that $R(n)$ and $R(n - 1)$ are calculated after performing power redistribution using Algorithm 2; thus these are only available as this algorithm is executed iteratively. If not, from among the remaining sub-carriers, the poorest sub-carrier is removed and the process is repeated. The iterations continues until a point is reached, say with the number of sub-carriers m , where $m \times R(m) < (m + 1) \times R(m + 1)$. This implies that the optimal set of sub-carriers are the $(m + 1)$ sub-carriers in the prior to last iteration and, $R(m + 1)$ is the maximum common rate to be used.

```

Input: Sub-carrier  $EVM_i$ , Rate  $j \in S$ , Rate threshold  $EVMTH_j$ ,
          EVM to Power Mapping
Output: Rate  $j$ , List  $L_{j,1}$ , List  $L_{j,2}$ 

foreach Rate  $j$  do
  foreach Sub-carrier  $i$  do
     $diff_i = EVMTH_j - EVM_i$ ;
    Map  $diff_i$  to power  $p_i$ ;
    if  $p_i > 0$  then
       $W =+ p_i$ ;
      Put  $p_i$  in List  $L_{j,1}$ ;
    end
    else
       $w =+ \text{mod}(p_i)$ ;
      Put  $\text{mod}(p_i)$  in List  $L_{j,2}$ ;
    end
  end
  if  $w \leq W$  then
    return Rate  $j$ ,  $L_{j,1}, L_{j,2}$ ;
  end
  else
    return Rate  $j - 1, L_{j-1,1}, L_{j-1,2}$ ;
  end
end

```

Algorithm 2: Rate selection using JPRA-CR

Once the maximum common rate and the sub-carrier set to be used have been determined as above, the receiver performs the power redistribution. It transfers power from the *fading-immune* sub-carriers, to the affected ones. Algorithm 3 presents the steps taken. The considered sub-carriers are divided into two classes (lists). The first list contains the sub-carriers with the excess power, while second list contains sub-carriers that require more power for meeting the EVM threshold for the chosen rate. Note that these lists were already returned by Algorithm 2 previously and are now used by Algorithm 3 to perform power redistribution. Upon executing this algorithm, the receiver first subtracts the excess power from the sub-carriers in List-1. It then adds the extra power as required to each sub-carrier in List-2 (in decreasing order of the required power). If total extra power is more than the required power then left over power is distributed evenly among all the sub-carriers.

Computational Complexity: The run time for JPRA-CR is $O(n^2 \cdot l)$, where n is the number of sub-carriers and l is the number of available data rates. In brief, JPRA-CR (i.e., Algorithm 1) iterates

over the n sub-carriers and executes algorithm 2 in every iteration. The latter includes a nested loop which executes over all the sub-carriers and over all the l transmission rates, thus, requiring $n \cdot l$ running time. Since, algorithm 3 (whose complexity is $O(n)$), is executed outside the loop in Algorithm 1, the time complexity of JPRA-CR is $O(n^2 \cdot l)$.

```

Input: Initial sub-carrier Power  $p'_i$ , List  $L_{j,1}$ , List  $L_{j,2}$ 
Output: Adjusted sub-carrier Power  $p''_i$ 

foreach sub-carrier in List  $L_{j,1}$  do
   $x =+ p_i$ ;
   $p''_i = p'_i - p_i$ ;
end
Sort List  $L_{j,2}$  in decreasing order;
foreach Sub-carrier in List  $L_{j,2}$  do
   $p''_i = p'_i + p_i$ ;
   $y =+ p_i$ ;
end
if  $x > y$  then
   $l = (x - y) / n$ ;
  where  $n$  is total number of sub-carriers;
  foreach sub-carrier in List  $L_{j,1}$  and  $L_{j,2}$  do
    Update  $p''_i =+ l$ ;
    return  $p''_i$ ;
  end
end

```

Algorithm 3: Power Redistribution after selecting rate through JPRA-CR

4.2 Minimizing Transmission Air Time with JPRA-MT

The transmission time of a packet depends on the modulation and coding rate used, Higher modulation schemes lead to shorter packet times. In turn, shorter packet transmission times lead to throughput improvements. Ideally, in order to minimize the packet air time one would select the highest modulation on all the sub-carriers. However, due to frequency selectivity in fading all sub-carriers might be unable to support this high rate.

