Practical, Distributed Channel Assignment and Routing in Dual-radio Mesh Networks

Aditya Dhananjay
New York University
aditya@cs.nyu.edu

Hui Zhang
Tsinghua University
zhanghui04@mails.tsinghua.edu.cn

Jinyang Li
New York University
jinyang@cs.nyu.edu

Lakshminarayanan Subramanian
New York University
lakshmi@cs.nyu.edu

ABSTRACT

Realizing the full potential of a multi-radio mesh network involves two main challenges: how to assign channels to radios at each node to minimize interference and how to choose high throughput routing paths in the face of lossy links, variable channel conditions and external load. This paper presents ROMA, a practical, distributed channel assignment and routing protocol that achieves good multi-hop path performance between every node and one or more designated gateway nodes in a dual-radio network. ROMA assigns non-overlapping channels to links along each gateway path to eliminate intra-path interference. ROMA reduces inter-path interference by assigning different channels to paths destined for different gateways whenever possible. Evaluations on a 24-node dual-radio testbed show that ROMA achieves high throughput in a variety of scenarios.

Categories and Subject Descriptors
C.2.2: Computer Communications Networks
General Terms
Algorithms, Design, Performance
Keywords
Wireless, Routing, Channnel Assignment

1. INTRODUCTION

Wireless mesh networks comprised of nodes having multiple radios (multi-radio mesh networks) have the potential to perform significantly better than single radio mesh networks. Since every node operates its radio on the same channel in a single-radio mesh network, a forwarding node interferes with the two subsequent nodes along any multi-hop path, drastically reducing the end-to-end throughput [7, 23]. A multi-radio mesh can eliminate such intra-path interference if potentially interfering links are operated on non-overlapping channels. Another important advantage of multi-radio networks is the ability to use many non-overlapping channels in the same physical region. As a result, there is less inter-path interference among multiple flows in a multi-radio mesh, resulting in higher aggregate throughput.

While there has been significant work on multi-radio mesh protocols [16, 20, 30, 32, 4, 29, 13], realizing the full potential of multi-radio mesh networks in real-world settings has remained a challenging problem. Real-world deployments, especially in urban settings, pose many practical challenges and constraints that affect both the design and performance of a multi-radio protocol. To the best of our knowledge, only a few of protocols [32, 31, 20] have been implemented and even fewer have been evaluated on a testbed of reasonable scale [31, 20].

Each node in a multi-radio network can be equipped with only a few radios. Commodity radios operating in the same frequency band interfere within close proximity (up to 18 inches). Since there are only two frequency bands (2.4 and 5.2 GHz) for use by 802.11 today, a physically compact node is restricted to using only two radios per node. Thus, a multi-radio protocol should perform well on a dual-radio mesh but also be extensible to handle more than two radios per node, should additional orthogonal frequency bands become available.

Channels must be assigned carefully to reduce interference in the network. However, when there are only a few radios at each node, it is not feasible to optimize for all paths simultaneously. Fortunately, not all paths are equally important. Most mesh deployments today have a few pre-specified gateway nodes and users care most about achieving high throughput on multi-hop paths from each non-gateway node to a gateway. To take advantage of such traffic patterns, each node should choose routes and channel assignments together to optimize for its gateway paths: when done correctly, one can construct multi-hop gateway paths consisting of high quality links operating on non-overlapping channels and also reduce inter-path interference among paths to different gateways.

In this paper, we present the design, implementation and evaluation of ROMA, a distributed routing and channel assignment protocol that achieves high end-to-end performance for gateway paths in a dual-radio mesh. In ROMA, each gateway chooses a channel sequence, e.g. $c_1, c_2, ...$, to guide other nodes' channel assignment. Specifically, a node

SIGCOMM'09 August 17–21, 2009, Barcelona, Spain.

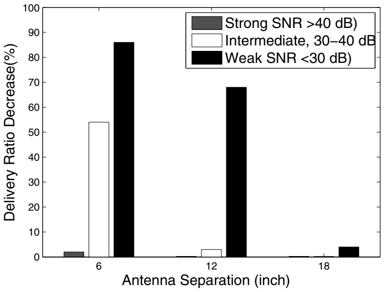

Figure 1: The reduction in packet delivery ratio at the receiving radio for various received signal strengths when a closeby radio transmits. The receiving and transmitting radios operate on channel 40 (5.22 GHz) and 165 (5.825 GHz), respectively. Their antennas must be separated by 18 inches to prevent the transmitter from interfering with a weak received signal.

i hops away assigns channels c_i and c_{i+1} according to the corresponding gateway's sequence. Since each sequence consists of distinct, non-overlapping channels, all gateway paths avoid intra-path interference. ROMA reduces inter-path interference as multiple interfering gateways try to use different channels in their channel sequences. Although proposals that perform joint channel assignment and routing exist [32, 4], ROMA is the first distributed joint protocol that addresses real world challenges such as lossy and highly variable channel conditions. In particular, ROMA contributes a novel measurement-driven path metric that takes into account link delivery ratios, fluctuations in link quality as well as external load. This path metric allows ROMA to choose multi-hop paths with good performance.

Using a detailed evaluation of ROMA on a 24-node dual-radio testbed, we show that ROMA achieves high end-to-end throughput; Paths with three or more hops have a median throughput of 4.1 Mbps, a mere 7% drop in performance compared to that of single-hop paths. ROMA's median aggregate throughput reaches 14.8 Mbps with three gateway nodes, which is 1.4× what is achieved when restricting all nodes to use a common channel and 2.1× what is achieved when assigning identical channels to all nodes.

2. CONSTRAINTS AND CHALLENGES

This section describes our problem setting and outlines important practical constraints and challenges.

2.1 The Case for Dual Radios

A multi-radio node forwards packets by simultaneously transmitting and receiving on different radios. Although there are many orthogonal channels (3 for 802.11b/g, 13 for 802.11a), it is challenging for a multi-radio node to use different channels from the same frequency band because a node's transmitting radio might interfere with its receiving radio, unless the two radios are separated by a sufficient distance. In order to understand these radio separation constraints, we performed the following experiment with two mesh nodes: We used one node to receive packets sent from a laptop while the other node was simultaneously transmitting packets. The receiving radio operates on channel 40 (5.2 GHz) and the transmitting radio is on channel 165 (5.825 GHz). We varied the physical distance between the two mesh nodes. Furthermore, we also changed the distance between the laptop and the receiving radio to vary the received signal strength. Figure 1 shows that, in order to prevent the transmitting radio from interfering with a relatively weak received signal, their antennas must be separated by at least 18 inches. Similar results are obtained with a variety of different cards and chipsets [11, 34, 2].

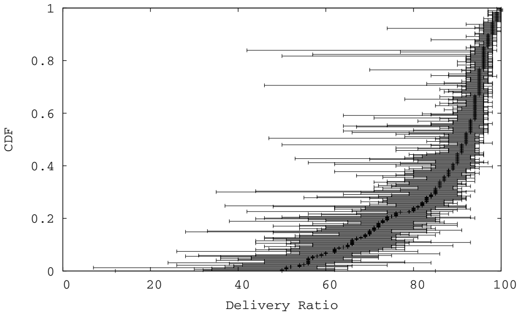

Figure 2: The distribution of the median link delivery ratios on channel 6 on a 24-node testbed. The errorbars correspond to the 20%,80% delivery ratio on the same link observed over a 20 second period.

To avoid interference in the same frequency band, one could ensure antenna separation using long pigtails [11], USB cables [13, 30] or Ethernet connections [31]. However, the resulting increase in node size is non-trivial and would significantly limit node placement, especially in indoor settings. As there is no interference among channels in different frequency bands, we can build compact dual-radio nodes by operating a node's two radios on 802.11a and 802.11b/g channels. Any 3-radio compact mesh node is bound to have interference across simultaneously sending and receiving radios because at least two of them have to operate on the same frequency band. To maintain the deployment advantage of compact nodes, we focus on dual-radio mesh networks. Our protocol can also be extended to work with more than two radios at each node (Section 3.6).

