RDCS: Routing Driven Channel Selection in Multi-Radio Ad-hoc Networks

David Murray
Murdoch University
D.Murray@murdoch.edu.au

Michael Dixon
Murdoch University
M.Dixon@murdoch.edu.au

Terry Koziniec
Murdoch University
T.Koziniec@murdoch.edu.au

ABSTRACT

Despite over a decade of research in multi-hop ad-hoc networks, a fundamental performance limitation remains largely unsolved. Significant performance problems occur when transmissions are hopped over multiple wireless nodes utilizing the same frequency. Although this is a contention problem, it is exacerbated by numerous other ad-hoc specific issues. The solution, which has been attempted at many layers in the ISO networking model, is to turn multi-hop ad-hoc networks from single channel networks into multi channel networks. RDCS (Routing Driven Channel Selection) is a channel selection mechanism that operates with the OLSR (Optimized Link-State Routing) protocol. It circumvents multi-hop performance problems by enabling multi-radio 802.11 mesh nodes to intelligently utilize a range of frequencies.

Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Network Architecture and Design—*Wireless Communications*

General Terms

Design, Experimentation, Performance

Keywords

Mesh Networks, Ad-hoc Networks, Routing, Channel Assignment, Channel Selection, 802.11

1. INTRODUCTION

The integration of 802.11 technology in laptops, phones and PDAs has created a demand for WLAN connectivity everywhere. In spite of this demand, infrastructure 802.11 networks are constrained by short transmission ranges. Large 802.11 wireless networks require many Access Points (APs), each with individual wired Ethernet connections to provide the network backbone. This is the wireless paradox; each 'wireless' AP must be individually wired. Wiring each AP creates a number of deployment problems. The 100m cabling limitation of Ethernet is insufficient, the cost of fiber is restrictive, and WAN technologies such as ADSL are both costly and slow. Furthermore, perhaps the biggest hindrance to deployment is the administrative difficulty of laying cables through public and private property. Wireless 802.11 mesh networks have the potential to solve this problem. By replacing the wired backbone link with a series of wireless links, mesh technologies can create inexpensive 802.11 backbones enabling large, truly wireless 802.11 networks.

Such networks have a number of applications. Current WLANs could be extended into open spaces where cabling issues had previously restricted deployment. More ambitious applications include; improving developing world communications where wired infrastructure is insufficient or disaster recovery when existing infrastructure has been destroyed. 802.11 mesh networks could also be used in residential areas to form community based wireless networks. Despite these attractive applications; throughput degradations currently limit the size and scalability of mesh networks. When transmissions are hopped across multiple APs the performance is significantly lowered.

This paper is structured as follows; firstly we will investigate throughput degradation in multi-hop 802.11 ad-hoc networks. Following this we will explore prior work which has attempted to address this problem. This leads us to our solution, RDCS, which is compared and contrasted with prior work. Finally, we will discuss implementation issues, followed by our results and conclusion.

1.1 Problem Definition

Multiple factors make throughput degradation in multi-hop 802.11 wireless mesh networks especially severe. Firstly, radios are half-duplex which means that to relay a packet, the radio must first receive and then transmit the packet. With only a single radio, the process of receiving and then transmitting the frame are two autonomous operations and the channel time required to perform this procedure is doubled. Furthermore, traditional WLANs use non-overlapping channels for adjacent APs, however, in multi-hop ad-hoc networks every node must use the same channel to participate; significantly increasing contention. The final catalyst of these contention problems is that the RF energy from 802.11 is often double the transmission range [1]. Consequently, a single transmission can prevent a large number of nearby nodes from simultaneously transmitting. Theoretically, only every fourth node in a string topology can transmit concurrently [7].

