

Distributed Channel Assignment in Multi-Radio 802.11 Mesh Networks

Bong-Jun Ko
IBM T. J. Watson Research
bongjun_ko@us.ibm.com

Vishal Misra
Columbia University
misra@cs.columbia.edu

Jitendra Padhye
Microsoft Research
padhye@microsoft.com

Dan Rubenstein
Columbia University
danr@ee.columbia.edu

Abstract— To increase the utilization of the available frequency channel space in 802.11-based wireless mesh networks, recent work has explored solutions based on multi-radio stations. This paper reports on our design and experimental study of a distributed, self-stabilizing mechanism that assigns channels to multi-radio nodes in wireless mesh networks. We take a modular approach by decoupling the channel selection decision from the data forwarding mechanism, which makes our solution readily applicable to real-world operation when used with emerging multi-radio routing solutions. We demonstrate the efficacy of our protocol on a real-world, 14-node testbed comprised of nodes, each equipped with an 802.11a card and an 802.11g card. We show via extensive measurements on our testbed that our channel assignment algorithm improves the network capacity by 50% in comparison to a homogeneous channel assignment and by 20% in comparison to a random assignment.

I. INTRODUCTION

In multi-hop wireless networks, the management of radio resources (e.g., transmission power control, frequency channel selection, routing, etc.) has a tremendous impact on the performance of the entire system. For instance, the transmission power of wireless devices impacts the topology (i.e., the wireless connectivity between nodes), the interference level, and quality of the wireless links. An effective channel selection mechanism can increase the spatial reuse of the available wireless channels, and improve the overall network capacity. A routing mechanism can choose paths (i.e., a set of wireless links) optimized for certain end-to-end performance metrics (e.g., delay, throughput, hop-count, etc.).

An optimal solution to the radio resource management would clearly need to simultaneously consider all radio resources. However, achieving the stability of a good joint solution is extremely difficult as managing one resource type greatly impacts the management of other resources, and therefore it requires the global status of the network such as traffic matrix and current link condition to be considered. In many cases, such information is very difficult to collect or changes rapidly in reality.

A more practical approach is to *decouple the whole radio resource management problem into those of optimizing for individual resource types*, and in this paper, we focus on the *wireless channel assignment problem in multi-hop wireless networks*. In addressing the channel selection problem, we leverage on the emergence of multi-radio networks, particularly in wireless mesh networking applications, and on recent advances in the multi-hop routing mechanism for multi-radio networks [3].

This material was supported in part by a NSF grant No. CNS-0411047 and a grant from Microsoft Research, and is continuing through participation in the International Technology Alliance sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsoring organizations.

The sensitivity and adaptivity of the channel selection depends on the dynamics of the underlying parameters considered by the decision mechanism. For instance, a channel selection based on quickly-changing network status, such as link quality and traffic condition, might be able to quickly adapt to the dynamics of network condition. However, such an approach would frequently change the connectivity between nodes, and can thus adversely impact the end-to-end transport performance of the network when used in conjunction with a dynamic routing mechanism that also attempts to adapt to the dynamic condition of the network links.

Therefore, an important question that needs to be answered in designing a channel selection mechanism is: *What information should the channel selection be based on?* Considering that the adaptivity to the dynamic network condition can be achieved by end-to-end routing/transport mechanisms such as the one in [3], our answer to the above question is as follows:

- The channels should be selected based only on locally available information.
- The assignment of the channel should be based on the physical structure of the network rather than on the dynamic network condition.
- The change in channel assignment should not frequently alter the connectivity between nodes, rather providing a stable channel environment for the end-to-end routing mechanism.

Note that our approach effectively decomposes the time-scale of resource control into the quickly-changing component (handled by routing) and the slowly-changing one (handled by channel assignment), and helps avoiding aforementioned undesirable cross-layer impact between routing and link layer resource control. Our goal here is therefore to *provide a diverse and quickly-stabilizing channel configuration* based on physical topology of network nodes that the end-to-end routing mechanism can exploit.

This paper is the report of our experimental study of a distributed channel assignment mechanism in multi-radio multi-hop networks. More specifically, we apply our lightweight, self-stabilizing channel selection strategy to wireless mesh networks where nodes are equipped with multiple 802.11 wireless transceivers, experiment it on a mesh network testbed, and measure the impact of our channel selection strategy on the end-to-end transport capacity of the network.

