This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

M-Polar: Channel Allocation for Throughput Maximization in SDR Mesh Networks

Eugene Chai and Kang G. Shin
Real-Time Computing Laboratory
Department of Electrical Engioneering and Computer Science
The University of Michigan, Ann Arbor, MI 48109-2121, U.S.A.
{zontar, kgshin}@eecs.umich.edu

***Abstract*—In traditional wireless networks, nodes use only a single channel per radio interface, thus limiting the overall channel diversity of the network. This restriction is due to the inherent limitations of commercially-available RF devices. With the advent of high-bandwidth software-defined radios (SDRs), we now have the option of assigning multiple contiguous, independent channels to a single wireless interface. This new-found opportunity raises an important question: how do we assign contiguous channels to nodes in order to maximize overall network throughput? This question lies at the often-ignored intersection of single-radio-multi-channel and multi-radio-multi-channel assignment schemes.**

In this paper, we develop a protocol that assigns contiguous channels with the goal of evenly spreading the load across the multiple channels. Neighboring nodes greedily adjust their channel ranges according to channel conditions to achieve an overall pattern of partially-overlapping bandwidths that maximizes the network throughput. The end-result is a network that can dynamically adapt its bandwidth usage to the network load and the conditions of the different channels. The proposed protocol is evaluated with a prototype built upon the USRP as well as with detailed simulation.

I. Introduction and Motivation

Current channel-assignment schemes fall into two main categories: Single-Radio-Multi-Channel (SR-MC) and Multi-Radio-Multi-Channel (MR-MC). The focus of SR-MC is the proper scheduling of a single interface across multiple channels, while that of MR-MC is a suitable approximate solution to the graph coloring problem. With regards to network capacity enhancement, both solutions suffer from a limitation of current wireless hardware: two radios can only communicate if their spectra overlap completely. In this paper, we design and implement a Single-Radio-Contiguous-Channel (SR-CC) channel-allocation scheme, called *M-Polar*, that enables communication with only partial spectrum overlap between radios.

SR-CC enables a node with only a single interface to *simultaneously access multiple contiguous independent channels* by channelizing the accessible spectrum. This extra flexibility offers many significant advantages over other single-radio schemes as follows.

1. Low-Overhead Interference Avoidance: Fig. 1 illustrates the ease with which an SR-CC network can avoid noise or interference in a hypothetical situation. In this example, the network consists of several SR-CC nodes arranged in a straight line, with a vertical box denoting the spectrum occupied by each node. Spatio-frequency regions with high interference power are marked. These high-interference regions may be due to transmissions from non-cooperative networks or primary user activity in the case of cognitive networks. With SR-CC, it is possible to maintain an uninterrupted link between every pair of adjacent nodes by judiciously selecting the spectrum occupied by each node. Hence, communication throughout the network is maintained with no channel switch overhead even in the face of severe interference.

The same situation, if handled by SR-MC, will resemble that shown in Fig. 2. Due to signal interference, the original string topology will now be split into separate factions. A gateway node maintains connectivity between these two factions by switching between the two channels, as shown. Considering that commercial 802.11b cards have channel-switch times ranging from $20\mu s$ to $200\mu s$ [4], a significant communication overhead will be incurred by this gateway node.

2. Dynamic Spectrum Utilization: The simultaneous multi-channel access capability of SR-CC greatly simplifies the on-demand scaling of spectrum utilization. Fig. 1 shows the ease with which a network employing SR-CC schemes can collectively exploit a frequency spectrum larger than that available to any one node. SR-CC can increase or decrease the bandwidth used by the network: a larger bandwidth allows nodes to be "spread out" across the spectrum and will reduce the load on any single channel, but may increase the path length; a smaller bandwidth will increase the node density on each channel and decrease the path length, but the synchronization overhead will be increased. The exact bandwidth use will depend on the channel load and capacity determined by individual nodes spread across the spatio-termporal-frequency domain. More importantly, this bandwidth adjustment can be done with minimal interruption to on-going traffic flows since connections between neighbors are uninterrupted as long as a portion of their spectra overlap.

3. Increased Single-Hop Throughput: SR-CC enables simultaneous access to multiple contiguous channels with only a single radio. Communication with neighboring nodes can be distributed over multiple channels. Hence, for a fixed number of neighbors, the use of SR-CC schemes results in fewer contending neighbors per channel, as compared to SR-MC schemes. Unlike SR-MC, this reduction in channel contention overhead does not incur a corresponding rise in channel-switching overhead since no channel-switches are necessary if neighboring nodes maintain some common spectra between them. This, coupled with the low-overhead interference avoid-

978-1-4244-5837-0/10/$26.00 ©2010 IEEE

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

ance capability of SR-CC, results in a significant throughput improvement.

SR-CC shares many of the advantages of MR-MC, most notably its multiple-channel access capability. However, the hardware requirement of SR-CC is significantly lower, since only a single radio interface is required. The number of channels accessible to a SR-CC node is limited only by its CPU capacity while for an MR-MC node, a separate hardware interface is needed for each channel. The lower hardware and power requirements of SR-CC, coupled with its multi-channel capability, makes it suitable for platforms ranging from cellphones to cellular base stations.

In the rest of this paper, we describe the design and implementation of M-Polar which has the following salient features:

- channelization of available spectrum using only a single radio interface,
- channel adjustment with very minimal interruption to active transmissions between neighbors,
- multi-channel communications using only a single radio without any channel-switching overhead, and
- dynamic adaptation of overall network bandwidth usage to traffic load.

We evaluate M-Polar using a prototype built upon the USRP platform, as well as with detailed simulations.

The paper is organized as follows. Section II introduces the theoretical basis behind our channel adjustment heuristic before describing the adjustment protocol in Section III. This is followed by a description of the PHY design of M-Polar in Section IV. Sections V and VI present the results of real-world evaluation and simulation, respectively. We conclude in Section VIII.

II. Problem Formulation and Theoretical Basis

We consider a network $(\mathcal{N}, \mathcal{C})$ where $\mathcal{N}$ is the set of wireless nodes and $\mathcal{C}$ is the set of available channels. Let r be the number of virtual interfaces of each node $n \in \mathcal{N}$. To avoid trivial solutions, we assume that $|\mathcal{N}| > |\mathcal{C}|$. The ascending ordered sequence of *active channels* assigned to each node is defined as $\mathbf{C_i} := (c_1, c_2, \ldots, c_r)^T$. Any two nodes $i, j \in \mathcal{N}$ are *aligned* if $\mathbf{C_i} = \mathbf{C_j}$.