2.2 Problem Setting and Challenges

The basic problem we address is: given a dual-radio mesh network, how does a distributed protocol assign channels and select routes that achieve high end-to-end performance?

The channel assignment challenge: Multi-radio networks achieve high performance by assigning non-overlapping channels to eliminate harmful intra-path interference and reduce inter-path interference whenever possible. For a single multi-hop path, one can easily assign channels to eliminate intra-path interference: each forwarder uses two distinct channels to communicate with its previous and next hop neighbor. Channel assignment becomes much more challenging if it is to reduce interference for *all* paths under arbitrary traffic patterns, since each node has only a few radios (two in our case), far fewer than the number of available non-overlapping channels.

Most prior proposals either use a centralized assignment algorithm or require all nodes to operate one of its radios on a common channel. Unfortunately, neither approach is satisfactory. Centralized algorithms without the use of a common channel cannot adapt robustly to cope with net-

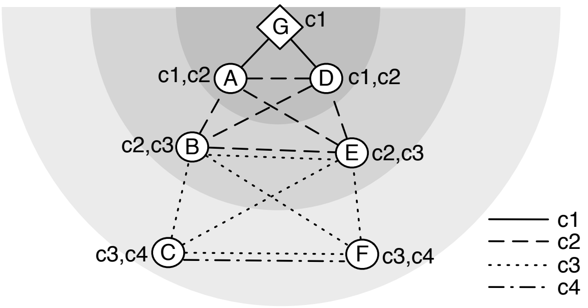

Figure 3: By assigning the same set of channels for all nodes residing on the same routing level, ROMA eliminates intra-path interference while preserving many cross links between paths.

work changes. Assigning a common channel to all nodes maximizes network connectivity but causes half of all links to operate on the same channel (the common channel), resulting in increased intra and inter-path interference. For a mesh network that provides Internet access to many clients, we can exploit the predominant traffic pattern to optimize the performance of gateway paths only. To do so, a multi-radio protocol should jointly choose gateway paths and channel assignments to ensure that each gateway path consists of high quality links operating on distinct channels and that paths to different gateways use different available channels whenever possible.

Routing challenges: Since intra-path interference is unlikely with careful channel assignment, the throughput of a multi-hop path in multi-radio networks is limited by its worst performing link. It is difficult to estimate link quality: not only are links lossy, but loss rates also vary across different timescales [3, 12]. Figure 2 characterizes the link delivery ratios in our 24-node testbed on channel 6 (2.437 GHz). The delivery ratio of a link is the fraction of successfully received broadcast packets during one second. 20 measurements were taken for each link over a duration of 20 seconds. Figure 2 plots the distribution of the median delivery ratio with error bars at 20% and 80%. As shown in the graph, the delivery ratio of many links fluctuates over a short time period. Such fluctuations can lead to suboptimal routes and unnecessary route changes. In addition to loss variations, practical mesh deployments have to share the 2.4 GHz frequency band with many types of popular devices (for example, cordless telephones) and potentially, a large number of access points. As a result, there is often significant external competition for the channel. Ignoring such competing traffic result in paths that under-perform.

3. THE DESIGN OF ROMA

In this section, we introduce the design of ROMA (Routing over Multi-radio Access Network), a distributed protocol that chooses routing path *and* channel assignment together to optimize the path throughput between every node and a few gateways. We first motivate the basic idea of ROMA before describing its design details.

3.1 The simplified scenario

We discuss ROMA's main idea in a simplified setup where the network has a single gateway with one radio and all other nodes have two radios. In the simplified case, ROMA aims to assign channels to eliminate intra-path interference for all routing paths to a single gateway radio (similar to [32]). In the single gateway radio case, inter-path interference is not an issue since competing flows contend for the same gateway resource. In the multiple gateway case, ROMA also aims to reduce inter-path interference for paths destined to different gateway radios (Section 3.5).

With a single gateway, we can view all nodes as residing on different levels based on their path length to the gateway. Intuitively, the network forms a ring-like pattern emanating from the gateway, as shown in Figure 3. If we assign all nodes within the same level the same two channels and let nodes on adjacent levels share one channel in common, all paths to the gateway can operate on distinct channels, eliminating intra-path interference. In Figure 3, nodes A,D belong to the same level and assign themselves channels c_1, c_2 and nodes B,E use channels c_2, c_3. As a result, all multi-hop gateway paths, such as G-A-B-C, can use different channels for each of their links.

It may appear counter-intuitive why ROMA assigns the same channels instead of different ones to nodes at the same level. This design choice is based on two considerations: first, there is little performance benefit in assigning different channels to nodes at the same level since all paths ultimately compete with each other at the first hop of the gateway. Second, there is an advantage in having the same channel assignments at the same level because it preserves many cross links between paths to the same gateway. In the example of Figure 3, links such as A-E and B-E would not exist had nodes at the same level assigned themselves different channels. When assigning channels with the goal of avoiding interference, a multi-radio protocol tends to make the network topology less dense, reducing chances for opportunistic routing [9, 10]. By assigning the same channels to nodes at the same level, ROMA preserves cross links whenever possible. These cross links are useful in opportunistic routing and also help a node adapt to link condition changes quickly. For example, if the existing route G-A-B-C degrades, node C can use a different route such as G-A-E-C without having to change its channel assignment.

ROMA's basic design: The basic channel assignment strategy works as follows. Each gateway associates a channel sequence with each of its radios, e.g. $c_1, c_2, c_3, \ldots$. Every non-gateway node in the network participates in a distributed routing protocol to discover its best gateway path that is likely to yield good throughput according to ROMA's path metric. Since a gateway's channel sequence is propagated along with routing information in periodic route announcement messages, a node calculates the best path to the gateway, such that the channels along this path satisfy the gateway's channel sequence. Thus, the node can assign appropriate channels according to its gateway path length and the gateway's channel sequence. A node i hops away from the gateway assigns channels (c_i, c_{i+1}) to its two radios. For example in Figure 3, node C finds its best gateway path to be G-A-B-C and assigns channels (c_3, c_4) according to G's channel sequence. Since the same sequence is used in assigning channels, nodes within the same hop distance away from the same gateway radio end up using the same channels, thus realizing the desired ring-like configuration shown in Figure 3. Section 4 describes the distributed operations of ROMA in more details.

The simplified description of ROMA ignores many important details. First, how does each node choose paths consisting of high quality links with little external load? (Section 3.2 and Section 3.3). Second, each dual-radio node must operate its two radios at different frequency bands. How do we take into consideration of the topological differences between 802.11a and 802.11b/g channels? (Section 3.4). Third, in the presence of multiple gateways, how should each gateway choose its channel sequences to improve the aggregate throughput of the network? (Section 3.5). Last, how do we extend ROMA to work with more than two radios at each node? (Section 3.6).

3.2 Calculating link metric

The most popular link metric today is ETX [14, 15] and its extension ETT [7, 16]. Both metrics explicitly measure the delivery ratio of a link and ETT is a scaled version of ETX to adjust for different link level transmission rates. In our initial implementation of ROMA, we found that using ETT often leads to suboptimal routing paths. Our improved link metric incorporates two additional factors to estimate the quality of a link: link variation and external load.

Many links exhibit highly variable delivery ratios on short timescales. As a result, a path that should perform well judging from its individual links' ETTs often has low actual throughput because the delivery ratio of some link along the path happens to incur higher than average loss rate. We modify the way standard ETT is calculated to account for delivery ratio variations. In ROMA, each node keeps track of two variables, p_a and p_v, which are exponentially weighted moving averages (EWMA) of the average and mean deviation of the link delivery ratio. Our method for calculating p_a and p_v is inspired by the technique used for estimating RTT in TCP [18]. Specifically, let p be the latest measurement of the delivery ratio, p_a and p_v are updated as follows:

$$p_a = p_a + g \cdot (p - p_a)$$
$$p_v = p_v + g \cdot (|p - p_a| - p_v)$$

The parameter g represents the gain factor in EWMA and is set to 0.2 in our implementation. Intuitively, if the delivery ratio of a link has high variations as indicated by a large p_v, that link is more likely to exhibit much lower than average delivery ratio during the actual data transmissions. We penalize links with high variations by calculating ETT to be $\frac{1}{r} \cdot \frac{1}{(p_a - p_v)*(p'_a - p'_v)}$, where r is the link level transmission rate and p'_a and p'_v are the average and mean deviation of the link delivery ratio in the reverse direction. The higher the link variation, the larger the corresponding ETT metric.