Our contributions are as follows:

- We develop a provably self-stabilizing channel assignment algorithm with which each node greedily selects a channel that minimizes its local objective function using only local information. This algorithm is accompanied with fully-distributed protocol that each node concurrently performs to ensure the convergence of the global assignment.
- We provide a practical channel assignment solution for

802.11-based multi-radio wireless mesh networks where nodes are equipped with limited number of single-channel 802.11 transceivers. Our solution provides a good balance between two conflicting goals of channel assignment, namely the node connectivity and the channel diversity.

- We implement and evaluate the performance of our distributed solution in a 14-node mesh network testbed in which each node is equipped with two off-the-shelf 802.11 interfaces. We measure the aggregate throughput capacity of the testbed network in which nodes are assigned channels using our channel selection mechanism, with the MR-LQSR protocol of [3] being used as the end-to-end routing mechanism.

We observe that our channel assignment improves the aggregate throughput of the network on average by 50% compared to the case when all nodes are assigned to the same channels, and by 20% over when channels are assigned at random. This is an impressive result particularly since the utilization of the channels assigned by our mechanism is quite limited due to heavy spectrum usage by external traffic in other access points around our testbed, yet we find a small improvement in spectrum utilization results in a marked improvement in network throughput.

This paper is organized as follows. After reviewing related work in Section II, we describe the mesh network system architecture and our network model in Section III. Then we present our channel selection mechanism in Section IV, and report the performance evaluation results in Section V. Section VI concludes the paper.

II. RELATED WORK

Much of the recent work in multi-channel 802.11 routing has looked at jointly solving the channel assignment and routing problem. A heuristic solution is looked at in [12], an algorithmic approach that optimizes for throughput is considered in [1], and an approach that preserves network connectivity for QoS is explored in [13]. These are centralized solutions that assume the availability of a global network view (e.g., traffic demand, nodes' status, etc.). In contrast, our modular approach decouples the channel assignment and routing problem separately, with both being solved in a fully distributed manner.

Raniwala et al. [11] propose a distributed channel assignment algorithm for 802.11-based multi-radio mesh network and perform an experimental evaluation. However, the network architecture in [11] is designed for mesh networks specifically used for the wireless Internet access applications, and their channel assignment algorithm works only for routers whose connectivity graph is a tree. In their assignment mechanism, the channel assignment to nodes positioned higher in the tree affects all nodes lower in the tree hierarchy. In contrast, our algorithm can operate on any arbitrary network structure, where every mesh node performs the same assignment task in a fully-distributed manner.

Ramachandran et al. [10] propose a centralized channel assignment algorithm which is performed by a central server that periodically collects dynamically-changing channel interference information. The comparison to this work is of particular interest as their channel selection method takes into account dynamically changing network status (i.e., interference), while our channel assignment is based on more static information (i.e., physical topology). In summary, the performance gain of our mechanism observed in the real-world testbed experiment appears similar to what is shown in their simulation results. This suggests the efficacy of our solution since our distributed

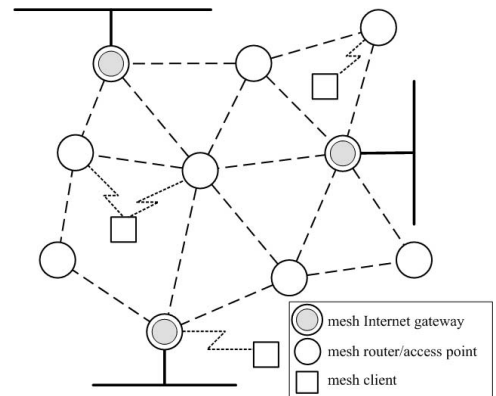


Fig. 1. Wireless mesh network architecture

mechanism requires only localized interaction between nodes, and does not need to be performed many times once it stabilizes, thus incurring much less overhead in performing the channel assignment than their approach.

In [8], Mishra et al. explore the possibility of utilizing partially-overlapping wireless channels in 802.11 access points, and show that intelligent assignment of non-orthogonal channels increases overall channel utilization and the system performance. We integrate their observation into the design of our assignment algorithm, and investigate the efficacy of a simple interference model with our testbed experiment.