Given a node $i \in \mathcal{N}$, the upper bound on the capacity of any of its channels $c \in \mathbf{C_i}$ is given by Shannon's capacity formula

$$C_i^c := W \log_2 \left(1 + \frac{P_i^c}{\sigma_c + \sum_{j \in \mathrm{NBR}_i^c} h_{ij} P_j^c}\right)$$

where W is the maximum channel capacity, NBR_i^c is the set of neighbors of node i on channel c, h_{ij} is the channel gain between nodes i and j (i.e., the ratio of the received power at i to the transmitted power from j), P_i is the transmit power of i and σ_c is the noise on channel c. In this paper, we adopt the normalized channel capacity by setting $W = 1$. For simplicity, we assume that the channel is symmetric, i.e., $h_{ij} = h_{ji}$. For notational convenience, we define

$$x_i^c := 1 + \frac{P_i^c}{\sigma_c + \sum_{j \in \mathrm{NBR}_i^c} h_{ij} P_j^c} \quad \text{and} \quad \mathbf{x_i} := (x_i^{c_1}, \ldots, x_i^{c_r})^T$$

where $\mathbf{x_i}$ is the corresponding *channel state vector*

Consider any two connected nodes $i, j \in \mathcal{N}$. The maximum possible capacity available to node i (and similarly, node j) is defined to be the *total capacity* of i (j). The capacity of the overlapping channels is defined to be the *effective capacity* between i and j. This is illustrated in Fig. 3.

The normalized total capacity of the node $i \in \mathcal{N}$ is therefore given by

$$\Gamma(\mathbf{x_i}) := \log_2 \pi(\mathbf{x_i})$$

where $\pi(\mathbf{x_i}) := \prod_{c \in \mathbf{C_i}} x_i^c$. Similarly, the normalized effective capacity of of node i is

$$\Lambda(\mathbf{x_i}) := \log_2 \rho(\mathbf{x_i})$$

where $\rho(\mathbf{x_i}) := \prod_{c \in \mathbf{C_i}} \mathrm{I}\{|\mathrm{NBR}_i^c| \neq 0\} x_i^c$

An M-Polar node has to balance two opposing requirements: (a) maximizing the effective capacity between itelf and its neighbors and (b) probing other channels to discover new transmission opportunities to maximize the total capacity to its neighbors. The effective capacity of an M-Polar node represents its current capacity while the total capacity is an upper bound on its potential capacity.

A capacity increase cannot be achieved if only one of these constraints is considered in isolation during channel adjustment. Fig. 5 illustrates this fact with nodes in a string topology. We only focus on nodes A, B and C, each tuned to channels 2-5 and with neighbor sets {D,B}, {A,C} and {B,E} respectively. We focus on the total and effective capacity of C. If the active channels of these nodes are arranged as shown in Fig. 5(c), the effective capacity of C will increase, since the total number of nodes on channel 5 has decreased. The total capacity of C will also increase since node C can now gain access to channel 6 with only 3 other nodes while giving up access to channel 2 which has 5 other nodes. In order to attain this capacity increase, the effective capacity of C must first be decreased by shifting its active channel set to cover channels 3 to 6 as shown in Fig. 5(b). If channel assignments are made to only maximize the effective capacity, then this adjustment step will not be taken, since it decreases the effective capacity of C in the short term. It follows that the improved channel assignment cannot be reached.

On the other hand, if channel assignments are made only to maximize the total capacity, the resultant active channels of A, B and C will be that shown in Fig. 6. In this case, the total capacity of each node will be maximized at the expense of its effective capacity, resulting in a significant drop in overall network capacity.

Hence, we specify the channel adjustment policy of M-Polar as an optimization problem:

$$\arg\max_{\mathbf{x}} \sum_{i \in \mathcal{N}} \Lambda(\mathbf{x_i}) \quad \text{s.t.} \quad \begin{array}{l} \forall i \in \mathcal{N}, \quad \Gamma(\mathbf{x_i}) \geq 0 \\ \mathbf{x} = [\mathbf{x_1}, \ldots, x_{|\mathcal{N}|}]^T. \end{array} \tag{1}$$

The associated Lagrangian is given by

$$L(\mathbf{x}, \lambda) = \sum_{i \in \mathcal{N}} \Lambda(\mathbf{x_i}) + \sum_{i \in \mathcal{N}} \lambda_i \cdot \Gamma(\mathbf{x_i}) \tag{2}$$

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

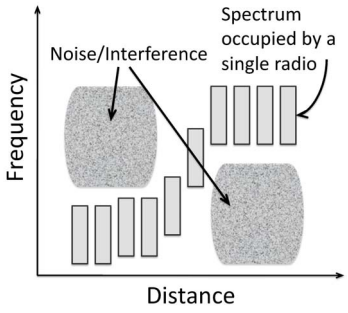

Fig. 1. Interference avoidance with SR-CC. By judiciously selecting the spectrum of each node, the string topology can maintain overall network connectivity with low overhead in the presence of crippling interference

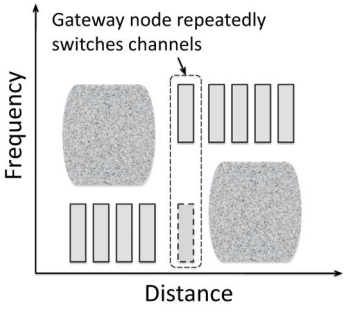

Fig. 2. Interference avoidance with SR-MC. Since complete spectrum overlap is necessary for pair-wise communication, the gateway node has to constantly switch between the two active spectra

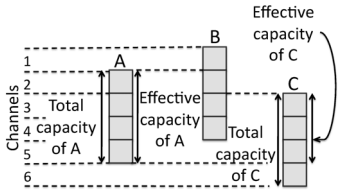

Fig. 3. Example of an assignment of contiguous channels to nodes A, B and C. The total capacity of A is equal to its effective capacity while the total capacity of C is greater than its effective capacity.