In Gupta and Kumar's paper [5], they demonstrated that

IWCMC'09 June 21-24, 2009, Leipzig, Germany

for a given number of nodes (n), each of which is communicating with another node, the maximum possible throughput, considering optimal node placement and communication pattern, is $\theta(\frac{1}{\sqrt{n}})$. Furthermore, it was also shown that, given a random node placement and communication pattern, the throughput available to each node follows the function $\varpi(\frac{1}{n\sqrt{logn}})$ [5] . These degradations in throughput were confirmed by Li et al [7] who used the combined methods of simulation and practical experimentation. This evidence suggests that as the number of nodes increase, the bandwidth available to individual nodes approaches zero. As stated in our problem description, the large interference range of 802.11 wireless mesh networks means that only every fourth node in a chain can operate simultaneously. In practical and simulated studies, Li et al [7] found that the random contention mechanism of 802.11 led to poor scheduling across a chain of nodes. The conclusion was that poor scheduling exacerbated the multi-hop problem such that only one in every seven nodes in a wireless chain could transmit concurrently.

To overcome the multi-hop throughput problem, wireless nodes must utilize multiple frequencies. If nodes could utilize orthogonal frequencies they could dramatically reduce contention; allowing simultaneous transmissions across multiple channels. The solution to this problem has been investigated at multiple layers in the OSI networking model. Our review will begin with MAC and network layer techniques.

2. PRIOR WORK

2.1 MC-MACs

New Multi Channel MAC (MC-MAC) schemes aim to take advantage of multiple channels by negotiating a schedule of transmissions across non-overlapping channels [13, 11]. Some schemes use multiple radios [13] while others operate with only a single radio [11]. The 802.11s mesh networking standards group was originally considering a mechanism similar to [11] known as CCF (Common Channel Framework). However, it was removed due to interoperability concerns with existing 802.11 DCF (Distributed Coordination Function). We suggest that this concern may be generalized for all of the modified MC-MACs.

As an alternative to creating a new multi channel MAC layer, we propose that the traditional 802.11 MAC could be used with multiple ad-hoc mode radios. This shifts the problem to a higher layer, requiring a dynamic channel assignment mechanism. To concurrently use multiple channels a new mechanism must assign channels to radios in a manner that forms a connected yet channel diverse mesh network. Before delving further into channel assignment, we will investigate the related issue of routing.

2.2 Routing Protocols and Routing Metrics

If a new MAC introduces interoperability difficulties, this problem must be addressed at higher layers. Routing protocols must therefore be redesigned to incorporate channel information into the metric. Currently, the majority of ad-hoc routing protocols can be divided into two categories, on-demand routing and proactive routing. Studies have suggested that on demand routing protocols are more suited to smaller and mobile ad-hoc networks whereas proactive routing protocols are more suited to larger and low mobility networks [14]. OLSR is a proactive, link-state routing protocol with a large user base and demonstrated scalability. Subsequently, we chose OLSR as the routing protocol to base our channel assignment mechanism around.

A variety of routing metrics have been proposed for ad-hoc networks. ETX (Expected Transmission Count) [2] is a reliability metric calculated through the transmission of hellos in both directions. The higher the percentage of successful hellos, the lower the metric. ETX was extended to ETT (Expected Transmission Time) [3] by adding a link bandwidth measure. While ETT improves upon ETX, neither of these metrics can incorporate the effect of channels.

The first metric to include channel information was WCETT (Weighted Cumulative Expected Transmission Time) [3]. WCETT improves on ETT by counting the number of times each channel is used on a given path. Repeated use of the same channel is punished by a higher metric. MIC (Metric of Interference and Channel-switching) [14] supersedes WCETT. In addition to capturing the number of interfering links in the same flow, it can also capture the effect of nearby interfering flows.

Ideally, all routing protocols, proactive and on-demand, should use MIC or WCETT. However, a number of problems prevent their use in current routing protocols. In the case of WCETT and MIC, they are non-isotonic metrics. Isotonicity is required for proactive link-state routing protocols such as OLSR and thus, they are currently unable to be used. Isotonic metrics maintain the cost of a given link when prepended by a common path [14]. Put simply, routing loops can form if the cost of a given link can be changed based on the channel of a previous link. A further issue that prevents even the implementation of ETT in OLSR is that there is currently no unified way to obtain the transmission bit rate for each neighbor from the wireless driver. This leaves ETX as the best metric which can currently be used with OLSR.