Our design of the distributed protocol is motivated by our previous theoretical work on a fully-distributed, self-stabilizing protocol for replica placement [7]. In this work, we propose a distributed replica placement scheme with which identical replicas are placed “far” from one another. We adapt the distributed approach of assigning replicated items to the wireless channel utilization problem, where the channel is the replica. This paper focuses much more on the practicalities of this problem that are ignored in the theoretical work.

III. ARCHITECTURE AND MODEL

In this section, we present the wireless mesh network architecture and the network model considered throughout the paper.

A. Mesh network architecture

Figure 1 depicts a generic wireless mesh network architecture. The mesh network consists of three types of wireless nodes: 1) **mesh gateways** are the relaying points of the traffic between wireless mesh network and external wired network, 2) **mesh routers** are 802.11 wireless nodes (typically stationary) that act as the wireless access points of the mesh clients, and form a multi-hop wireless network infrastructure, and 3) **mesh clients** are end-user 802.11 nodes, each connected to at least one (typically only one) mesh router to have their packets forwarded from/to mesh gateways or other mesh clients.

Note that a mesh gateway can also serve a dual function and also act as a mesh router, when desired. This dual role is useful when network flows within the mesh network is not only between mesh clients and mesh gateways, but also between pairs of mesh clients. Community surveillance, emergency service, and community resource sharing are potential mesh network applications that would generate a good amount of client-to-client traffic as well as client-to-gateway traffic [2].

Our focus is on how to utilize the 802.11 channels within the wireless network of mesh routers, allowing us to ignore the mesh clients and mesh gateways.

B. A Simple Model

Our model of the previously described mesh architecture consists of a set of N nodes, $V = \{1, 2, \dots, N\}$. There are K wireless channels, $1, \dots, K$, whose *frequency spectrum can possibly overlap*. A node i can communicate with some other node j only if i and j share a common channel assigned to some of their interfaces. In this case, a node i is said to be a *neighbor* of node j if i is within a distance from j within which i can correctly decode the packets transmitted by j .

A *channel interference cost function* (simply referred to as *cost function* in the remainder of the paper), $f(a, b)$, provides a measure of the spectral overlapping level between channels a and b . The interference cost function is defined in such a way that $f(a, b) \geq 0$ and $f(a, b) = f(b, a)$, where a value of 0 indicates that channels a and b do not interfere with each other. Also, $f(a, b)$ decreases as the the spectral distance between channels a and b grows.

Note that our definition of interference cost function is *not* meant to represent the actual wireless interference level mutually experienced by two nodes in spatial domain. Rather, it can be interpreted as the first-level approximation of interference level between two (overlapping) *channels* in spectral domain, and hence the symmetry argument can hold.

Though our channel selection mechanism works for any interference cost function, we use in our experiment in Section V a relatively straightforward cost function which linearly decreases with the spectral distance between two channels with a single tunable parameter δ :

$$f(a, b) = \max(0, \delta - |a - b|), \quad (1)$$

where a large δ models a more heavily overlapping spectral separation between adjacent channels, and small δ represents a less overlapping channel space.

A node i belongs to a node j 's *interference set*, S_j , of node j if there exists a node (either i , j , or possibly a third node k) for which transmissions from i can be corrupted by transmissions from j . The problem of accurately determining the interference patterns within a network is a difficult one [9]. In our experiments, we use the heuristic from [5] that assumes that all other nodes within three hops of node i are in node j 's interference set.

IV. DISTRIBUTED CHANNEL ASSIGNMENT

In this section, we describe our distributed channel assignment mechanism, beginning with our distributed channel selection strategy, which is then used as the baseline of our complete solution to channel assignment for multi-radio nodes with a limited number of single-channel interface cards.