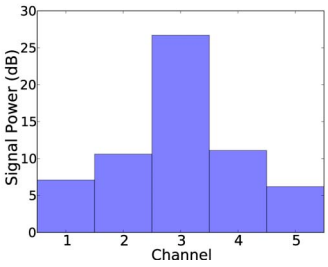

Fig. 4. Signal power measured on each channel in response to transmissions on channel 3

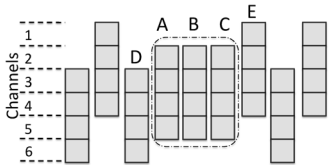

(a) Nodes A, B and C initially occupy channels 2 to 5. The neighbor sets of A, B and C are {D,B}, {A,C} and {B,E} respectively.

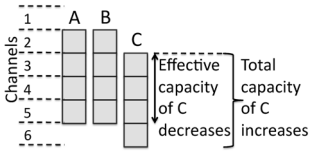

(b) Active channel set of C is shifted to channels 3 to 6. The total capacity of C increases while its effective capacity decreases

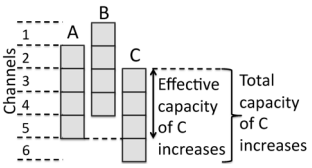

(c) Active channel set of B is shifted to channels 1 to 4. The effective capacity of C is now increased beyond its initial effective capacity

Fig. 5. Channel adjustment example. After the active channels of nodes B and C have been adjusted, the effective capacity of C is increased.

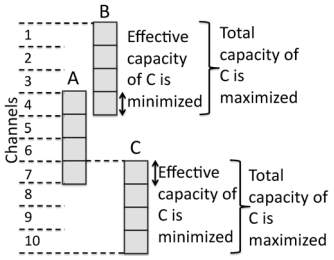

Fig. 6. Channel assignment if only the total capacity is maximized.

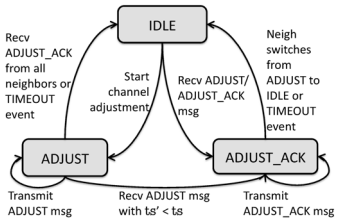

Fig. 7. State-transition diagram for the channel adjustment protocol

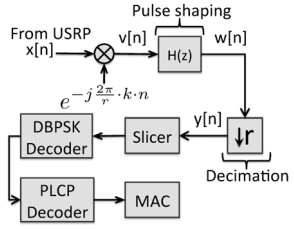

Fig. 8. Processing path for a single channel.

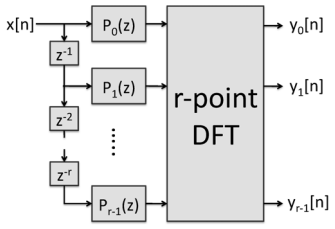

Fig. 9. Optimized polyphase DFT filterbank used in M-Polar PHY. $P_k(z)$ is the k^{th} polyphase decomposed filter of $H(z)$

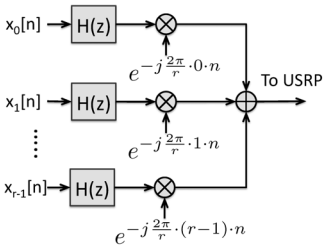

Fig. 10. M-Polar PHY transmit filter design.

along with the partial differential w.r.t. $\mathbf{x_i}$

$$\frac{\partial L}{\partial \mathbf{x_i}} = \frac{\rho'(\mathbf{x_i})}{\rho(\mathbf{x_i})} + \lambda_i \frac{\pi'(\mathbf{x_i})}{\pi(\mathbf{x_i})} \quad (3)$$

From Eq. (3), we can perform a greedy approach to maximizing the effective capacity of an M-Polar by ensuring that in each adjustment step, the normalized decrease [increase] in effective capacity is within a constant factor of the normalized increase [decrease] in total capacity. λ is the weighing factor that defines the relative importance of the effective and total capacities. This observation serves as a basis for the channel adjustment heuristic that will be described in the next section.

III. M-Polar Channel Adjustment Protocol

We now present a distributed, iterative channel adjustment protocol based on (3). M-Polar adopts a greedy iterative approach to maximizing the Lagrangian in (2) with (3) as the cost function. In each iteration, M-Polar constructs several *pseudo-active channel lists*, each representing a potential active channel list. The pseudo-active channel list that results in the greatest increase in the Lagrangian (2) is chosen as the next active channel set. Table I lists the variables used by the channel adjustment protocol.

TABLE I
Variables used by the channel adjustment protocol

Variable	Description
NBR^c	List of neighbors on channel c
t_{neigh}	Neighbor list timeout
t_{chan}	Channel adjustment timeout
s_c	SNR measured on channel $c \in \mathbf{C_i}$
n_{SNR}	No. of SNR measurements exchanged in a single data frame
$\mathbf{C_i}$	List of active channels of node i
$\mathbf{S_i}$	SNR of all channels from the point of view of node i
l_{thresh}	Adjustment threshold

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

A. Neighbor Discovery and Management

Before M-Polar can compute $\rho(\cdot)$ and $\pi(\cdot)$, the set of neighbors NBR^c on each of its channels must be known. Neighbor discovery is achieved via overhearing. For each accessible channel c, M-Polar maintains a neighbor list NBR^c. Each entry $(i,t) \in \mathrm{NBR}^c$ contains both the ID of the neighbor i and the most recent time t when a frame from neighbor i was received. Entries older than t_{neigh} are purged.

For each non-active channel, $c \in \mathcal{C} \setminus \mathbf{C_i}$, M-Polar tracks the average SNR from all its neighbors that are active on c. These SNR measurements are used if the node does not have any corresponding local measurements (e.g. it has never accessed some channel and thus does not know the state of that channel). Any inaccuracy in the neighbors' SNR values will be quickly negated by subsequent channel adjustments.

Each M-Polar node attaches the most recent SNR measurements from up to n_{SNR} active channels onto each transmitted frame. The value of n_{SNR} is set to achieve a balance between the extra overhead and speed of information exchange. In our implementation, we set $n_{\mathrm{SNR}} = 5$. The receiving node retrieves these n_{SNR} values and updates the average SNR information of all its non-active channels The SNR of active channels are updated via direct channel measurements by M-Polar.

B. Channel Adjustment Algorithm

Given any active channel list $\mathbf{C_i}$ of some node i, the channel state vector is $\mathbf{x_i} = (x_i^{c_1}, \ldots, x_i^{c_r})$ with $x_i^{c_k} \in \mathbf{S_i}$ and $c_k \in \mathbf{C_i}$. We also define the *pseudo*-active channel list $\mathbf{D_i}$ as the list of channels where $0 < |\mathbf{D_i} \cap \mathbf{C_i}| < |\mathbf{C_i}|$ with the corresponding *estimated* channel state vector $\mathbf{y_i}$ where some elements $y^c \in \mathbf{y_i}$ are the average SNR of non-active channels. $\mathbf{y_i}$ represents the estimated channel state of the potential active channels to which M-Polar can shift.