The throughput of a link is reduced by competing traffic on the same channel. In a multi-radio mesh, a node can potentially find an alternative route on different channels with less competition. We explicitly measure the external load of a link by having each node continuously snoop the medium to record non-ROMA packets received, including those that fail the MAC-level CRC check. Based on the transmission rate and size of received packets, a node estimates the fraction of time a channel is occupied by external transmissions. Our calculation underestimates the actual external load since some interference (e.g. overlapping channel interference, non-802.11 interference) do not result in packet reception. Similar to the modified ETT, a node keeps track of both the average (L_a) and mean deviation (L_v) of its measured external load and calculates $L = L_a + L_v$ ($0 \le L \le 1$). The load of a link between two nodes is the maximum of the estimated external load (L) at both nodes.

In ROMA, the link metric is represented by a pair of values, (ETT, L), which collectively characterize the performance of a link due to loss and external load.

3.3 Choosing routes

In a multi-radio mesh, there is tension in choosing between shorter paths with lower total transmission overhead (i.e. smaller ETT sums) and longer paths consisting of better performing links on different channels. WECTT [16] and SIM [13] are two path metrics that resolve such tension by using a linear combination of the two factors. ROMA's path metric (M) extends the SIM metric [13] to take into account external load and is calculated as follows:

$$M = (1 - \beta) \cdot S + \beta \cdot T \tag{1}$$
$$where\ S = \sum_{i \in path} ETT(i) \tag{2}$$
$$T = max_s \sum_{i \in Seg_s} ETT(i) \cdot (1 + L(i)) \tag{3}$$

The extended SIM metric (M) is a linear combination of path overhead (S) and performance (T) with parameter β balancing the tradeoffs between them. The path overhead S is approximated by the sum of expected transmission time along the path. The path performance (T) is characterized by the estimated service interval of the bottleneck path segment and a smaller T corresponds to better performance. A path segment (Seg) consists of one or more links that interfere with each other on overlapping channels. In the common case, links operate on distinct channels along the gateway path and thus form path segments of length one.

The original SIM metric estimates the service interval of the bottleneck path segment using the sum of ETT along that segment [13]. In ROMA, the estimated service interval of a link is its ETT weighted by the observed external load, i.e. $ETT \cdot (1 + L)$. When the external load increases from zero to near 1, the estimated service interval doubles. This weighting approximates the expected service interval when there is a single competitor that transmits as fast as possible and the underlying MAC fairly divides the channel time among competing nodes. The weighting does not accurately capture the service interval for multiple external competing senders. Nevertheless, we find it has worked well for ROMA in practice.

Previous work [16, 13] choose the parameter β empirically. Here, we present an analysis that bounds the path performance (or overhead) for any given β. For any chosen path, its total transmission time (S) is greater than that of the worst path segment, i.e. $S \ge \max_s \sum_{i \in Seg_s} ETT(i) \ge T/2$. Furthermore, the total transmission time (S) cannot exceed the product of the total number of path segments (h) and the estimated service interval of the bottleneck segment (T), i.e. $S \le hT$. Therefore, we obtain $T/2 \le S \le hT$. Substituting this inequality back to Equation (1) and simplifying, we obtain the best and worse possible path performance (or overhead) for any M:

$$\left(1 - \beta + \frac{\beta}{h}\right) S \le M \le (1 + \beta) S$$

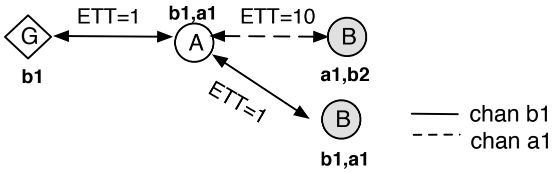

Figure 4: B has two possible assignments according to G's sequence: (a_1, b_2) and (b_1, a_1). B uses $(b_1, a1)$ even though this assignment causes intra-path interference because the quality of A-B link on channel b_1 is much better than that on channel a_1. We assume no external load.

$$(\frac{1+\beta}{2})T \leq M \leq ((1-\beta)h+\beta)T$$

The above inequalities make the tradeoff between path performance (T) and overhead (S) explicit. In a small mesh network where most paths have no more than 4 hops ($h \leq 4$), the performance of any chosen path is at least $\frac{1+\beta}{2(4-3\beta)}$ of the best performing path with an identical path metric. Likewise, the chosen path's overhead is at most $\frac{4(1+\beta)}{4-3\beta}$ times the lowest overhead. Since ROMA's primary goal is to achieve good path throughput, we set $\beta = 0.8$ so that any chosen path's performance is at least $\frac{1+0.8}{2(4-3*0.8)} = 0.56$ of the best possible performance. The corresponding path overhead will be no more than 4.5 times the lowest possible overhead.

3.4 Using 11b/g and 11a channels robustly

A compact dual-radio node must assign channels from different frequency bands in order to avoid the cross-channel interference that result from radios operating within close range of each other. In ROMA, such a constraint is easily achieved by letting each gateway use a sequence of alternating 802.11b/g and 802.11a channels, e.g. $b_1, a_1, b_2, a_2, ...$ As a result, each node along the gateway path can operate its two radios on a 802.11b/g channel and a 802.11a channel without interference.

There is considerable topological differences between 802.11a and 802.11b/g channels. In particular, links under 802.11b/g channels have much longer range than those on 802.11a channels due to the fact that 802.11b/g is at a lower frequency band and that 802.11b standard permits lower and more robust transmission rates (e.g. 1 or 2 Mbps). To operate robustly despite such topological differences, a node must be flexible in following a gateway sequence so that it is not forced to communicate using a very weak or non-existent link. Figure 4 gives an example. Node B is connected to A over a very weak link on channel a_1 but a perfect link on b_1. If it were to assign channels strictly according to G's channel sequence $(b_1, a_1, ...)$, it would be forced to communicate with A over the weak a_1 channel. ROMA deals with this situation by giving a node the alternative to assign identical channels as its previous hop neighbor.

The path metric in Equation(1) helps a node decide whether to advance to the next tuple in the channel sequence or to use the same assignment as its previous hop. When using the same assignment, a node's gateway path incurs intra-path interference, resulting in a bigger T. In Figure 4, successive links G-A and A-B operate on the same channel if B's channel assignment is (b_1, a_1), thus $T = 2$ and $M = 0.1 * S + 0.9 * T = 2$. On the other hand, if B is to use (a_1, b_2), $T = \max(1, 10) = 10$ and M=10.1. Because the first assignment achieves a gateway path with smaller path metric, node B ends up using channels (b_1, a_1).

3.5 Choosing channel sequences

The simplified design in Section 3.1 focuses on a single gateway with one radio. In the single gateway radio case, ROMA's main objective is to eliminate intra-path interference along all gateway paths. In scenarios where a gateway has more than one radio or multiple gateways co-exist, gateways should choose channel sequences carefully to reduce inter-path interference.

Aggregate throughput is improved when flows destined for different gateways' radios utilize different channels. A single gateway with two radios uses two sequences $b_1, a_1, b_2, a_2, b_3, a_3$ and $a_2, b_3, a_3, b_1, a_1, b_2$ so that flows with three or fewer hops destined for different radios at the same gateway do not interfere. When there are multiple gateways, it is impossible to assign sequences so that flows using different sequences do not interfere because there are only three non-overlapping 802.11b/g channels. We use a simple heuristic targeted at minimizing the interference of first hop transmissions at different gateway radios since the first hop performance is typically the bottleneck with multiple flows. Upon startup, a gateway node scans all three 802.11b/g channels for a period of time to learn of existing gateways' channel assignments. Subsequently, it chooses gateway sequences whose first hop channels differ from those chosen by potentially interfering gateways. For example, if a new gateway overhears the channel sequences of another gateway as $b_1, a_1, b_2, a_2, b_3, a_3$ and $a_2, b_3, a_3, b_1, a_1, b_2$, it will choose to use sequences $b_2, a_4, b_3, a_5, b_1, a_6$ and $a_5, b_1, a_6, b_2, a_4, b_3$. Since a gateway largely ignores the actual network topology in choosing channel sequences, the resulting channel sequence may not be optimal. Nevertheless, we find that this simple approach works well in practice.