3. CHANNEL SELECTION

While channel aware metrics are critically needed, we believe that the most pronounced performance gains will be realized with a channel selection mechanism that can assign frequencies to radios in an intelligent manner. Channel selection must balance two opposing objectives; interference and connectedness. Interference is the number of nodes sharing the same frequency in the same area. Connectedness is provided when every node has a working path to every other node. The goal of channel assignment is to minimize interference while maintaining the networks connectedness. Prior approaches can be divided into centralized approaches, where channel assignment is performed by a single server, and distributed approaches, where each node makes decisions based on locally available information.

3.1 Centralized

Early approaches to channel assignment [12, 8] used a centralized server to assign channels to nodes. These contributions were predominately theoretical, formulating the problem in graph theoretic terms and assuming that an input to the algorithm is the unit disk graph or G=(V,E), where the graph G has V (vertices), representing nodes, and E (edges), representing possible links. Most of the centralized schemes use variations on graph coloring algorithms such as weighted conflict graphs to assign colors (channels) to edges (links) in a manner that maintains connectivity yet simultaneously minimizes the number of nodes sharing the same channel [12, 8]. As the complexity of finding a solution is exponential [12, 8], to enable these algorithms to terminate in polynomial time, heuristic/approximate approaches are used.

The theoretical nature and lack of experimentation with centralized approaches led to a number of oversights. The unit disk graph, which is an assumed input for many of the previously mentioned channel assignment algorithms, is difficult to obtain in a practical environment. Furthermore, few of these schemes provided mechanisms to recover from network partitions.

3.2 Distributed

Recent papers [6, 9] have proposed distributed channel assignment mechanisms which are significantly more experimental; often discussing the physical architecture and network design and producing experimental results via a testbed. Hyacinth [9] is a channel assignment scheme designed to provision Internet access across the mesh. The approach is similar to 802.1D STP (Spanning Tree Protocol) because the network formed is a tree structure spanning away from the gateway. This approach provides fast reactions to topology changes and link failures. Routing is also simplified as all traffic gets routed upstream towards the gateway. However, the tree based approach makes operation in the presence of multiple gateways problematic. Furthermore, a STP design will concatenate paths and flows reducing the available bandwidth. Traffic not destined for the gateway, will travel upstream when more direct and less congested paths may exist elsewhere.

Ko et al [6] designed a channel selection mechanism that can support multiple gateways, does not build a STP tree and ensures mesh connectivity. For brevity, we refer to this work as DCA [6]. Like Hyacinth, DCA is based on a two radio architecture, however, in DCA, one radio operates on a default common control channel and the other radio is used to provide channel diversity. The non-control radio switches between channels in the 2.4 GHz spectrum with the goal of individually lowering the nodes local interference. DCA was tested in a 14 node testbed with the Microsoft LQSR routing protocol and the WCETT metric. DCA [6] operates independently from the routing protocol enabling it to operate with any routing protocol. The design choice to use a common control channel provides resiliency against topology changes, however; the drawback of a common control channel is greater contention. Schemes such as Hyacinth which use dynamic assignment on every radio will suffer from less contention than DCA.

4. RDCS

RDCS (Routing Driven Channel Selection) is a distributed channel selection algorithm written in Ruby. Like Hyacinth [9] it uses two switchable interfaces to provide maximum channel diversity. It includes mechanisms to stabilize on channel assignments and also react to resolve network partitions when connectedness is lost. As RDCS is coupled with a link-state routing protocol, the network is fully routed and capable of operating in the presence of multiple gateways. RDCS makes intelligent channel selections by leveraging information from the wireless driver, the OLSR routing table and also by exchanging information with neighboring nodes.