For each pseudo-active channel list, we approximate the gradient of the Lagrangian in (3) using

$$\pi'(\mathbf{x_i}) = \frac{(\pi(\mathbf{y_i}) - \pi(\mathbf{x_i}))}{\|\mathbf{y_i} - \mathbf{x_i}\|},\ \rho'(\mathbf{x_i}) = \frac{(\rho(\mathbf{y_i}) - \rho(\mathbf{x_i}))}{\|\mathbf{y_i} - \mathbf{x_i}\|}$$

The pseudo-active channel list that has the largest positive gradient will be selected as the new active channel list. The pseudo-code detailing this adjustment algorithm is shown in Algorithm 1.

For each pseudo-active channel list, M-Polar can choose to shift the current list of active channels by up to $\pm(r-1)$ channels. In lines 6-13, M-Polar computes the channel state vector for the current pseudo-active channel list. If the state of the channel is available in the cache (either via direct measurements or information exchange), then that value is used. Otherwise, the average of the current channel state is used.

With this new channel state vector, the approximate value of (3) is computed and stored in a list (lines 14-19). Pseudo-active channel lists with negative gradients are ignored since they represent a decrease in the effective and total capacity.

Finally, in lines 21-27, M-Polar selects the pseudo-active channel list with the highest Lagrangian gradient and returns its corresponding channel shift if this gradient exceeds a predefined threshold. The use of this threshold reduces the possibility of channel adjustments due to spurious bursts of noise. If no pseudo-active channel list can increase the overall effective and total capacity, then the current active channel list is not changed.

Algorithm 1 Single Step Adjustment

```
Require: x_i, C_i, S_i
 1: r ← |C_i|
 2: L ← {}
 3: x̄ ← Average(x_i)
 4: for d = −r + 1 to r − 1 do
 5:    y ← ()
 6:    for c ∈ C_i do
 7:       s ← GetChannelState(S_i, c + d)
 8:       if s ≠ NULL then
 9:          y ← y ∪ {s}
10:       else
11:          y ← y ∪ {x̄}
12:       end if
13:    end for
14:    α ← (π(y) − π(x_i))/(‖y − x‖ · π(x_i))
15:    β ← (ρ(y_i) − ρ(x_i))/(‖y_i − x_i‖ · ρ(x_i)
16:    l ← β + λ · α
17:    if l > 0 then
18:       L ← L ∪ {(l, d)}
19:    end if
20: end for
21: if |L| > 0 then
22:    (l′, d′) ∈ L s.t. ∀(l″, d″) ∈ L, l′ ≥ l″
23:    if l′ > l_thresh then
24:       return d′
25:    end if
26: end if
27: return 0
```

C. Full Protocol Description

M-Polar synchronizes its channel adjustment invocation with neighboring nodes to ensure that (a) no two adjacent nodes adjust their channels simultaneously, and (b) no channel adjustment results in a disconnection of any neighbor. In this subsection, we describe the distributed protocol used by M-Polar to achieve an contiguous channel assignment with maximal capacity.

During channel adjustment, M-Polar can be in any of three states: IDLE, ADJUST and ADJUST_ACK. Fig. 7 summarizes the possible transitions between these states. The relevant fields of the frame header pertaining to these state messages are: `msg`, `timestamp` and `id`. The `msg` field holds the state, i.e., IDLE, ADJUST_ACK or ADJUST, while the `timestamp` field contains the time at which the M-Polar node transitions to the state specified in `msg`. The `id` field holds the identifier of the node associated with the `msg` state. This identifier may not necessarily belong to the node from which the frame originated. M-Polar can forward state information from neighboring nodes to accelerate the propagation of synchronization information among neighbors.

We refer to M-Polar nodes in the ADJUST state as *active* nodes. Nodes in all other states are considered *inactive*.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

1) IDLE State: When M-Polar is in the IDLE state, it only monitors the channel (in addition to the normal frame transmission and reception). From this state, it can transition into either ADJUST or ADJUST_ACK states. If M-Polar receives an ADJUST message from any neighboring node, it transitions to the ADJUST_ACK state and transmits an acknowledgement of its state change. Alternatively, M-Polar can switch to the ADJUST state to prepare to shift its own channels. Channel shifting can then only occur after all neighbors acknowledge their transitions to ADJUST_ACK.

2) ADJUST_ACK State: M-Polar enters this state only when it receives an ADJUST state message from a one-hop neighbor. While in this state, M-Polar does not adjust its active channel list. M-Polar sets the `msg`, `timestamp` and `id` fields of all subsequent transmitted frames to ADJUST_ACK, the time at which the transition occurs and the ID of the node to which it is acknowledging, respectively. This repeated state broadcast increases the speed of state propagation among neighbors. An M-Polar in this state will transition to the IDLE state under either of two conditions: an IDLE message is received from the active node or the active node switches to the ADJUST_ACK state. The former occurs when the neighboring node completes its channel adjustment. The latter indicates a resynchronization event: two neighboring M-Polar nodes are active and the active node with the later `timestamp` value transitions to the ADJUST_ACK state. This can occur if, for example, ADJUST messages are lost due to frame corruption.

3) ADJUST State: M-Polar transitions to this state just before it shifts its active channels. M-Polar sets the `msg`, `timestamp` and `id` fields to ADJUST, the time at which this transition occurred and its own identifier respectively. While in this state, M-Polar transmits the ADJUST message in each frame header and waits until an ADJUST_ACK message is received from each one-hop neighbor. At any point, if an ADJUST message is received from some neighbor with an earlier `timestamp` value, M-Polar transitions to the ADJUST_ACK state. If all ADJUST_ACK messages have not been received after t_{chan}, M-Polar aborts the channel adjustment attempt and transitions to the IDLE state. This is done to avoid deadlocks due to channel desynchronizations.