Adaptive modulation and coding schemes have been proposed to cope with frequency selectivity (e.g., [16]). With such schemes, the bit-rate for each sub-carrier is selected based on its CSI (Channel State Information). However, power redistribution is not performed. Power redistribution helps in maximizing the total number of bits transmitted by the sub-carriers in a symbol duration. Let us consider a toy example with two sub-carriers A and B, initially with default power allocations. A can only support BPSK while B can support QPSK. In total A and B can carry 3 bits since BPSK and QPSK can modulate 1 and 2 bits, respectively. If one completely shuts down A and transfers its power to B, B may be able to use 16QAM with the extra power. Thus, the total number of bits transmitted in a symbol duration is now 4 instead of 3 (a 16QAM symbol maps on to 4 bits). Thus, we gain 1 extra bit per symbol by doing power redistribution. With a large number of sub-carriers, where each sub-carrier has a different EVM and requires a different level of extra energy to support a higher modulation scheme, the gains can be (and actually are) more significant.

Formalizing the problem: The problem of appropriately allocating powers (and thus, transmission bit-rates) to sub-carriers in order to minimize the packet air-time can be formally defined as follows: Lets assume that there are N sub-carriers $x_1, x_2 \dots x_N$, and $R+1$ (one for each rate level and a null rate level for no assignment)

bit-rate levels for each sub-carrier. On each sub-carrier only one specific bit-rate out of the $R+1$ available bit-rates can be used; for each bit-rate $j \in x_i$ there is a corresponding *profit* $r_{i,j}$, which is the bit-rate itself and a weight that corresponds to the power $p_{i,j}$, needed to achieve this bit-rate. We estimate this power using our measurements from the calibration phase. Maximizing the profit, essentially translates to using the highest bit rates on the sub-carriers; as one can easily see, this maximizes the number of bits transmitted per symbol and as a consequence, minimizes the transmission air time. However, the maximization of the profit as above is constrained by the total power budget available (P).

The above problem can be mapped to the well known Multiple Choice Knapsack Problem (MCKP). In MCKP, we are given m classes N_1, N_2, \dots, N_m of items that are to be packed in a knapsack of capacity C . Each item $j \in N_i$ has a profit $a_{i,j}$ and a weight $w_{i,j}$, and the problem is to choose at most one item from each class such that the profit sum is maximized without the weight sum exceeding C .

With our problem, a sub-carrier x_i corresponds to a class, that consists of items that correspond to the $R + 1$ bit rates. The profit is the bit-rate value $r_{i,j}$ and weight is the required power $p_{i,j}$. The total capacity of the knapsack is the total transmission power P , available to the transmitter.

Using this mapping, our problem can be formulated as the following MCKP:

$$\text{Maximize} \quad \sum_{i=1}^N \sum_{j \in x_i} r_{i,j} k_{i,j} \quad (3)$$

$$\text{Subject to} \quad \sum_{i=1}^N \sum_{j \in x_i} p_{i,j} k_{i,j} \leq P \quad (4)$$

$$\sum_{j \in x_i} k_{i,j} = 1, \quad i = 1, \dots, N$$

$$k_{i,j} \in \{0, 1\}, \quad i = 1, \dots, N, \quad j \in x_i$$

Here $k_{i,j}$ is a slack variable for each $r_{i,j}$; if a particular $r_{i,j}$ is selected to be included in the knapsack it is 1 and otherwise 0. The total number of items considered is $n = \sum_{i=1}^N R + 1 = N \cdot (R + 1)$.

Solving the problem: The MCKP problem is NP-hard [11] as it contains the knapsack problem as a special case [11]. However, it has been shown that it can be solved in pseudo-polynomial time [7]. Pisinger [15] proposes an algorithm (MCKNAP) which has been shown to outperform other algorithms for solving the MCKP problem. The algorithm first extends the partitioning algorithm of Balas and Zemel [4] for the original knapsack problem, to the case of MCKP. The partitioning algorithm identifies a set of classes (called the *core*) that include the optimal solution with high probability; this can be thought of as a reduction phase. A dynamic programming algorithm is then applied on the core in order to identify the optimal solution in pseudo linear time. This can be thought as an expanding phase, since it might add classes that do not belong to the core. The resulting solution defines the classes and the corresponding elements that are included in the knapsack (i.e., the elements with $k_{i,j} = 1$). Based on the assigned rates, the total transmit power is allocated among the carriers.