In deployed mesh access networks, each node acts as both a mesh forwarder and an AP for unmodified single-radio clients. It is desirable for the link between a client and its associated AP to follow the corresponding gateway channel sequence to avoid intra-path interference. If other words, if a mesh node operates on channels (a_1, b_1) and uses channel a_1 for its gateway path, it should make its clients preferentially associate itself on channel b_1 instead of a_1. This can be achieved in the case of unmodified clients using techniques similar to those described in [28].

3.6 Extending beyond dual-radio nodes

ROMA can be extended naturally to handle m-radio nodes ($m \geq 3$). One can build such nodes today by separating antennas far apart enough to avoid interference. In the future, extra bands other than the 2.4 GHz and 5 GHz ranges could open up for use by commodity 802.11 radios so that one can build compact m-radio nodes.

For a compact m-radio mesh, a gateway sequence should be made up of channels from alternating bands such that each node will assign channels from distinct bands. To extend ROMA to the m-radio case, we follow a channel assignment approach similar to the dual-radio case. Given a gateway channel sequence, if a node's best gateway path uses channel c_i to its previous hop neighbor, it would assign channels $c_{i-m/2}, .., c_i, .., c_{i+m/2}$ to each of its radios based on the gateway sequence. Such an assignment preserves dense network connectivities among nodes that follow the same gateway sequence. While this is a natural extension to the multi-radio case, one unaddressed challenge is how such an extension compares with alternative assignment strategies such as letting each node follow multiple gateway sequences.

4. DISTRIBUTED OPERATIONS

The previous section describes ROMA in a static setup. The simplicity of the design allows for a robust distributed implementation that can continuously adapt to changing network topologies due to channel condition changes and node churn. This section presents the distributed operations of the ROMA protocol. In Section 4.1, we describe how a node updates its routing tables, assuming it has already found its gateway path and assigned channels correspondingly. In Section 4.2, we show how a node can efficiently explore its neighborhood to find a better gateway path when its current choice of gateway path and channel assignment has become suboptimal due to topological changes.

4.1 Link measurement and route propagation

Once a node has found its best gateway path, it switches to the assigned channels, called *home* channels, according to the gateway's sequence. Each node continuously monitors its neighboring links' conditions on the home channels and propagates its current gateway path and path metric to keep other nodes updated.

Link quality measurement: A node can directly communicate with a subset of all its potential neighbors (i.e.*actual* neighbors) whose home channels overlap with the node's own. The previous hop of a node's gateway path is one such actual neighbor. Each node snoops the medium to measure external load and periodically broadcasts probe packets every T_{bcast} seconds from all radio interfaces to measure the delivery ratio of links to its actual neighbors. Instead of broadcasting probes one at a time, as suggested in [14, 7], the delivery ratio is measured more accurately when nodes send a burst of probe packets in quick succession. We have noticed that, for some links in our testbed, the loss rate observed on a long burst of transmissions is much higher than that of a very short burst and stabilizes when the burst length exceeds 20 packets. As a result, broadcasting probes one at a time causes a node to dramatically over-estimate the actual delivery ratio during data traffic forwarding. We were able to deterministically reproduce this observation on CMU's wireless emulation testbed [19], suggesting that delivery ratio differences due to burst size could be a common problem, at least for Atheros-based WiFi cards. We find that using a burst size of 20 packets results in robust delivery ratio measurements.

Route propagation: Each node periodically announces its current channel assignment, gateway path and the corresponding gateway sequence on its home channels. The gateway path consists of a series of forwarding nodes, their channel assignments as well as the link metrics (in terms of ETT and external load) between successive forwarders. A gateway node announces a path of length zero to indicate its gateway status. A node without any gateway path announces a path length of ∞ on the default 802.11a and 802.11b/g channels.

A node processes received advertisements from all interfaces and stores extracted node and link information in a partial local link table [15, 7]. The link table contains the list of known nodes with their corresponding home channels as well as the link metrics between them. In order to distinguish new information from old ones, each node or link entry in a route announcement is associated an increasing sequence number generated by the originating node.

A node continuously updates the path metric of its current gateway path. Furthermore, if a node finds a better gateway path in its link table that does not require it to change home channels, it immediately switches to the new route.

Algorithm 1 Select the most promising neighbor for investigation every T_{invs}.

```
for all $x$ of my potential neighbors do
  for $c \in x$'s home channels do
    if $c \notin$ my channel assignments then
      $m \leftarrow$ metric of link between me and $x$ on channel $c$
      if $m$ does not exist then
        $m \leftarrow min$ //assume the best link metric
      else
        //if link is measured before, weigh its metric by age
        $m \leftarrow m + (min - m) * 0.1$
      end if
      estimate my path metric via $x$ using $x$'s gateway path metric and $m$
      remember $(x, c)$ corresponding to the best path metric
    end if
  end for
end for
if the best estimated path metric is less than a threshold of my current gateway metric then
  investigate the link to $x$ on channel $c$.
end if
```

4.2 Discovering better routing paths

As the underlying network topology changes, a node may need to use different home channels for its best gateway path. To converge to the best possible channel assignment and gateway path, a node performs temporary channel switches to explore alternative gateway paths efficiently.

The goal of temporary channel switch is to investigate the quality of a specific link on a foreign channel with the hope that it will yield a better gateway path than the current one. To maximize the chances that a temporary switch finds a better path, a node chooses the most "promising" link among all potential neighbors for investigation. Two types of information are needed in order to choose the most promising link: the current home channels of potential neighbors and their gateway path metrics. ROMA employs network-wide gossip to inform each node of its potential neighbors' channel assignments as well as their gateway path metrics. Algorithm 1 gives the pseudo-code for finding the most promising link for investigation. Every T_{invs} seconds, a node estimates the link delivery ratio to a neighbor (x) on one of x's home channels. If the node has never investigated that link before, it optimistically assumes the best link metric. Otherwise, it uses the existing link metric discounted by the age of that information. The node then uses the estimated link metric and x's gateway path metric to estimate the metric of the potential gateway path via x. The most promising link for investigation is one with the best estimated gateway path metric. If the estimated path metric is less than a threshold of the node's current gateway path, the node starts the actual investigation by switching to the foreign channel.

Notation	Meaning	Default value
T_{bcast}	Periodic probe/route broadcast interval	15 sec
N_{probe}	The burst length of probes	20 pkts
T_{invs}	Periodic investigation interval	150 sec
T_{dur}	The amount of time a node dwells on a foreign channel	0.1 sec

Figure 5: ROMA's protocol parameters and their default values.

Upon switching to the foreign channel, a node immediately sends a burst of N_{probe} broadcast probes destined for the neighbor x under investigation. The neighbor x replies with N_{probe} probes of its own upon receiving the investigative probes. A node waits for no more than T_{dur} seconds on the foreign channel to wait for the neighbor's replies before switching back to its home channel. When the investigation is finished, a node computes its best gateway path by taking into account the new link information and changes its home channels if the new path requires a different channel assignment. We use a modified Dijkstra's algorithm to search for the best path in a link table, similar to that in [13].

Figure 5 summarizes the list of protocol parameters in ROMA and their default values used by our prototype. These parameters control the tradeoff between how quickly ROMA can discover the best channel assignment and its overhead during normal operations. Channel re-assignment and temporary switches are potentially disruptive for ongoing multi-hop flows. Fortunately, these are infrequent events in our testbed (Section 6.6).

5. IMPLEMENTATION

We have implemented ROMA using the Click modular router toolkit [21] as a Linux kernel module. Our implementation re-uses part of the software infrastructure originally developed for the Srcr [7] routing protocol. The kernel module directly invokes the functions exported by the wireless driver to change channels. We use the Madwifi 0.9.3.3 driver [1] with a bug-fix to ensure that channels are changed upon command. Upon receiving the command to change channel, the driver waits until its current transmit queue is drained by the radio hardware before switching to the new channel. ROMA uses source routing when forwarding data traffic. The sender specifies the sequence of forwarders and the channels to be used, and intermediate nodes in the path forward packets based on the specified source route.