IV. M-Polar Physical Interface and Design

Conventional PHY implementations utilize all accessible bandwidth on an interface for a single broadcast channel. With SDRs, it is possible to split the accessible bandwidth into multiple contiguous channels. This is especially useful since the SDR can often access a larger bandwidth than is required by state-of-the-art wireless protocols. For example, each GSM channel is 200kHz wide while a USRP can sample a channel of up to 8Mhz, enabling up to 40 channels to be accessed simultaneously.

In this section, we describe the multi-channel design of M-Polar, as well as several key challenges that must be addressed before multi-channel communication is practical.

A. Receiver & Transmitter Design

Fig. 8 shows the processing path used for a single channel in M-Polar. Incoming samples for the channel is first modulated to obtain a baseband signal. A digital pulse-shaping filter $H(z)$ is then applied to the stream before it is decimated by a factor of r. One key requirement of this pulse-shaping filter is its ability to minimize inter-symbol interference. This is particularly important to M-Polar since adjacent channels often experience cross and co-channel interference. Here, we make use of the root raised cosine (RRC) filter. The slicer partitions the stream of samples into Differential Binary Phase Shift Key-ed (DBPSK) symbols, and sums the energy of all samples in each symbol to obtain the receive signal strength. The DBPSK decoder then converts the phase shifts between adjacent symbols into binary bits before feeding them into the PLCP decoder. The start of the frame is demarcated by a known `sync` field in the PLCP header. If such a field is detected, the rest of the frame is decoded and sent to the MAC layer for further processing.

This design incurs too high a processing overhead to be practical for decoding multiple channels. Instead an equivalent, optimized form of the receiver based on uniformly modulated DFT filter-banks as shown in Fig. 9 is used instead.

The transmitter is necessarily designed to enable independent transmissions on multiple contiguous channels. Each channel has a separate processing path which is a simple reciprocal of the receiver design. However, given the asynchronous and bursty nature of transmissions across separate channels, the non-optimized channel filterbank as shown in Fig. 10 can be used with minimal overhead.

B. Dependent Channel Control

The use of a single RF transceiver limits M-Polar to be either transmitting or receiving on all channels at the same time. For example, consider an M-Polar node that uses 4 channels c_1, c_2, c_3 and c_4. Assume that M-Polar detects a transmission only on c_2 and c_4 while it has frames queued for transmission on c_1 and c_3. M-Polar can determine the intended destination of the transmissions only after entire frames on c_2 and c_4 are received. If the frames are destined for a neighbor, then M-Polar will have missed a transmission opportunity. However, without further information, M-Polar will have to adopt this conservative approach to avoid missing frames destined for itself.

M-Polar addresses this issue by using *cut-through MAC header decoding*. On any particular channel c_i, as soon as the MAC header is successfully decoded, M-Polar updates its channel state information before the rest of the packet is completely received. If the frame is destined for the local node, the RF transceiver is held in receive mode to ensure proper frame reception. Otherwise, the transmission on c_i can be safely ignored. Since the size of the header is only a small fraction of the overall frame size, channel-state decisions can be made up to an order-of-magnitude earlier with cut-through header decoding.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

C. Accurate SNR Measurements

M-Polar relies on accurate channel SNR measurements for correct operation. One difficulty with such a metric is the ambiguity between signal and noise components given only aggregate channel energy measurements. There exists a large body of research on SNR estimation. For example, the methods in [6] and [10] can be used for AWGN and time-varying fading channels, respectively.

M-Polar uses a *data-assisted* method for SNR measurements. This is based on the principle that any transmission with a decodable PLCP `sync` field carries a valid frame and should not be treated as noise. The `sync` field is decoded before the rest of the PLCP and MAC headers, allowing M-Polar to quickly detect the presence of a signal, even if the rest of the frame is subsequently corrupted and dropped. In the event that the `sync` field is corrupted, M-Polar assumes that noise is being received.

V. Operational Verification via Implementation & Experimentation

In this section, we demonstrate the operational behavior of M-Polar using two USRP nodes placed 1.5m apart in an indoor environment. M-Polar is implemented using a modified version of GNURadio 3.1.3, and is tested on the USRP platform. CSMA is used as the channel access protocol. We compare the throughput achieved by M-Polar to the throughput achieved without any form of channel adjustments.

A. Cross-Channel Interference

We configure two USRP nodes, each using five 200kHz channels without M-Polar, as a sender-receiver pair. Frames are transmitted on only channel 3 (the center channel) and the signal power is measured on all five channels. Fig. 4 shows the signal power measured at the receiver.

The received power on channel 3 is 27dB while the power on adjacent channels drop off sharply to reach only 7dB on channels 1 and 5. If also transmit on channel 1, we can expect the SNR on channel 1 (assuming the same noise floor for all channels) to be 20. Hence, by transmitting on one of the channel pairs (1,3), (2,4) or (3,5), the interfering signal power will be kept to a sufficiently low level.

B. Throughput Without M-Polar

We then measure the throughput between two nodes without M-Polar. Each node has five 200kHz channels and throughput is measured for approximately 4000s at three different frequency ranges: 2.4308 to 2.4318GHz, 2.4312 to 2.4322GHz and 2.4316 to 2.4326GHz. The per-minute and the average throughput over the entire measurement interval is shown in Fig. 12. The averaging is done using a weighted moving average filter with the weight set to 0.1.

Observe that the throughput at overlapping frequency ranges can fluctuate significantly. Furthermore, signal fading at different frequency ranges occur at different time periods. At 2.4308 to 2.4318GHz (Fig. 12(a)), average throughput drops below 150bytes/s 1233s and 2400s. This can also be seen in Fig. 12(b) between 2735s to 3250s and 3310s to 5100s. It is clear that signal fading can persist for an extended period of time. Hence, when one frequency range is experiencing a persistent signal fade, it is important to quickly switch to a different, better performing range of frequencies

C. Throughput with M-Polar

The previous experiment is repeated using M-Polar over five 200kHz channels. The initial active channel list is tuned to the frequency range spanning 2.4308 to 2.4318GHz and M-Polar adjusts the active channel list at most once a minute to limit the rate of channel adjustments. The results of this experiment are shown in Fig. 11. The triangular markers in Fig. 11 denote the instances where M-Polar performed a channel adjustment step.

At the start of the experiment, the channel state cache only contains SNR information on the initial 5 channels. For each of those unknown channels, the SNR is taken to be equal to the mean SNR of these initial channels. In this case, this is an over-estimate of the actual SNR, resulting in a series of channel adjustments (up to 970s) as M-Polar probes the state of the other channels. The transmission between the two nodes is never broken due to the maintainence of at least one common channel.