Computational Complexity: The computational complexity of MCKNAP is $O(n + c \sum_{x_i \in c} r_i)$ for a minimal core c ; thus, the algorithm executes in linear time for small cores and pseudo polynomial time for large cores. The proofs are available in Section 7 of [15].

Carrier Frequency	2427MHz
RF Bandwidth	10MHz and 20MHz
Number Of Sub-carriers	64
Modulation Schemes	BPSK, QPSK, 16QAM
Payload Length	1470 bytes
Tx Power	19dBm
MAC Protocol	CSMA/CA

Table 2: MAC/PHY parameters for JPRA implementation

5. PERFORMANCE EVALUATION

In this section, we examine the performance of JPRA via extensive real testbed as well as simulation experiments.

5.1 Testbed Setup and Implementation

Our testbed utilizes the Wireless Open-Access Research Platform (WARP) developed at Rice University. The WARP platform consists of three main components; (a) A Xilinx Virtex-II Pro FPGA, (b) A 2.4/5GHz Radio Board, which supports wideband applications such as OFDM, and (c) A 10/100 Ethernet port, which serves as the interface between the board and the wired Internet. The Xilinx module implements the MAC and PHY layer protocols. MAC protocols can be implemented in C, while the PHY layer protocols are implemented within the FPGA fabric using MATLAB Simulink. The current physical layer design uses an OFDM implementation that is loosely based on the PHY layer of the 802.11a standard. The WARP/MAC/WARPPHY modules provide basic building blocks towards implementing more advanced MAC/PHY protocols. The basic configuration of WARP's OFDM design is given in Table 2.

We conduct our experiments on a six node WARP indoor testbed. We randomly selects three sender/receiver pairs, and send saturated traffic with packets of 1472 bytes. We consider cases where the connections interfere with each other and cases where they don't. Each experiment lasts for 6 minutes and the reported results are the averages of 5 runs unless otherwise stated. Each node is connected to a laptop through an Ethernet switch, which acts as a controller. The controller is also used for analyzing and collecting traces during the experiment.

Implementation: We implement JPRA on top of WARP's standard OFDM design, which runs on the board's FPGA in real time. We run JPRA at the MAC layer to calculate the rates to be used and the power redistribution. The receiver communicates this information to the transmitter through ACKs or NACKs (as mentioned in Section 4). We modify the transmitter to perform power redistribution and rate adaptation per sub-carrier, upon obtaining this information from the receiver. We also implement (for comparison), a standard EVM based rate adaptation scheme, which selects a single rate for all the sub-carriers. This scheme referred to as *Standard rate*, selects the appropriate rate for each packet by comparing the average packet EVM with the threshold values in Table 1. Standard rate adaptation does not perform any kind of power adaptation. Finally, for comparison purposes as well, we also implement FARA as described in [16]. In brief, FARA selects a rate based on the EVM or SNR threshold for each sub-carrier. In our implementation, FARA examines the individual sub-carrier EVM and selects the rate which can be supported. If a carrier cannot satisfy even the basic (minimum) rate it is left unused (no symbols are mapped on to that sub-carrier).

5.2 Experimental Results

Effects of power redistribution on RSSI value: We first want to ensure that a power redistribution with JPRA does not affect key

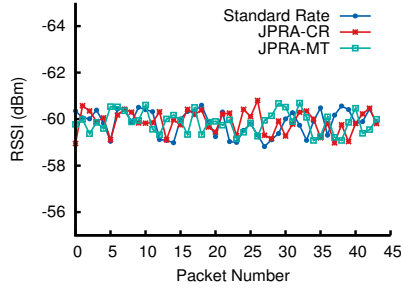


Figure 6: RSSI values per packet




Figure 7: Normalized network throughput with 10MHz channel width.