6. EVALUATION

This section demonstrates the following points about the performance and behavior of ROMA:

1. ROMA eliminates intra-path interference and the throughput of multi-hop gateway paths is comparable to that of single hop paths (Section 6.2).
2. ROMA achieves good aggregate throughput in the presence of many active flows and multiple gateways by utilizing many non-overlapping channels within the same physical area. (Section 6.3).
3. Incorporating link variation and external load in the path metric helps ROMA choose better multi-hop routing paths (Section 6.4, 6.5).

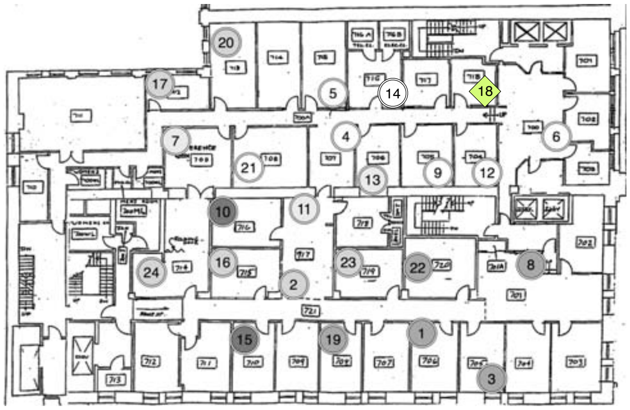

Figure 6: Node placement in the testbed. The top and bottom half of the offices are separated by thick concrete walls that greatly attenuate the signals. An example set of gateway paths and channel assignment is shown with node-18 as the gateway. Darker nodes indicate longer gateway paths.

4. ROMA chooses stable channel assignments and gateway paths while remaining adaptive to changes in the underlying topology. By preserving cross-links among paths to the same gateway radio, a node can sometimes change its gateway path without changing channel assignment (Section 6.6).

6.1 Experimental Setup

All our experiments are conducted on a 24-node dual-radio testbed. The testbed spans a single floor of a tall office building in densely populated downtown Manhattan. The placement of nodes in our testbed is shown in Figure 6. There are many visible 802.11b/g access points and a few 802.11a access points. Although there is rarely traffic of significant volume on 802.11a channels, we have frequently observed high volumes of background traffic on 802.11b/g channels, often keeping the channel busy for up to 85% of the time.

Each testbed node is a small form-factor device the size of a typical home wireless router. Each node is powered by a Geode processor with 128 MB RAM and has two 802.11a/b/g radios with the Atheros 5212 chipset. All nodes run Linux with Click kernel patch (2.6.19). Since we purchased the nodes in two batches over the course of a year, half of the 24 nodes are slow with 233 MHz processors while the other half are fast with 498 MHz processors. The forwarding capacity of slow nodes is limited by the CPU and saturates at only slightly over 6 Mbps. To avoid CPU bottlenecks, we use autorate adaptation [6], but fix the maximum link level transmission rate to 6 Mbps. Since the slowest transmission rate of 802.11a is 6 Mbps, this effectively disables rate adaption for radios operating on 802.11a channels. We plan to upgrade the slow nodes in the future to lift the maximum rate restriction.

In all experiments, a gateway uses channel sequences (40,6,50,11,60,1) and (11,60,1,40,6,50) with its radios operating on channels 40 and 11, unless specified otherwise. By default, we measure the throughput of UDP flows using the `iperf` tool. With a maximum link level transmission rate of 6 Mbps, the maximum achievable end-to-end path throughput is 5.5 Mbps, due to 802.11 protocol overheads. We con-

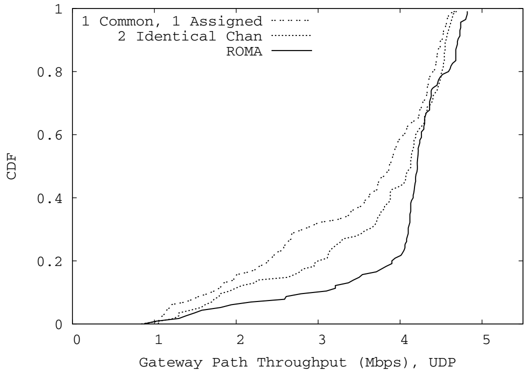

Figure 7: The cumulative distribution of gateway path UDP throughput in ROMA as compared to the configurations using 2 identical channels or 1 common, 1 assigned channels. The experiments use a single gateway and activate one UDP flow at a time.

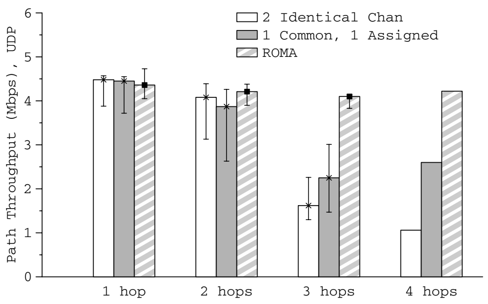

Figure 8: The median UDP throughput of gateway paths with different hopcounts for the same experiments in Figure 7. The errorbars correspond to 20% and 80% throughput. No error bars are plotted for 4-hop paths because there were only a few of data points.

figure the UDP source to send at a fixed rate of 6 Mbps with 1300-byte packets for a duration of 120 seconds and measure the end-to-end throughput (i.e. goodput) of UDP flows.

We compare ROMA against two alternative channel assignment strategies, *2 identical channels* and *1 common, 1 assigned channel*. In the 2 identical channels configuration, all nodes use two fixed channels (channel 40 and 6). Due to its simplicity, the identical channels configuration is frequently used in many multi-radio mesh deployments [16, 13]. We model the 1 common, 1 assigned channel configuration after the proposal in [20]. Each node operates one of its radios on the common channel (channel 40). A node assigns the other radio to the 802.11b/g channel (channel 1, 6 or 11) with the fewest interfering nodes so long as one of the node's immediate neighbors is also on the chosen channel. When links are lossy, the notion of neighbors is fuzzy; we consider two nodes as immediate neighbors if the link between them has greater than 50% delivery ratio in both directions. Like [20], we assume nodes within three hops of a node x interfere with x. In our testbed, the three hop neighborhood of a node covers most of the network. We did not implement the the distributed channel assignment protocol of [20]; instead, we measure the network topology on the common channel 40 and use it to calculate each node's channel assignment offline before the start of each experiment. In both identical and common channel configurations, nodes still rely on ROMA's path metric to find the best gateway paths.

6.2 Single flow throughput

We first examine the performance of individual gateway paths. In each run of the experiment, we randomly pick one of the 24 nodes to act as the gateway. We start ROMA on all nodes at the same time and wait for 5 minutes for the protocol to converge. Typically, a node finds its best gateway path quickly, within one or two investigative switches. In each experiment, we initiate a UDP flow from the gateway to each of 23 non-gateway nodes, one at a time. There were five experiment runs in total.

Figure 7 shows the CDF of the path throughput in ROMA compared to that in the identical and common channel configurations. The median path throughput of ROMA is 4.2 Mbps, similar to that of the identical channel configuration (4.1 Mbps), and better than that of the common channel configuration (3.8 Mbps). Since most gateway paths consist of only 1 or 2 hops, the median throughput is high because two hop paths do not suffer from intra-path interference in all configurations. Figure 7 also shows that ROMA achieves significantly higher throughput in the lower percentiles. In particular, the 20-percentile throughput of ROMA is 4 Mbps, compared to 3.1 Mbps for the identical channels configuration and 2.5 Mbps for the common channel configuration.