A. Channel selection algorithm

Consider the following channel selection algorithm with which each node continually seeks to greedily improve its current choice of channel:

Algorithm 1: ChannelSelection(node i)

Input

S_i : Set of nodes in i 's interference range.

c_j : The channel of each node $j \in S_i$

c_i : i 's current channel

begin procedure

for all $k = 1, \dots, K$,

$F(k) \leftarrow \sum_{j \in S_i} f(k, c_j)$.

if $F(c_i) > F(k)$ for any $k = 1, \dots, K$, **then**

$c_i \leftarrow k_{min}$ where $k_{min} = k : F(k) \leq F(k') \forall k' = 1, \dots, K$

end if
end procedure

In other words, a node i greedily selects a channel that satisfies its local objective of *minimizing the sum of interference cost* from the set of nodes within its interference range.

It is not intuitively obvious that this distributed channel selection process is self-stabilizing, i.e., that nodes continually looking to improve on their local interference cost will eventually converge to a stable channel allocation; one node's channel change can increase some other node's interference level, and cause the other node to change its channel, and so forth. However, we can show that indeed this process does stabilize.

Theorem 1: If every node selects its channel following Algorithm 1, within a finite number of channel changes by nodes, the channel assignment reaches a stable state where nodes cease changing channels.

Proof: We provide here a sketch of the proof (details can be found in [6]). The core part of the proof is that a monotonicity property of the global channel configuration holds whenever a node changes its channel. More specifically, whenever any node i changes its channel to decrease its current interference cost, $F(c_i)$, the total interference level of all nodes, $F = \sum_{i \in V} F(c_i)$ decreases as well. This is because, node i changes its channel only when its interference cost can decrease. But because of the symmetry in the interference cost function, the aggregate interference that other nodes in S_i "gain" due to i 's channel change is always smaller than what they "lose" (otherwise, i would not have changed its channel). Since i 's channel change does not affect nodes outside i 's interference range, the aggregate interference level F of all nodes should decrease whenever any node changes its channel. ■

An interesting artifact that the above sketch of proof reveals is that each node's greedy choice to improve its local objective results in the improvement in global objective of total interference level, and eventually leads to a channel assignment in which all nodes are satisfied with their channel choice.

An implicit condition for Algorithm 1's stabilization is that, however, when a node changes its channel, it has the correct knowledge about other nodes' current channels. However, since nodes can change their channels multiple times with the algorithm, this condition is often broken in practical networking environments where there is latency in exchanging information between nodes. Therefore, a sophisticated coordination must be performed to ensure the correctness and convergence of the algorithm's operation (otherwise, deadlocking or livelocking condition can easily occur).

We employ a distributed, mutual-exclusion protocol, which is similar to what we develop in our previous theoretical work [7], to guarantee the stabilization of the algorithm in a fully-distributed manner. In this protocol, nodes perform an asynchronous, distributed mutual exclusion operation by exchanging three-way handshake messages with other nodes in their interference range to ensure the correctness of channel information and to prevent simultaneous channel changes. As opposed to total ordering mechanisms that are typically used to serialize distributed operations, our protocol enables the concurrent execution of the algorithm by multiple nodes in the network, thus achieving much higher level of parallelism. More detailed description of this distributed protocol can be found in [6] and [7].

B. Applying to 802.11 Multi-radio network

Now we provide our complete solution to the channel assignment for multi-radio 802.11 networks. Since the commodity 802.11 transceivers are single-channel, half-duplex devices, a wireless node with k 802.11 transceivers can communicate simultaneously on up to k wireless channels, where the number of transceivers is typically smaller than the number of available channels (e.g., 12 orthogonal channels in 5GHz band, and 3 in 2.4GHz band).

The fact that the number of transceivers per node is limited raises the trade-off between two conflicting objectives of the channel diversity and the node-connectivity: While channels need to be assigned to a node in such a way that wireless interference from other nodes at the same channel can be minimized, a node should also share some common channels with its neighbors so that the original network communication graph (determined by the physical location of nodes) should not be partitioned.

We strive to find a good balance between the channel diversity and the network connectivity with the following assignment rules:

- 1) One interface of each node is dedicated to a default channel common to all nodes. This ensures a basic connectivity between neighboring nodes.
- 2) The assignments to each of the remaining interfaces are performed using Algorithm 1, with one exception: the selected channel must be one of those already assigned to some neighbors within communication range. This prevents the interface from being assigned a useless channel.
- 3) The distributed protocol for the algorithm's stabilization remains the same: use the three-way handshake when changing to a channel selected from this limited subset of channels.