Figure 8: Normalized network throughput with 20MHz channel width.

network functionalities like carrier sensing. Typically the RSSI value is used in most of the wireless systems to determine the amount of energy on the channel. It is used as an indicator of whether the channel is idle or not. In WiFi, packet detection is based on RSSI values as well. Thus, we measure the RSSI values per packet with and without JPRA. Figure 6 depicts a part of a representative trace that we collected with Standard Rate adaptation, JPRA-CR and JPRA-MT. As one can see the RSSI values are similar in all scenarios, thus leading us to believe that JPRA will leave key higher functions such as carrier sensing or packet detection unaffected. The reason behind this observation can be explained as follows. The RSSI value is essentially the average energy present on the antenna during the reception of the preamble. Since JPRA is not applied on the preamble, we do not observe any significant RSSI variations.

Network throughput: Next, we examine the total network throughput. Figure 7 compares the throughput with JPRA with that using Standard Rate Adaptation and FARA. The throughputs are normalized by the highest throughput value among all observations (highest throughput = 1). This provides an immediate assessment of the relative performance with the schemes that we compare. The channel bandwidth in this experiment was 10MHz. We observe that JPRA-CR outperforms standard rate adaptation by up to 28%. As one may expect, FARA outperforms JPRA-CR, since it opportunistically employs higher rates on a subset of sub-carriers. However, the gains are marginal (< 8%); further, JPRA-CR only requires feedback corresponding to the common rate. JPRA-MT outperforms both standard rate and FARA by, above 42% and 20% respectively. The main source of gain for JPRA-MT in comparison to FARA and Standard rate adaptation, is the per sub-carrier power redistribution. The power redistribution allows the transmission of a higher number of bits per symbol as discussed earlier. Note here that we show the 95% confidence intervals for all results.

Wider channels exhibit a higher degree of frequency selectivity. We also perform the same set of experiments with a 20 MHz channel; this channel width is typical in WiFi systems. With the 20 MHz channel the observed normalized throughput with JPRA-CR is 43% higher as compared with that of the Standard Rate adaptation (Figure 8). However, FARA still outperforms JPRA-CR by about 7%. With JPRA-MT we observed higher gains with the 20MHz channel width. In particular, our results (Figure 8) indicate that JPRA-MT delivers up to 75% more network throughput than Standard Rate Adaptation. It also outperforms FARA by 35%. Thus, *per sub-carrier power adaptation enhances network performance in comparison to schemes without power adaptation*. With the increase in channel width, gains also increase. This becomes especially noteworthy, given the increasing number of technologies that use wide

channels (e.g., 802.11n with 40 MHz channels and white space networking with channels of at least 100MHz).

Stability of rate: One of the main reasons for the design of JPRA-CR is the stability of rate that it can offer. In addition JPRA-CR requires a much simpler transceiver design since it chooses a single rate for a given transmission which can be supported on all or a chosen subset of sub-carriers; therefore a single rate or modulation is applied to all the sub-carriers like standard OFDMA. In contrast FARA and JPRA-MT in the worst case use l (total number of rates) different modulations on the different sub-carriers for a given transmission, which leads to a more complex transceiver design. In particular the transceiver has to implement additional bookkeeping to track the rates on the different sub-carriers with these other schemes.

To reiterate, a single rate is chosen, which all sub-carriers can support. Since some carriers might experience fading while others may not, on average with power redistribution, the maximum common rate that can be supported is not expected to change significantly from one transmission to the next (even though the sub-carriers that experience fading might be different). In order to examine the performance of JPRA-CR with respect to rate stability, we created 10 diverse links in our lab by changing the positions of the sender and the receiver. The sender transmits 1472 byte size packets for 3 minutes and we monitor the throughput on the receive side. We also log the number of rate changes applied by the transmitter. The results are shown in Figure 9. We plot the average number of rate changes over 3 trials with JPRA-CR, Standard Rate, JPRA-MT and FARA. As one can see JPRA-CR resulted in up to 27% fewer rate changes in comparison to Standard Rate Adaptation, it also outperforms the other schemes. JPRA-MT performs up to 8% and 14% better than FARA and Standard Rate Adaptation. With FARA and JPRA-MT, we consider a rate change to have occurred if the transmission bit-rate on any sub-carrier is changed between packet transmissions. We believe this is reasonable since the receiver has to now decode packets with the new rate on the specific sub-carrier. Note here that FARA also performs better than standard rate because it shuts down sub-carriers which are in deep fades, these are the carriers where the fluctuations are most likely. *To summarize, power redistribution helps in stabilizing rate changes since it improves performance on sub-carriers experiencing selective fading (which results in increased rate changes)*.