In Figure 8, we examine how throughput differs for paths with different hopcounts. Unlike the identical or common channel configurations, there is little throughput degradation in ROMA as hopcount increases because ROMA assigns non-overlapping channels along all links in a gateway path. Even for paths with 3 or more hops, ROMA achieves a median of 4.1 Mbps, a 7% drop in performance compared to that of single hop paths (4.4 Mbps). In the identical channel configuration, any path with 3 or more hops suffers from intra-path interference, causing the throughput to be reduced by more than half. In the common channel configuration, some 3-hop paths consist of only one link on the common channel, thus avoiding intra-path interference. However, since the network has the densest connectivity on the common channel, the majority of three hop paths require 2 links on the common channel and thus suffer from intra-path interference.

Interestingly, we observe that the average and 20-percentile throughput of 1 and 2-hop paths in ROMA are better than that in the identical and common channel configurations, as shown in Figure 8. For example, the 20-percentile throughput of 2-hop paths in ROMA is 3.9 Mbps, compared to 3.1 Mbps for the identical channel configuration. This is because with the identical channel configuration, many nodes choose 1 or 2-hop gateway paths involving links with mediocre delivery ratios. There exist alternative 3 or 4-hop paths over links with high delivery ratios, but these paths are not chosen because they involve links that interfere with each other on the same channel, therefore resulting in worse path metrics. We also notice that the 1 and 2-hop paths in the identical channel configuration outperform those in the common channel configuration. When all nodes use 2 identical channels, the network is densely connected with many high quality links to choose 1 and 2-hop paths from. Since ROMA assigns channels according to a

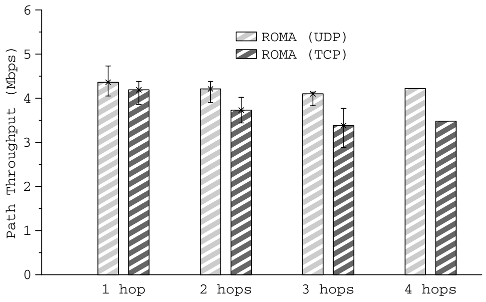

Figure 9: The median TCP vs UDP path throughput for the single gateway, single flow experiments. Error bars correspond to 20% and 80% throughput. No error bars are plotted for 4-hop paths because there were only a few data points.

node's best gateway path, ROMA retains those high quality links necessary for constructing good gateway routes. By contrast, since the common channel configuration in [20] is designed to optimize for the all-pairs traffic pattern, nodes assign channels independently of any routing paths, resulting in the loss of some high quality links that are useful for choosing good gateway routes.

TCP Throughput: We repeat the same single-flow experiment to evaluate the throughput of TCP flows. Achieving high throughput over multi-hop wireless paths is known to be difficult, since the increase in RTT variance as well as the increased loss rates has a detrimental impact on TCP [16, 5]. However, from Figure 9, we observe that TCP flows in ROMA achieve only marginally lower throughput when compared to the UDP flows, even for longer paths. For example, in 3-hop (and 4-hop) paths, the median TCP throughput is 3.38 Mbps (and 3.48 Mbps), which is merely 17% lower when compared to the median UDP throughput of 4.10 Mbps (and 4.22 Mbps). This shows that ROMA is able to identify paths that consistently exhibit high performance, leading to high throughput over UDP as well as TCP.

Route stretch: ROMA uses the parameter β to balance the tradeoff between path performance and overhead. We set up one node (node-18) as the gateway and vary β to study its effect on the average path length as well as the single-flow throughput from the gateway to each of the 23 non-gateway nodes. We use three different values of β $(0.7, 0.8, 0.9)$ and run the experiment 3 times for each of them. Across the 3 runs, the average path length is 1.98, 2.1 and 2.63, and the total number of 4+ hop paths is 5, 12 and 21, for β=0.7, 0.8 and 0.9 respectively. We observe that as β is decreased, many of the high-performance 4+ hop paths are replaced by sub-optimal 3 hop paths. In particular, the average 3-hop path throughput is 3.8, 3.98 and 4.29 Mbps for β=0.7, 0.8 and 0.9 respectively. In summary, higher values of β improve performance at the cost of increased path overhead, while smaller β values tend to penalize longer paths at the cost of lower performance.

6.3 Aggregate throughput of multiple flows

In addition to achieving better single flow performance, ROMA also improves the aggregate performance of multiple flows by utilizing multiple non-overlapping channels within the same physical region. We measure the aggregate throughput of multiple randomly chosen simultaneous flows for two network configurations: (a) 1 gateway, 3 flows; (b) 3 gateways, 9 flows. In each experiment, we start ROMA on all nodes and wait for 5 minutes to allow the routes to stabilize, before starting the flows. We repeat each set of experiments 60 times, using randomly selected gateways and traffic sinks. In experiments with one gateway, the gateway sets its channels to (40,11). In the 3-gateway case, the channels used by the gateways are (40,11),(60,6) and (50,1).

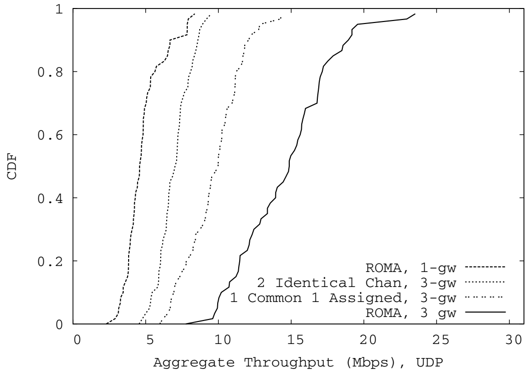

Figure 10: Aggregate UDP throughput for networks with 1 or 3 gateways.

Figure 10 shows the cumulative distribution of the aggregate throughput of ROMA. In the 1-gateway experiment, the median aggregate throughput is 4.58 Mbps, while in more than 90% of the runs, the throughput is greater than 3.46 Mbps. In the 3-gateway experiment, the median aggregate throughput further increases to 14.8 Mbps, with more than 90% of the runs resulting in an aggregate throughput of more than 10.7 Mbps. However, we notice that in a large fraction of runs, the set of randomly chosen flows do not always utilize both radios on all of the gateways. For example, in many runs of the 1-gateway 3-flow experiment, all 3 flows route to the same radio on the gateway. This explains why the median throughput in the 1-gateway, 3-flow experiment is less than two times the median throughput of the single-flow case. One potential improvement is to explicitly balance the routes chosen to different gateway radios, as is done in [26].

We repeat the 3-gateway, 9-flow experiment with the identical and common channel configurations. The median aggregate throughput in the identical channel configuration is 7 Mbps while with the common channel configuration, the aggregate throughput improves to 10 Mbps. This is because the 3 gateways in the common channel configuration utilize all 3 orthogonal 802.11b/g channels as compared to only 1 802.11b/g channel in the case of the identical channel configuration. Since ROMA uses 3 802.11b/g and 3 802.11a channels among 3 gateways, as opposed to only 1 802.11a channel in the common channel configuration, it achieves the highest aggregate throughput. In particular, ROMA's median aggregate throughput is 1.4× and 2.1× that of the common and identical channel configurations, respectively.

The presence of multiple gateway radios on non-overlapping channels is not the only reason for increased aggregate throughput. As the number of gateways increases, the average path length between a node and its nearest gateway decreases significantly, thereby reducing path overhead and increasing aggregate throughput. In particular, for the

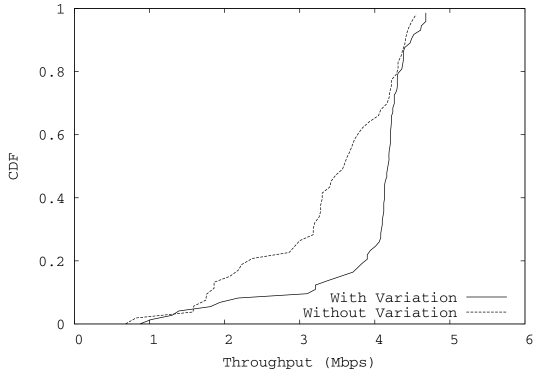

Figure 11: UDP throughput for paths with $\geq$ 2 hops for the single gateway, single flow experiments. Incorporating link variation improves performance significantly over many paths.

3-gateway case, a majority of gateway paths are 1-hop paths. Since the gateways use different channels for their radios, flows destined to different gateway radios do not interfere with each other, leading to high aggregate throughput.