At 970s, M-Polar adjusts the active channels to span 2.4312 to 2.4322GHz. At this point, M-Polar is not able to find an adjustment that will result in a increase in the total or effective capacity, and does adjust the active channels any further until 1890s. At this time, all SNR entries of non-active channels have expired and have been purged from the SNR cache. M-Polar resumes its channel probing process and makes an adjustment into a set of channels experiencing prolonged signal fading. M-Polar continues to be stuck in this local minima until 2600s. Occasional sub-optimal decisions such as this one is unavoidable because we are only using two nodes. M-Polar thus loses the benefits of SNR information that would otherwise be exchanged with other non-aligned nodes.

At 2600s, M-Polar escapes from this local minima and manages to achieve a throughput of approximately 220 bytes/s for the rest of the experiment. The remaining channel adjustments are made in response to temporal fading events.

Note that the mean throughput that M-Polar settles at after 2600s (220 bytes/s) is a 46% increase from the initial throughput (153 bytes/s). This throughput increase is not possible without careful channel adjustments, as is evident from a comparison between Figs. 12 and 11.

VI. Simulation

We evaluated the behavior of M-Polar in large networks with the aid of detailed simulation, and describe the results of an NS-2.33 implementation of M-Polar on a network of 25 nodes arranged in a grid.

A. Setup

Table II shows the parameters used in the simulation. The transmission parameters are so chosen that the SNR values at

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

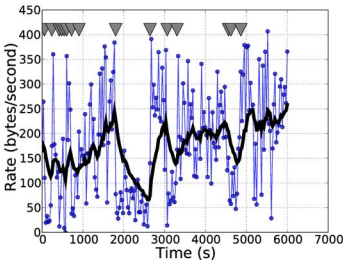

Fig. 11. Achievable throughput with M-Polar. The triangular markers denote the time instances at which channel adjustment are carried out.

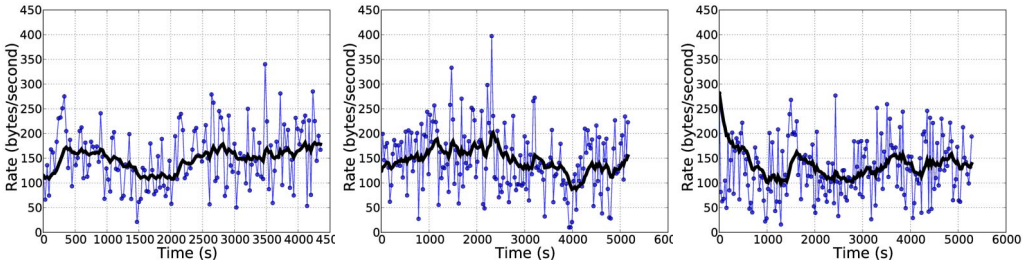

(a) 2.4308 to 2.4318GHz (b) 2.4312 to 2.4322GHz (c) 2.4316 to 2.4326GHz

Fig. 12. Achievable throughput at three different channel ranges. Each M-Polar node uses five 200kHz channels

TABLE II
SIMULATION PARAMETERS

Parameter	Value
No. of nodes	25
Node placement	Grid
Node separation	50m
Noise floor	1.559×10^{-11}
Transmit power	0.78
No. of channels per node	4
Total no. of channels	30 (indexed from 0-29)
Noise model	Zero-mean Gaussian
Simulation duration	4000 seconds

each M-Polar node is similar to those obtained via real-world experiments.

Here, we study the results of three experimental setups using M-Polar. Each simulation is repeated 20 times and we present typical results of these runs. The first has the same channel noise on all 30 channels. This will demonstrate the ability of M-Polar to adapt channel assignments to the perceived load on each channel. In the second experiment, we adjust the channel noise such that channels with higher indices have greater noise power. This will demonstrate the ability of M-Polar to adapt to noisy channels, in addition to network load. In our third simulation study, each node only uses a single channel with bandwidth equal to the sum of the 30 channels used in the previous two simulations. All other parameters in this simulation follow those in Table II. This simulation will show that network throughput under contiguous channel assignment by M-Polar is greater than that of network with only one high bandwidth channel.

B. Metrics

The following three metrics are used to measure the performance of M-Polar.

Per-hop throughput measures the improvement of single hop throughput with M-Polar, between every pair of one-hop neighbors. Since M-Polar performs link-level channel assignment, this metric directly measures the impact of contiguous channel assignment by M-Polar.

Path length shows the overall change in path length due to channel adjustments. We assume an ideal routing algorithm: given any two nodes, the path between them consists of a series of one-hop links that form the shortest route between them. Recall that M-Polar uses frame overhearing to determine its set of one-hop neighbors. Channel fading can cause intermittent periods of high packet loss, resulting in the expiration of neighbor information before channel adjustment. Hence, it is possible for path lengths to change. The aggregate path length statistics presented do not include one-hop neighbors to avoid an unneccesary bias.

End-to-end throughput measures the overall effect of channel assignment. This throughput is measured over the shortest path between any two nodes. An increase in per-hop throughput does not necessarily increase the end-to-end throughput, since the per-hop improvement may come at the expense of longer path lengths, higher loss rate, etc., due to channel adjustments. Similarly, we ignore any one-hop neighbors to avoid any bias.

C. Grid Topology, Identical Channel Noise

Fig. 13 shows the channel placement of each node at the end of the simulation. Each node is denoted by a (`id`, `channel`) pair, where `id` is the node identifier and `channel` indicates the highest channel ID of the active channel list of node `id`. The circles centered at each node are drawn with a size proportional to the value of `channel`: the greater channel index, the larger the diameter of the circle.

The M-Polar nodes arrange themselves into a "channel gradient" with nodes 4 and 9 occupying the lowest channels. In order to understand this resulting channel assignment, consider any node i in Fig. 13 with active channel list $\mathbf{c_i}$. Once node i begins transmission on its active channel list, the signal power on these channels increase. Nodes two-hops and farther away (i.e. non-neighbors) from i on the same channels c will thus experience a decrease in total capacity due to the increased noise power. These nodes can increase their total capacity by reducing the number of overlapping active channels with $\mathbf{c_i}$. This process is repeated as more nodes fix their own active channel lists, thus achieving the overall channel gradient.