Power savings: In our next experiments, we examine the possible power saving gains with JPRA-CR. In particular, we seek to answer the following question “Can we use lower total transmission power in order to achieve a target PDR?”. We set a target PDR of 90% and we vary the total transmission power of the sender. On an average over 5 trials we observe that with JPRA-CR we need 4.5dBm less power compared to Standard Rate adaptation in order

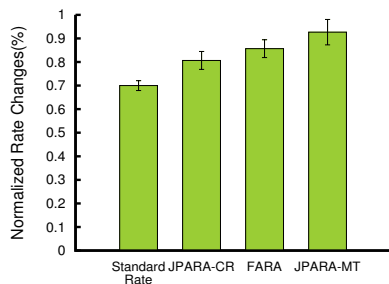


Figure 9: Normalized number of rate changes with 20 Mhz.




Figure 10: Network throughput with packet size 512 bytes.

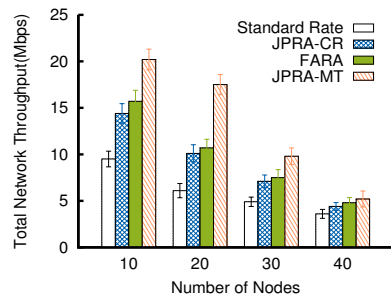


Figure 11: Network throughput with 1472 bytes packet size .

to achieve the preset target delivery ratio. These gains are important, since potentially JPRA-CR can reduce interference; moreover, it can lower energy costs since it requires less transmission power to achieve a certain performance.

Overheads: A receiver needs to send feedback information to the transmitter to perform rate and power reallocation. WARP’s CSMA/CA MAC layer uses standard 802.11 DATA and ACK packet formats, but allows us to piggyback rate and power information on the ACK packets (as with 802.11n ACK packets [3]). Bit-rate information is encoded using 6 bits (up to 64 values). Each sub-carrier power value is a signed real value between 0 and 2.0, which can be represented using 7 bits. When channel remains stable for multiple packets we just send a normal ACK (without any additional information) which tells the receiver to use power and rate values from the previous instance. For 48 data sub-carriers to encode bit-rate and power information we require 48×6 bits and 48×7 bits, respectively. In a worst case scenario, JPRA-MT incurs a 5.2% overhead and JPRA-CR incurs a 2.5% overhead compared to Standard Rate adaptation for a payload size of 1472 bytes. Note that the corresponding overhead increase with FARA is 2.4%. In our testbed the average overhead that we observed for JPRA-MT and JPRA-CR was 4.1% and 1.40%, respectively. The average overhead of FARA was 1.82%, which is more than JPRA-CR since JPRA-CR is more stable and requires less feedback. We wish to point out here that our throughput results, already account for this overhead i.e., the gains are in spite of the overhead .

5.3 Simulations

In order to evaluate our scheme in larger settings we perform simulations in NS-3. In order to have a realistic OFDM PHY layer model, we use a modified version of NS-3 which includes a detailed physical layer model called PhySim-Wifi[14]. PhySim-Wifi includes a physical layer implementation of the OFDM PHY specification for the 5 GHz band as well as a wireless channel emulation. Specifically in our simulations we use NS-3.9 with the PhySimWiFi 1.1 module. Key Physical and MAC layer parameters are given in Table 3 . We further modified PhySim-Wifi in order to implement JPRA-CR, JPRA-MT and FARA. Standard rate adaptation is already implemented in NS-3 but we modified it to use EVM instead of SNR as rate adaptation metric. We implemented these schemes in a similar fashion as we implemented them on our testbed. We repeat the calibration phase in the simulator in order to calculate per sub-carrier EVM, EVM to sub-carrier Power mapping and EVM to rate thresholds. EVM to sub-carrier power mapping follows a similar third degree polynomial curve that we observed on our testbed.