6.4 Effect of incorporating link variation

As described in Section 3.2, we penalize links that exhibit highly variable losses by incorporating the mean deviation of measured delivery ratios in the ETT calculation.

In order to understand the effect of incorporating variance in the link metric, we repeat the single flow experiments using the original ETT metric without the deviation penalty. Since 802.11b/g links exhibit greater variability than 802.11a links, we consider only those paths that are of length ≥ 2, since they contain at least one 802.11b/g link. Figure 11 shows that incorporating variation improves the path throughput for a significant fraction of paths, because it enables ROMA to choose better and more stable paths. In particular, 85% of 2+ hop paths using the variation-aware metric achieve throughput of greater than 3.5 Mbps, while only 50% of 2+ hop paths achieve more than 3.5 Mbps without incorporating variation. Similarly, 75% of 2+ hop paths using the variation-aware metric achieve throughput of greater than 4 Mbps, while only 34% of 2+ hop paths achieve more than 4 Mbps without incorporating variation. We further observe that links with almost perfect delivery ratios tend to exhibit relatively low variability. As a result, a number of high throughput paths consisting of links with low variability are coincidentally chosen by variance unaware ROMA as well. The effect of incorporating variation is more pronounced for paths that need to use links with intermediate delivery ratios, since many of these links also tend to be highly variable. In summary, our results show that delivery ratio variation is an important consideration for choosing stable, predictable and high performance routes.

6.5 Effect of incorporating external load

To study the effect of external load on ROMA, we use a laptop as a controlled interference source to generate external load. We conduct this experiment in the middle of the night, where the measured real background traffic is negligible. We set up one gateway (node-10) and configure the other nodes as non-gateways. We start the interference source to transmit on channel 6 and measure the single flow UDP throughput from each node to the gateway. We vary the external load to occupy 10%, 40% and 100% of channel time and run the experiments with and without incorporating load into the path metric.

Figure 12, 13 and 14 compare the performance of load-aware ROMA with load-unaware ROMA, for different degrees of load. We observe that at low loads (10%), load-aware ROMA shifts some gateway paths to alternate unloaded paths, resulting in a small improvement in performance. At moderate load (40%), some links become lossier and as a result, even load-unaware ROMA shifts some paths. However, those paths whose delivery ratios do not change much remain on the loaded path. Load-aware ROMA, on the other hand, has completely avoided the loaded links in all paths, resulting in 70% increase in average path throughput over load-unaware ROMA. Finally, when the load is near saturation, the delivery ratios of most links on channel 6 deteriorate drastically, causing even load-unaware ROMA to switch all its paths away from saturated links. Under saturating load, both load-aware and load-unaware ROMA end up picking the same paths that avoid the saturated channel; hence, both achieve similar performance. In summary, load-awareness leads to significant performance improvements under moderate external loads.

6.6 Channel assignment over time

To understand how stable channel assignments are, we monitor the progress of ROMA over a 5-hour period with node-16 as the gateway. Figure 15 shows the route and channel changes in ROMA for every node. Over the 5-hour experiment, there are 76 route changes among the 23 non-gateway nodes, and 59 of these route changes involve channel changes as well. Upon startup, all non-gateway nodes perform at least one investigation and find their gateway paths. For a majority of nodes, their channel assignments and gateway routes remain stable. For example, nodes 2,10,19,21,22 do not change their gateway paths or channels after the initial route calculation. Some nodes (e.g. node 8) do not change channels after the initial assignment, but recalculate better gateway routes on the same channels. This demonstrates ROMA's advantage of maintaining cross-links in the topology: alternate routes can be found without requiring expensive channel changes.

We also find that some nodes (e.g. nodes 6,11,12,13,14) change their routes and channels twice within the span of a few minutes. This behavior has two causes: Initially when node A investigates node B on channel *c*, A has no information about the variation and load of link A-B, since *c* is a foreign channel. It optimistically assumes that the delivery ratio deviation is 0 and recalculates a new gateway route using channel *c*. After *A* has dwelled on channel *c* for a while, it learns that the link actually exhibits high variation or high load, and therefore, the metric of the new path degrades. As a result, *A* might perform another investigation and find a better path on a different channel. To prevent flapping, *A* remembers the measured link variation and load, so that it does not again incorrectly assume that the variability of the link is 0. Another cause is that when a node A changes its channels, it occasionally learns of additional potential neighbors, causing it to discover even better alternate routes on a different channel.

Overall, we observe that ROMA's channel assignment is stable over fairly long periods of time. By preserving cross-

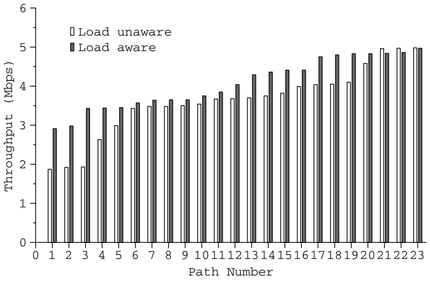

Figure 12: At 10% load, load-aware ROMA improves performance slightly.

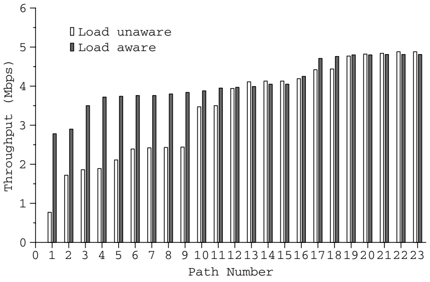

Figure 13: At 40% load, load-aware ROMA improves performance significantly.

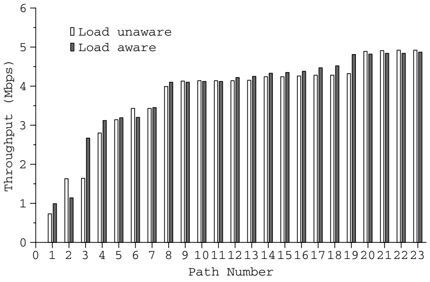

Figure 14: At 100% load, load-aware and load-unaware ROMA choose identical paths.

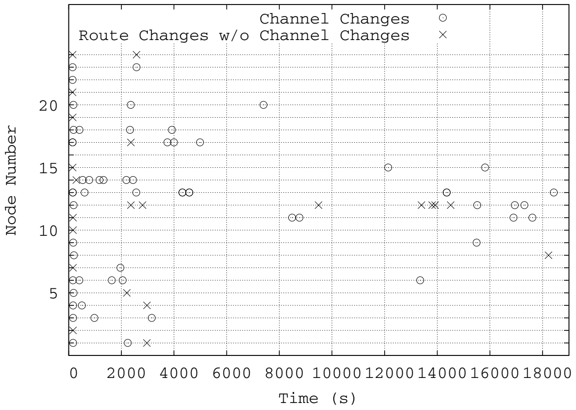

Figure 15: Route and channel changes over time

links, a node can sometimes switch to better new routes to cope with network topology changes without having to change its channel assignment.

7. RELATED WORK

ROMA builds upon a large body of prior work in multi-radio mesh protocols. Below, we summarize related work and point out their relationship to ROMA.

Centralized channel assignment: Centralized solutions aim to find the best combination of routes, channel assignment and transmission schedules given the network topology on *all* channels and the traffic pattern. Most centralized optimizations[4, 27, 25, 24, 33, 17] are evaluated in simulations and lack practical solutions for coordinating topology measurement and disseminating channel assignments.

In [30, 31], a centralized channel assignment algorithm (TIC) is implemented and evaluated on a 20-node testbed. TIC requires that all nodes operate one of its radios on a common default channel in order to coordinate topology measurement and disseminate channel switch commands. The channel assignment algorithm in [30] always takes into account of external load estimated using received packets.