This assignment strategy exhibits a few desirable properties. First, the common channel ensures the connectivity between neighboring nodes is preserved, even when some nodes have only one interface. The presence of the common channel provides further protection from topology changes due to nodes join, move, or failure. If all interface cards were utilized for the variable channel assignment, the network connectivity would be easily compromised even upon a slight fluctuation in the topology.

Second, the channels for additional interfaces are assigned in such a way that nodes share the same channel with some of their immediate neighbor nodes, while other non-neighboring nodes in the interference range are likely to be on different channels. This is a good allocation strategy that provides both the connectivity and small interference since the channels indeed shared by communicating neighbors will likely experience low interference level from other nodes in interference range.

Third, it has a positive impact on multi-channel routing protocols. Nodes under this assignment are likely to be grouped together on the same channel assigned to non-default interfaces with some of their neighbors. As a result, these groups of neighboring nodes on the same channel, with different groups likely being on different channels, are overlaid on top of prevalent links formed by the default channel among all nodes (one may view this as the "islands" of fast-channel links on the "sea" of common-channel links). Therefore, a routing protocol is given abundant choices of channel-diverse paths that consist of some express links (those formed by channels at additional interfaces) and some local links on the default channel.

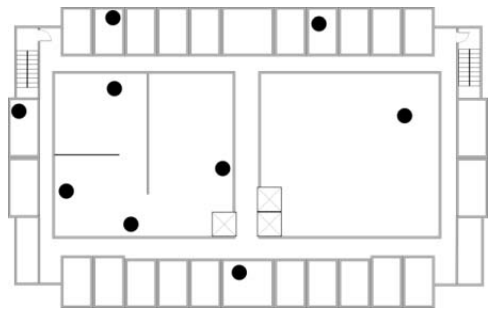


Fig. 2. Wireless mesh network testbed : showing node locations in one of three floors where mesh nodes are placed. the typical distance between closely nodes is about 10 meters.

V. PERFORMANCE EVALUATION

In this section, we report the performance evaluation results obtained via experiments on our mesh network testbed in which each node is equipped two 802.11 interfaces.

A. Testbed and Methodology

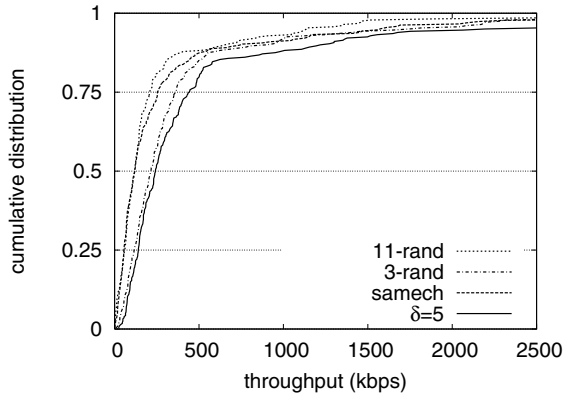
Our mesh network testbed consists of 14 nodes, placed throughout three floors of a multi-story building in Columbia University (see Figure 2). Each mesh node is equipped with two IEEE 802.11a/b/g wireless interface cards. To overcome the cross-channel interference between two 802.11 cards in physical proximity within a node (reported by others as well [3], [12]), we assign channels from two different frequency bands to the two cards, i.e., one from 2.4GHz 802.11g band, and the other from 5GHz 802.11a band.

We use MR-LQSR protocol with the WCETT metric [3] as the routing protocol in the experiment. The WCETT routing metric is designed to select network paths based on the estimated quality of each individual links and the overall channel-diversity of paths in multi-radio environments.