To create frequency selectivity, the sub-carrier gains $|h_n^2(t)|$ are generated based on path loss and fading. The path loss model is

Carrier Frequency	5320MHz
RF Bandwidth	20MHz
Number Of Sub-carriers	64
Modulation Schemes	BPSK, QPSK, 16QAM, 64QAM
Payload Length	512/1472 bytes
Tx Power	20dBm
MAC Protocol	CSMA/CA

Table 3: MAC/PHY parameters used in simulation

characterized by a standard model $|h_{PL}|^2 = K * 1/d^\alpha$ parameterized by $K = 40.14dB$ and the path loss exponent $\alpha = 3.5$. The fading component is implemented based on [20] and parameterized with the values provided in [24] (Table 1 of the reference). The resulting fading gains feature correlation in time and in frequency. The environment is further characterized by an RMS delay spread of 25ns.

Simulation setup: Our simulation topology consists of 40 nodes, arranged in a 8×5 grid. The inter-node distance in the grid is 20m. We consider randomly selected transceiver pairs (also referred to as traffic pairs) to send data packets. We consider 4 sets of experiments while varying the number of active links. In particular, we experiment with 5, 10, 15 and 20 traffic pairs, active simultaneously. Since the topology is fixed, a higher number of traffic pairs implies higher traffic intensity. We use CBR (constant bit rate) to generate packets of size 512 bytes and 1472 bytes. Simulation time for each run is 120 seconds, and our results are the average of 10 of such simulation runs.

Simulation results: Figures 10 and 11 present the results for network throughput with different node densities with each scheme. With JPRA-CR, the total network throughput is 71% higher as compared with Standard rate adaptation, when node density is 20. JPRA-MT improves the total network throughput by upto 185% in comparison to Standard rate and also registers a gain of 64% over FARA. Similar to our experimental results FARA outperformed JPRA-CR by 9% in simulations; again, this is due to per sub-carrier rate adaptation which leads to higher throughput. As the number of competing links increase, the gain for each scheme increases up to a certain traffic intensity. After this *sweet point* the gains are diminished since interference and contention become the dominant impairments.

With increased interference we observe an increasing number of packet collisions. Interference reduces the opportunity to gain from power redistribution since it affects all the sub-carriers.

6. CONCLUSIONS

In this paper, we make a case of adaptively distributing uneven power levels to OFDMA sub-carriers to cope with frequency selectivity in fading. To validate our thesis that this will provide significant throughput benefits, we design JPRA, for jointly selecting the transmission power and bit-rate for the sub-carriers that compose a channel. We design two versions of JPRA: (a) JPRA-CR, a simpler scheme, which while less effective in improving throughput, provides increased stability in the rate in use, simplicity in transceiver design and potential power savings and (b) JPRA-MT which provides significant throughput gains albeit with lower rate stability. We implement both variants on our WARP radio testbed and also perform simulations to showcase its benefits in larger settings. We show that JPRA-MT can provide upto a 35 % increase in throughput compared to the state of the art, FARA scheme that only performs per sub-carrier rate adaptation.

7. ACKNOWLEDGEMENTS

We thank the MobiHoc'12 reviewers and our shepherd, Dr. Jerry Chiang, for feedback on this paper. This work was supported in part by Cisco University grant and Multi-University Research Initiative (MURI) grant W911NF-07-1-0318.