Fig. 16 shows the per-hop, end-to-end throughput and the path length achieved during the simulation of M-Polar. Fig. 16(a) shows that, over a series of channel adjustments by M-Polar, the median per-hop throughput increases by 97% while even the 25^{th} percentile sees an 18% increase. This highlights the importance of reducing the channel contention overhead while increasing throughput fairness by reducing the number of nodes on each channel.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

While M-Polar improves fairness between one-hop neighbors on the *same* channel, it makes no such guarantees for neighbors that are farther away. Despite this limitation, Fig. 16(b) shows that M-Polar increases the end-to-end throughput. The median end-to-end throughput increases by 21% while the upper and lower quartiles increase by 67% and 12%, respectively.

Fig. 16(c) shows that channel adjustment by M-Polar does increase the length of the end-to-end paths. The median path length of all pairs of nodes increases from 2 to 4.1 while those of the upper and lower quartiles increase from 2 to 6 and 3, respectively. However, the increase in per-hop throughput and fairness is generally sufficient to overcome the effects of path length inflation.

D. Grid Topology, Non-Identical Channel Noise

For this simulation, the noise σ_i of channel i is set according to $\sigma_i = 2i \times \sigma_0$, where σ_0 is the channel noise used in the previous simulation (Section VI-C). Channels with larger indices will have higher noise power.

Fig. 14 shows the channel assignment at the end of the simulation. Here, we again see a similar gradient of channels assigned to the nodes by M-Polar. However, only a smaller range of channels is utilized by M-Polar, as seen in Fig. 15. Nodes that utilize higher-index active channels will experience greater interference and thus lower throughput than nodes with low-index active channels. Hence, to maximize effective and total capacity, M-Polar skews assigned channels towards those with smaller indices. This will decrease network capacity due to increased channel load and contention overhead. This is a case where M-Polar adapts channel use to the capacity of and load on each channel.

In spite of this, M-Polar can still increase the median per-hop throughput (Fig. 17(a)) by 48%. The upper quartile per-hop throughput increases by 118% while the lower quartile remains largely unchanged. It is no surprise that the variance in per-hop throughput here is greater than in the case with identical channel noises. Since the channels assigned to any node are necessarily contiguous, it is difficult for M-Polar to active perfect load balancing across all non-identical channels.

Because the range of active channels is now smaller, the path length inflation is lower as well. Fig. 17(c) shows that the median path length increases from 2 to 4 hops by the end of the simulation while the upper and lower quartiles increase from 2 to 6.1 hops and 3 hops, respectively.

E. Single High-Bandwidth Channel

Fig. 18 shows the performance of a network where each node is equiped with only a single high bandwidth channel: each node is now assigned a bandwidth equivalent to 7.5 times that of the M-Polar nodes in previous experiments. In spite of this, the achievable per-hop and end-to-end throughput of the non-M-Polar nodes are lower than those with M-Polar. Figs. 18 and 17 illustrate this fact: the median per-hop throughput without M-Polar is at under 200 KB/s while with M-Polar under non-identical channel noise achieves 300KB/s—a 50% difference. This relation can also be seen, albeit to a smaller degree, when we compare the median end-to-end throughput—118KB/s without M-Polar vs. 128KB/s with M-Polar. However, the absence of M-Polar reduces the path-length inflation: any increase in path length is mainly due to losses of update packets to channel fluctuations.

This observed behavior is the result of excessive channel contention overhead. Since only a single channel is used, the interfering set of neighbors of any node is significantly higher than when nodes are distributed over several contiguous channels. The increased channel capacity along with reduction of cross-channel interference and channel-switch overhead is mitigated by this increase in channel contention overhead. Furthermore, as discussed previously, the large set of interfering neighbors reduces the overall throughput fairness.

VII. Related Work

The channel assignment problem is known to be NP-Hard and a multitude of methods have been proposed to approximate the optimal assignment. Alicherry *et al.* [1] take a theoretical approach that uses a relaxed LP formulation to compute an optimal route and its associated channel assignment. Another theoretical proposal made by Xing *et al.* [11] make use of super-imposed codes. To minimize the channel switching overhead in single-radio networks, Vedantham *et al.* [9] propose a channel assignment scheme based on components (i.e., largest connected subgraphs). This minimizes the number of channel switches required for any flow, thus increasing overall throughput. Other notable channel assignment protocols include MMAC [8], SSCH [2], McMAC [3] and a traffic-aware assignment scheme [7].

The spectrum occupied by multiple channels can also be partitioned into time-spectrum blocks and assigned to nodes on an as-needed basis. CMAC [12] is one example of this. Here, neighboring nodes reserve a spectrum block for some period of time (i.e., a time-spectrum block) for data exchange. Moscibroda *et al.* [5] describe another scheme that assigns spectrum blocks of increasing widths to APs with larger numbers of connected clients.

VIII. Conclusion

The growing availability of software-defined radios opens the opportunity to many novel channel assignment schemes. In this paper, we have developed a working prototype of M-Polar and have shown that the achievable throughput by SDRs using M-Polar can even be greater than traditional single-radio-single-channel nodes with 7.5 times the bandwidth of the SDRs.

An exciting corollary from our results is that we do not require large-bandwidth SDRs in order to achieve high throughput. Instead, careful management of the spectrum using the physical layer control knobs made available by the SDR is more important.

Acknowledgements

The work described in this paper was supported in part by the NSF under Grant CNS-0721529

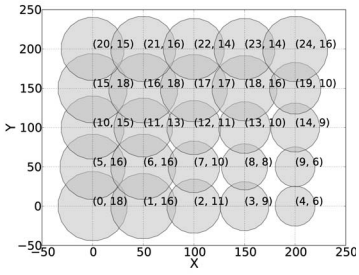

Fig. 13. M-Polar channel assignment with identical channel noise. Each node is tagged with a (id, channel) pair. The diameter of the circle centered at each node is proportional to channel.