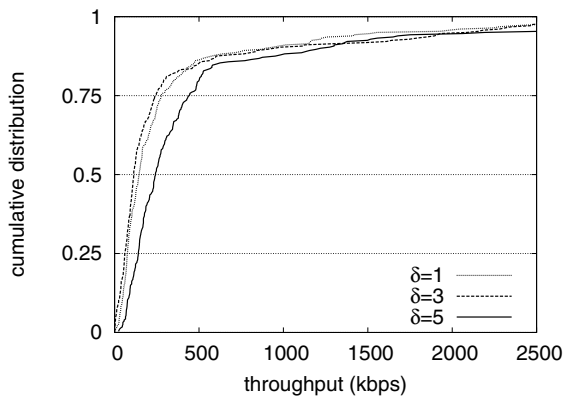
To evaluate the transport capacity of the network, we measure the aggregate throughput achieved by multiple simultaneous TCP flows. Each node in the network acts as a source of a TCP flow. Thus, we have 14 concurrent TCP flows (i.e., 14 source-destination pairs) in our network for each experiment. The destination of each TCP flow is chosen at random, while we avoided choosing single-hop destinations in order to eliminate the bias caused by the dominance of single-hop TCP flows [4]. We generated four different sets of such TCP flows, with each set denoted by **fset1**, **fset2**, **fset3**, and **fset4**, respectively. All flows start simultaneously, and each flow sends data as fast as TCP permits for 120 seconds, creating heavy loads on the network.

Our experiments showed that, in our testbed, we could achieve higher overall throughput (in all cases) when the 802.11a band provided the default channel and a varied channel assignment was implemented within the 802.11g band. We suspect this is due to the interference caused by heavy usage of 802.11b/g infrastructure network around our testbed, which reflects well the typical channel usage environment considering higher market penetration of 802.11b/g devices than that of 802.11a ones at the time of our experiment. The results in this section reports on the cases when one network card of each node is assigned to a common default channel (channel number 36 at 5180 kHz), and the other network card is assigned one of (possibly overlapping) 11 channels in 802.11g band in the following manner.

We first consider channel assignments generated by our distributed assignment using three different values for $\delta = 1$,



(a) Comparison to baseline allocation strategies



(b) Comparison between different choice of δ

Fig. 3. Cumulative distribution : throughput of individual TCP flows

3, and 5 in the interference cost function in Section III-B, and refer the respective assignments as $\delta=1$, $\delta=3$, and $\delta=5$. Note that $\delta=5$ reflects the fact that only channels separated by 5 channels (e.g., channels 1 and 6) are truly orthogonal in 802.11g band. Smaller values of δ allow our algorithm to aggressively assign partially-overlapping channels.

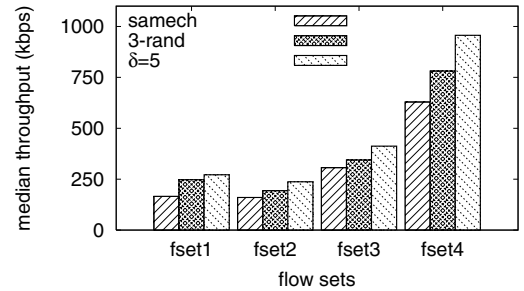
We consider three baseline channel assignment strategies to compare with our assignment:

- **samech** (same channel assignment): all nodes are assigned the same 11g channel.
- **11-rand**: each node is assigned one of 11 802.11g channels selected uniformly at random. This assignment corresponds to the case of $\delta = 0$ in the context of our cost function.
- **3-rand**: each node is assigned one of three orthogonal 802.11g channels (i.e., channel 1, 6, and 11) selected uniformly at random.

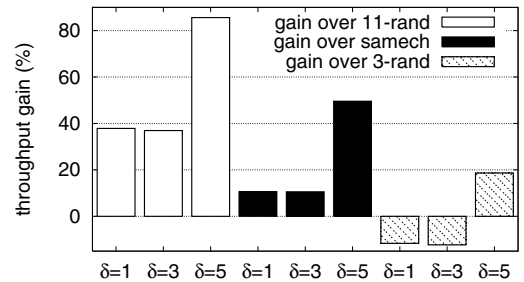
For each of the six channel assignment strategies described above (3 baselines + 3 for different δ values), we generate 5 different channel allocations. The different allocations for 11-rand and 3-rand are generated by using different random seeds, and for the **samech** assignment, we use channels 1, 3, 6, 9, and 11 in 802.11g radio for those 5 allocations. Thus we have 30 total channel allocations. For each channel allocation, we run the 4 traffic sets described earlier.