8. REFERENCES

- [1] Onoe rate control.
- [2] Wireless lan medium access control (mac) and physical layer (phy) specifications: High-speed physical layer in the 5 ghz band. *Part 11, Standard ed.,IEEE802.11 Working Group* (1999).
- [3] *Part 11 : Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. IEEE Computer Society, 2009.
- [4] BALAS, E., AND ZEMEL, E. An algorithm for large zero-one knapsack problems. *Operations Research* (1980).
- [5] BHARTIA, A., CHEN, Y.-C., RALLAPALLI, S., AND QIU, L. Harnessing frequency diversity in wi-fi networks. In *Proceedings of the 17th annual international conference on Mobile computing and networking* (2011), MobiCom.
- [6] BICKET, J. C. Bit-rate selection in wireless networks. Tech. rep., Master's thesis, MIT, 2005.
- [7] DUDZINSKI, K., AND WALUKIEWICZ, S. Exact methods for the knapsack-problem and its generalizations. *European Journal Of Operational Research* (1987).
- [8] HOLLAND, G., VAIDYA, N., AND BAHL, P. A rate-adaptive mac protocol for multi-hop wireless networks. In *Proceedings of the 7th annual international conference on Mobile computing and networking* (2001), MobiCom.
- [9] ISHIKAWA, H., FUJII, M., ITAMI, M., AND ITOH, K. Bi-directional ofdm transmission using adaptive modulation that spreads data symbols. In *Power Line Communications and Its Applications, IEEE International Symposium on* (2006).
- [10] JUDD, G., WANG, X., AND STEENKISTE, P. Efficient channel-aware rate adaptation in dynamic environments. In *Proceedings of the 6th international conference on Mobile systems, applications, and services* (2008), MobiSys.
- [11] KELLERER, H., PFERSCHY, U., AND PISINGER, D. *Knapsack problems*. Springer, 2004.
- [12] KITTIPIYAKUL, S., AND JAVIDI, T. Subcarrier allocation in ofdma systems: beyond water-filling. In *Signals, Systems and Computers. Conference Record of the Thirty-Eighth Asilomar Conference on* (2004).
- [13] MCKINLEY, M., REMLEY, K., MYSLINSKI, M., KENNEY, J., SCHREURS, D., AND NAUWELAERS, B. Evm calculation for broadband modulated signals. *64th ARFTG Conf. Dig* (2004).
- [14] MITTAG, J., PAPANASTASIOU, S., HARTENSTEIN, H., AND STROM, E. G. Enabling Accurate Cross-Layer PHY/MAC/NET Simulation Studies of Vehicular Communication Networks. *Proceedings of The IEEE* (2011).
- [15] PISINGER, D. A minimal algorithm for the multiple-choice knapsack problem. *European Journal of Operational Research* (1994).
- [16] RAHUL, H., EDALAT, F., KATABI, D., AND SODINI, C. G. Frequency-aware rate adaptation and mac protocols. In *Proceedings of the 15th annual international conference on Mobile computing and networking* (2009), MobiCom.
- [17] SADEGHI, B., KANODIA, V., SABHARWAL, A., AND KNIGHTLY, E. Opportunistic media access for multirate ad hoc networks. In *Proceedings of the 8th annual international conference on Mobile computing and networking* (2002), MobiCom.
- [18] SEN, S., SANTHAPURI, N., CHOUDHURY, R. R., AND NELAKUDITI, S. Accurate: Constellation based rate estimation in wireless networks. In *NSDI* (2010).
- [19] STEEL, R., AND TORRIE, J. *Principles and procedures of statistics*. 1960.
- [20] WANG, C.-X., PATZOLD, M., AND YAO, Q. Stochastic Modeling and Simulation of Frequency-Correlated Wideband Fading Channels. *IEEE Transactions on Vehicular Technology* (2007).
- [21] WILLINK, T. J., AND WITTKKE, P. H. Optimization and performance evaluation of multicarrier transmission. *IEEE Transactions on Information Theory* (1997).
- [22] WONG, C. Y., CHENG, R. S., D., K. B. L. R., MURCH, R. D., MEMBER, S., AND MEMBER, S. Multiuser ofdm with adaptive subcarrier, bit, and power allocation. *IEEE Journal on Selected Areas of Communications* (1999).
- [23] WONG, S. H. Y., YANG, H., LU, S., AND BHARGHAVAN, V. Robust rate adaptation for 802.11 wireless networks. In *Proceedings of the 12th annual international conference on Mobile computing and networking* (2006), MobiCom.
- [24] ZHAO, X., KIVINEN, J., AND VAINNIKAINEN, P. Tapped delay line channel models at 5.3 GHz in indoor environments. *Composites Part A-applied Science and Manufacturing* (2000).