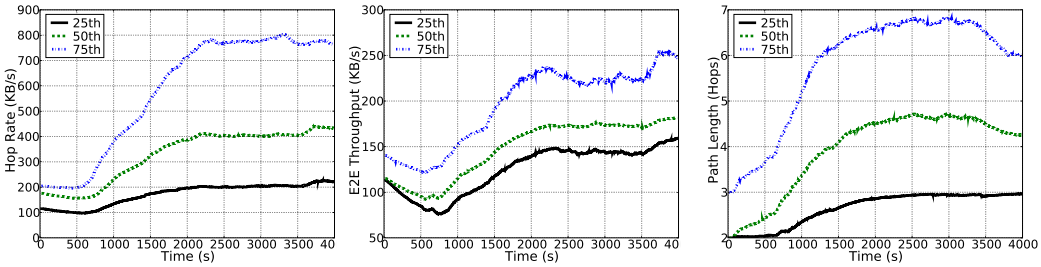

(a) Per-hop throughput (b) End-to-end throughput (c) Path length

Fig. 16. 25^{th}, 50^{th} and 75^{th} percentile performance of M-Polar with identical channel noise.

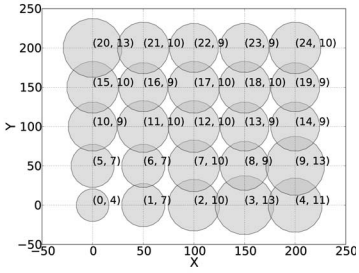

Fig. 14. M-Polar channel assignment with non-identical channel noise. Each node is tagged with a (id, channel) pair. The diameter of the circle centered at each node is proportional to channel.

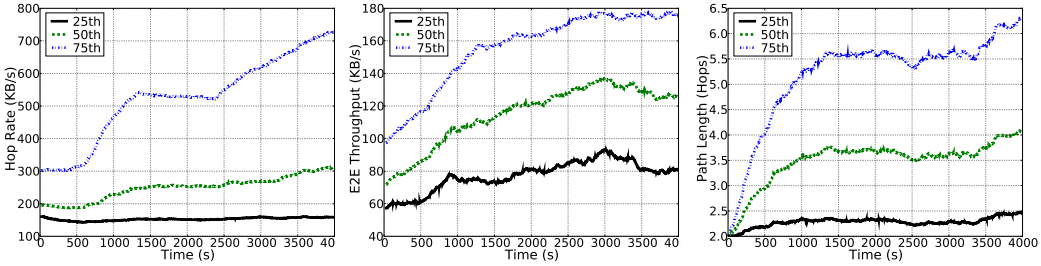

(a) Per-hop throughput (b) End-to-end throughput (c) Path length

Fig. 17. 25^{th}, 50^{th} and 75^{th} percentile performance of M-Polar with non-identical channel noise

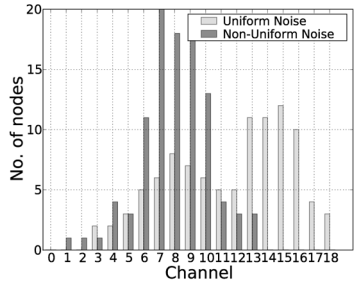

Fig. 15. Number of nodes on channels 1 to 18. No nodes are on any channel not shown here.

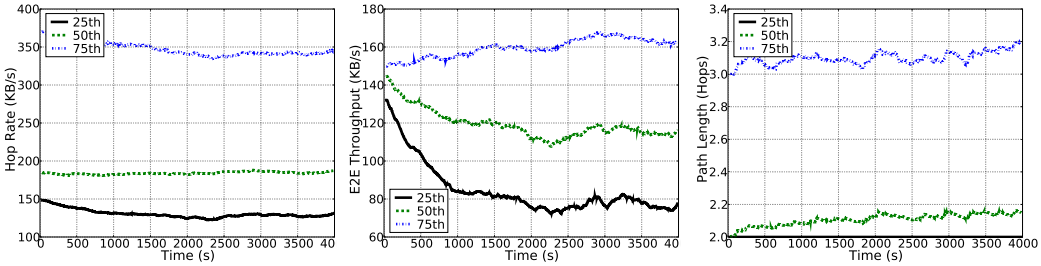

(a) Per-hop throughput (b) End-to-end throughput (c) Path length

Fig. 18. 25^{th}, 50^{th} and 75^{th} percentile performance with each node utilizing a single high-bandwidth channel

REFERENCES

[1] M. Alicherry, R. Bhatia, and L. E. Li, "Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks," in *Proc MobiCom '05*, 2005, pp. 58–72.

[2] P. Bahl, R. Chandra, and J. Dunagan, "Ssch: slotted seeded channel hopping for capacity improvement in ieee 802.11 ad-hoc wireless networks," in *Proc. MobiCom '04*, pp. 216–230.

[3] Hoi-Sheung, W. So, J. Walrand, and J. Mo, "Mcmac: A parallel rendezvous multi-channel mac protocol," *IEEE WCNC 2007*, pp. 334–339.

[4] A. Mishra, V. Shrivastava, D. Agrawal, S. Banerjee, and S. Ganguly, "Distributed channel management in uncoordinated wireless environments," in *Proc. MobiCom '06*. New York, NY, USA: ACM, 2006, pp. 170–181.

[5] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan, "Load-aware spectrum distribution in wireless lans," *IEEE ICNP 2008*, pp. 137–146.

[6] D. R. Pauluzzi and N. C. Beaulieu, "A comparison of snr estimation techniques for the awgn channel," *Communications, IEEE Transactions on*, vol. 48, no. 10, pp. 1681–1691.

[7] E. Rozner, Y. Mehta, A. Akella, and L. Qiu, "Traffic-aware channel assignment in wireless lans," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 11, no. 2, pp. 43–44, 2007.

[8] J. So and N. H. Vaidya, "Multi-channel mac for ad hoc networks: handling multi-channel hidden terminals using a single transceiver," in *Proc. MobiHoc '04*, pp. 222–233.

[9] R. Vedantham, S. Kakumanu, S. Lakshmanan, and R. Sivakumar, "Component based channel assignment in single radio, multi-channel ad hoc networks," in *Proc. MobiCom '06*, pp. 378–389.

[10] A. Wiesel, J. Goldberg, and H. Messer-Yaron, "Snr estimation in time-varying fading channels," *IEEE Transactions on Communications*, vol. 54, no. 5, pp. 841–848.

[11] K. Xing, X. Cheng, L. Ma, and Q. Liang, "Superimposed code based channel assignment in multi-radio multi-channel wireless mesh networks," in *Proc. MobiCom '07*, pp. 15–26.

[12] Y. Yuan, P. Bahl, R. Chandra, T. Moscibroda, and Y. Wu, "Allocating dynamic time-spectrum blocks in cognitive radio networks," in *Proc. MobiHoc '07*, pp. 130–139.