B. Evaluation Results

In Figure 3(a), we compare the CDF of the throughput of individual flows under $\delta=5$ channel assignment to that under



(a) Median throughput of samech, 3-rand, $\delta=5$



(b) Gain(in %) in median throughput of $\delta=1$, $\delta=3$, and $\delta=5$ over 11-rand, samech, and 3-rand

Fig. 4. Median network throughput

the baseline assignment. The x -axis in the figure represents the throughput in kbps, and the y -axis indicates the ratio of the number of flows that achieve the end-to-end throughput up to the value in x -axis.

The performance gain by our mechanism is clear. For instance, at the 75-percentile mark, the throughput by our algorithm with $\delta=5$ reaches 440 kbps, while the best of the other assignments (3-rand) achieve no more than 350 kbps. Also the benefit of $\delta=5$ assignment is the most noticeable in the region that the per-flow throughput is relatively small (up to 75% percentile mark, or below 500 kbps). These lower-rate flows typically traverse many hops¹, and hence are more adversely affected by high interference conditions. Hence, we see these flows over longer paths benefit significantly from a well-designed channel selection protocol.

In Figure 3(b), we compare the impact of the choice of δ for our channel assignment. It can be clearly seen that $\delta=5$ results in the best per-flow throughput among different cost function parameters. This result is an anticipated one since the cost function with $\delta = 5$ reflects the fact that channels separated by at least 5 channel(e.g., channel 1 and 6) are in fact orthogonal to one another.

We now look at the aggregate network throughput. In Figure 4(a), we compare the median aggregate network throughput of the best case of our channel assignment (i.e. $\delta=5$) to those of the two other baselines (i.e. samech and 3-rand) for each traffic set. We see that $\delta=5$ outperforms the best case among the baseline assignments (3-rand) by 20%.

To see more clearly the performance gain by our mechanism, we depict in Figure 4(b) the percentage improvement in median throughput achieved by our channel assignments (for three values of δ) over the three baseline assignments(11-rand,

¹Note that MR-LQSR routing protocol sometimes prefer paths with many more hops over those with smallest number of hops in the network in order to find channel-diverse routes.

samech, and 3-rand). Overall, our assignments significantly outperform 11-rand assignment by 40 to 80%, and samech assignment by 10 to 50%. The comparison to samech assignment is particularly interesting since it indicates that utilizing even partially overlapping channels exhibits better performance than tuning the interface cards of nodes to the same channel. This result also coincides with what is reported in [8] for access-point wireless networks.

Furthermore, the 50% improvement observed in $\delta=5$ assignment over samech assignment can be compared to what is reported in [10], in which channels are assigned to multi-radio nodes with the dynamic interference level taken into account. In their simulation study, their channel assignment scheme resulted in around 50% throughput gain on average (and around 25% when external interference is present) over the homogeneous assignment. Though it is admittedly hard to quantitatively compare their results with ours, this can tell us the effectiveness of our mechanism as our results show similar (or better) performance gain with less protocol overhead (recall that, once it stabilized, our channel assignment does not change until topology changes) in more realistic environment under heavy external interference.

The performance gain of our assignments is lower when compared to the 3-rand assignment. While the $\delta=5$ assignment outperforms 3-rand by 20%, lower values of δ perform worse than 3-rand. We suspect that the relatively good performance of 3-rand assignment is an artifact of our testbed setup, in which many nodes are densely clustered together around the center of the network. When only 3 channels are randomly assigned, nodes are likely to have some neighbor assigned on the same channel in this dense area, and the probability that the assigned channel is isolated becomes very low.

samech	11-rand	3-rand	$\delta=1$	$\delta=3$	$\delta=5$
10.1	3.0	15.1	16.3	11.6	22.7

TABLE I

CHANNEL UTILIZATION(IN %): PERCENTAGE OF TCP THROUGHPUT CARRIED ON 802.11G CHANNELS

Now we investigate how the channels are utilized in the experiments. Table I shows the utilization of the links assigned channels in 802.11g band. Here the utilization measures the percentage of end-to-end traffic carried on channels on links on the 802.11g band. For instance, under our channel assignment with $\delta=5$, the wireless links assigned to 802.11g channels collectively carried 22.7% of end-to-end traffic, while links of the common 802.11a channel contributed for the other 77.3%. The utilization of 802.11g channels is in large part lower due to the active usage of 802.11b/g infrastructure network around our testbed. Nevertheless, we see that our channel assignment (especially $\delta=5$) makes better use of 802.11g channels, and thus reduces the congestion on the common 802.11a channel. This statistics reveals an interesting fact that a small enhancement in channel utilization can result in a large improvement in overall network throughput, which speaks for the importance of efficient radio resource management.