Distributed channel assignment: Most practical multi-radio deployments choose not to perform sophisticated channel assignment but use identical channels on all nodes [16, 13, 8]. Assigning channels dynamically in a distributed fashion is a hard problem and only two known protocols do so [20, 32]. In [20], the authors propose a distributed channel assignment protocol that relies on a common channel across the network to ensure connectivity. Each node runs the distributed assignment protocol to select a channel for its other radio. The assignment prefers channels that are least used by a node's interfering neighbors. Routing is performed independently of channel assignment using the MR-LQSR protocol [16] with the WECTT metric. The protocol has been implemented and evaluated on a dual-radio testbed of 14 nodes. In comparison, ROMA does not require a common channel and can use channels more efficiently to eliminate intra-path interference and improve aggregate throughput.

Hyacinth [32] is a distributed assignment protocol that is closest in spirit to ROMA. Hyacinth explicitly builds a spanning tree rooted at each gateway node where each node independently chooses a channel for one of the radio interfaces to communicate with its children. Like ROMA, Hyacinth does not rely on a common channel to keep the network connected and optimizes channel assignments along routes between mesh nodes and a few gateways. ROMA differs from Hyacinth in several ways. One fundamental difference is that Hyacinth does not consider link losses and loss fluctuations, one of the most important factors affecting throughput in mesh networks. Routes consisting of highly lossy and fluctuating links are bound to perform poorly and Hyacinth cannot adapt to changing channel conditions other than node failures. Furthermore, Hyacinth's join/leave protocol for spanning tree construction can be fragile in lossy environments as it requires reliable delivery of protocol messages and accurate detection of node failures. Since Hyacinth has been primarily evaluated in simulations and a 9-node controlled testbed where there is no reported link loss, it is unclear how robust its performance is in real deployments. By contrast, ROMA explicitly incorporates link loss, loss variations and external traffic load in the link metric and can quickly adapt to changing channel conditions.

Route selection: WECTT [16] and SIM [13] are two path metrics that help routing protocols preferentially choose routes with less intra-path interference. ROMA uses the SIM metric for choosing paths and extends it to take into account link variations and external load.

While our modification of the ETT metric is similar in spirit to the mETX metric [22], there exists subtle differences. While mETX proposes a more accurate model to estimate the expected transmission count of a single packet for time-varying links, it captures the average-case scenario. By contrast, our modification is closer to model the worst-case scenario than the average case. The rationale for modeling the worst-case scenario follows from ROMA's goal of choosing routes where each link delivers good and predictable performance. In addition, computing mETX requires bit-level loss information from all corrupted packets. However, most packet losses do not result in any (corrupted) packet reception in our testbed.

8. CONCLUSIONS

Designing a high-performance multi-radio protocol faces many practical constraints and challenges (small node size, highly fluctuating link qualities, external load). ROMA is a distributed protocol that performs routing and channel assignment to achieve high end-to-end performance in a dual-radio mesh by eliminating intra-path interference and reducing inter-path interference. ROMA finds high-performance multi-hop paths by leveraging a new path metric that incorporates link variations and external load. ROMA also adapts well to network topology changes while choosing stable routing paths.

Acknowledgments

David Bindel contributed to an early design and Arthur Meacham wrote a simulator for ROMA and helped set up the testbed. We are particularly indebted to the Meraki Networks engineers, especially John Bicket and Sanjit Biswas, who provided invaluable advice and encouragement. We thank many people who have helped us improve this work: Frank Dabek, Robert Morris, Yair Sovran, Michael Paik, our shepherd Dina Katabi and the anonymous reviewers.

9. REFERENCES

[1] Madwifi. `http://sourceforge.net/projects/madwifi`.

[2] A. Adya, V. Bahl, J. Padhye, A. Wolman, and L. Zhou. A multi-radio unification protocol for ieee 802.11 wireless networks. In *IEEE BroadNets*, 2004.

[3] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. A measurement study of a rooftop 802.11b mesh network. In *Proc. ACM SIGCOMM Conference (SIGCOMM 2004)*, September 2004.

[4] M. Alicherry, R. Bhatia, and L. Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *Proc. ACM International Conference on Mobile Computing and Networking (MobiCom)*, 2005.

[5] H. Balakrishnan, V. N. Padmanabhan, and R. H. Katz. The effects of asymmetry on tcp performance. *Mobile Networks and Applications*, 4(3), Sept. 1999.

[6] J. Bicket. Bit-rate selection in wireless networks. Master's thesis, Massachusetts Institute of Technology, February 2005.

[7] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *ACM Mobicom*, 2005.

[8] S. Biswas. Meraki networks' next generation multi-radio mesh platform, 2008. private communication.

[9] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. In *ACM SIGCOMM*, 2005.

[10] S. Chachulski, M. Jennings, S. Katti, and D. Katabi. Trading structure for randomness in opportunistic wireless routing. In *ACM SIGCOMM*, 2007.

[11] C. Cheng, P. Hsiao, H. Kung, and D. Vlah. Adjacent channel interference in dual-radio 802.11 nodes and its impact on multi-hop networking. In *IEEE Global Telecommunications Conference (GLOBECOM)*, 2006.

[12] S. M. Das, H. Pucha, K. Papagianakki, and Y. C. Hu. Understanding wireless routing link metric dynamics. In *Proceedings of the 7th ACM SIGCOMM/USENIX Internet Measurement Conference*, 2007.

[13] S. M. Das, Y. Wu, R. Chandra, and Y. C. Hu. Context based routing: Technique, applications and experience. In *USENIX NSDI*, 2008.

[14] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proceedings of the 9th ACM International Conference on Mobile Computing and Networking (MobiCom '03)*, San Diego, California, September 2003.

[15] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. In *Proc. ACM SIGCOMM Conference (SIGCOMM 2004)*, September 2004.

[16] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*, 2004.

[17] P. Dutta, S. Jaiswal, and R. Rastogi. Routing and channel allocation in rural wireless mesh networks. In *INFOCOM*, 2007.

[18] V. Jacobson. Congestion avoidance and control. In *ACM SIGCOMM*, 1988.

[19] G. Judd and P. Steenkiste. Using emulation to understand and improve wireless networks and applications. In *USENIX NSDI*, 2005.

[20] B. Ko, V. Misra, J. Padhye, and D. Rubenstein. Distributed channel assignment in multi-radio 802.11 mesh networks. In *WCNC*, 2007.

[21] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click modular router. *ACM Transactions on Computer Systems*, 18(4), November 2000.

[22] C. Koksal and H. Balakrishnan. Quality-aware routing metrics for time-varying wireless mesh networks. *IEEE Journal on Selected Areas in Communications*, 24(11), 2006.

[23] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking*, 2001.

[24] M. Marina and S. Das. A topology control approach for utilizing multiple channels in multi-radio wireless mesh networks. In *Proceedings of Broadnets*, 2005.

[25] S. Merlin, N. H. Vaidya, and M. Zorzi. Resource allocation in multi-radio multi-channel multi-hop wireless networks. In *UIUC Technical Report*, 2007.

[26] V. Mhatre, F. Baccelli, H. lundgren, and C. Diot. Joint mac-aware routing and load balancing in mesh networks. In *ACM CoNext*, 2007.

[27] A. Mohsenian-Rad and V. Wong. Joint logical topology design, interface assignment, channel allocation and routing for multi-channel wireless mesh networks. In *INFOCOM*, 2008.

[28] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill. Designing high performance enterprise wi-fi networks. In *USENIX NSDI*, 2008.

[29] A. Prabhu, H. Gupta, and S. Das. Minimum-interference channel assignment in multi-radio wireless mesh networks. In *IEEE SECON*, 2007.

[30] K. Ramachandran and M. M. B. Elizabeth Belding, Kevin C. Almeroth. Interference-aware channel assignment in multi-radio wireless mesh networks. In *IEEE Infocom*, 2006.

[31] K. Ramachandran, I. Sheriff, E. Belding, and K. Almeroth. A multi-radio 802.11 mesh network architecture. *Mobile Networks and Applications*, 13(1-2):132–146, 2008.

[32] A. Raniwala and T. cker Chiueh. Architecture and algorithms for an IEEE 802.11-based multi-radio wireless mesh networks. In *IEEE Infocom*, 2005.

[33] A. Raniwala, K. Gopalan, and T. cker Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. 8(2), 2004.

[34] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a multi-radio mesh networking testbed. In *1st Workshop on Wireless Network Measurements (Winmee)*, 2005.