We also collected the protocol dynamics statistics (such as number of protocol messages per node, time to convergence, etc.). Though not shown here due to space limitation, the collected data suggests our protocol is very light-weighted, showing that overall network resource usage by our protocol is very low, and the convergence time of the protocol is quite small. This enables us to believe our distributed mechanism is suitable for the channel assignment task in large-scale wireless mesh networks.

VI. CONCLUSION

We presented a fully-distributed mechanism that assigns 802.11 channels to multi-radio nodes in wireless mesh networks. Our assignment mechanism stabilizes to a desirable channel configuration that strikes a good balance between network connectivity and channel diversity, and that routing protocols can exploit to provide better end-to-end system performance. Our design takes into account several constraints present in current 802.11 devices, and its distributed nature ensures it is sufficiently light-weight to be executed on large-scale mesh networks.

The modular design that decouples channel selection from data forwarding makes our solution readily available for real-world operation, providing a complete solution to wireless mesh networks in combination with existing routing protocols. We ran experiments on our wireless mesh network testbed and showed that our channel assignment can increase the capacity of wireless mesh network between 20% and 50% over other conventional channel selection mechanisms.

In the future, we plan to explore the impact of mesh clients to the performance of our assignment while expanding the size of the testbed. We anticipate even greater performance improvement in comparison to conventional method in larger settings. Also, a formal investigation into the time-scale decomposition of radio resource control across multiple layers will be fruitful in studying the stability and robustness of resource allocation mechanisms in multi-hop wireless networks.

ACKNOWLEDGEMENT

The authors would like to thank Kyung-Wook Hwang for helping us run the experiments reported in this paper.

REFERENCES

- [1] Mansoor Alicherry, Randeep Bhatia, and Li (Erran) Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *MobiCom '05*, pages 58–72, New York, NY, USA, 2005. ACM Press.
- [2] P. Bahl. Opportunities and challenges of community mesh networking, keynote @ mics workshop eth zurich, july 6, 2004.
- [3] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *MobiCom'04*, 2004.
- [4] M. Gerla, R. Bagrodia, L. Zhang, K. Tang, and L. Wang. TCP over wireless multihop protocols: Simulation and experiments. In *Proceedings of IEEE ICC '99*.
- [5] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE/ACM Transactions on Information Theory*, 46(2), March 2000.
- [6] Bong-Jun Ko, Vishal Misra, Jitendra Padhye, and Dan Rubenstein. Distributed self-stabilizing channel assignment in multi-radio 802.11 mesh networks. Technical report, Columbia University, November 2005.
- [7] Bong-Jun Ko and Dan Rubenstein. Distributed self-stabilizing placement of replicated resources in emerging networks. *IEEE/ACM Transactions on Networking*, 13(3):476–487, 2005.
- [8] A. Mishra, E. Rozner, S. Banerjee, and W. Arbaugh. Exploiting partially overlapping channels in wireless networks: Turning a peril into an advantage. In *IMC '05*.
- [9] J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of link interference in static multi-hop wireless networks. In *Short paper, IMC 2005*.
- [10] Krishna N. Ramachandran, Elizabeth M. Belding-Royer, Kevin C. Almeroth, and Milind M. Buddhikot. Interference-aware channel assignment in multi-radio wireless mesh networks. In *IEEE Infocom'06*.
- [11] Ashish Raniwala and Tzi cker Chiueh. Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. In *IEEE Infocom'05*.
- [12] Ashish Raniwala, Kartik Gopalan, and Tzi cker Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. *Mobile Computing and Communication Review*, 8(2):50–65, 2004.
- [13] Jian Tang, Guoliang Xue, and Weiyi Zhang. Interference-aware topology control and qos routing in multi-channel wireless mesh networks. In *MobiHoc '05*, pages 68–77, New York, NY, USA, 2005. ACM Press.