4.1 Channel Masking

Before making a decision as to which channel to use, RDCS will trim or mask the list of available frequencies. In RDCS, channels are deemed inappropriate for a number of reasons including: to avoid rogue nodes, to manage inter-radio interference and to prevent dual links forming between two nodes. Combined, these unusable frequencies are used as a mask against the list of usable frequencies. This section investigates the reasons for each frequency mask. While reading this section the reader should be mindful that a separate channel mask is created for each radio interface.

The rogue node avoidance mask searches for channels used by nodes that are not mesh nodes. A basic function of 802.11 devices is the ability to scan through alternate wireless channels and find other wireless devices. This information is stored in a scan cache by the wireless driver and is retrieved with the wireless-tools package. Nodes not using the specified SSID and not operating in ad-hoc mode are deemed rogue nodes and their frequencies are added to the rogue node avoidance mask. Sharing a channel with rogue APs can result in interference and less channel time to transmit frames, degrading throughput.

The inter-radio interference mask ensures adequate separation between two radios deployed in the same machine. During experiments, numerous research efforts [10, 3, 4] encountered interference problems with closely located radios. It was discovered that using two 802.11b cards in the same machine degrades performance so severely that the benefits of a multi-radio system are completely negated [10]. The results suggest that two co-located radios cannot send and/or receive simultaneously even if operating on orthogonal 2.4 GHz frequencies. Fuxjager et al [4] refers to this as the near-far problem and suggests a serious mismatch in prior channel selection mechanisms that neglected to address this problem. Our solution is to ensure that radios are adequately separated by utilizing different frequency bands. For each interface, RDCS will check the alternate interfaces on the router to determine which bands are in use. By creating the inter-radio interference mask, we can prevent co-located radios from using the same band. While users can define their own bands in the RDCS configuration file; RDCS.conf, our testing confirms that three non-interfering bands exist. These are the ISM band, the UNI-1 band and UNI-2 band.

The final mask is the dual-link mask which prevents two directly connected mesh nodes from acquiring the same channel assignment. Little benefit is attained from two nodes utilizing the same two channels. To prevent this from occurring, each node checks the channels used by its current 1-hop neighbors. Neighbors are listed in the OLSR links table. The channels used by directly connected neighbors form the dual-link mask. The combination of the rogue node, inter-radio interference and dual-link mask will be applied against the list of available frequencies. The remaining channels are valued on a channel cost metric.

4.2 Channel Costing

Potential neighboring nodes are located in the same manner that rogue nodes are discovered. The wireless drivers scan cache is searched for nodes in ad-hoc mode with the correct SSID. The channel mask is then applied against the list of channels containing mesh nodes. If mesh nodes are detected on the remaining usable channels, they are evaluated based on a simple cost metric. The formula shown in equation 1 states that the estimated capacity of a channel is equal to the average SNR (Signal-to-Noise-Ratio) of all nodes sharing a channel divided by the number of nodes. This formula will provide high (favorable) values to channels containing few nodes with high SNRs and low (unfavorable) values to channels containing many nodes with low SNRs. The following section will show that this formula is not the critical metric and subsequently we do not validate its accuracy. Once we have discussed how channel assignments are stabilized it will become clear that channel selection is really routing driven.

$$EstBW = \frac{SNRavg}{NoNodes} command. \tag{1}$$

4.3 Stability and Optimization

The ripple effect is the potential knock-on effect that occurs when a channel change causes a cascade or ripple of channel reassignments throughout the network [10]. Some schemes circumvent this problem by allowing a finite number of channel changes [6] or terminating the assignment algorithm after an approximate channel assignment has been found [12, 8]. These will not suffer from the ripple effect because their channel assignment mechanism has terminated. However this renders the algorithm incapable of responding to future network/topology changes. In Hyacinth [9], topology changes or breaks en route to the gateway simply require downstream nodes to search for and connect to another upstream node.

If RDCS was left in the currently described form, this problem would prevent the formation a stable channel assignment. The constant channel reassignments would be problematic for routing protocols and cause repeated Dijkstra recalculations. RDCS needs a way to deal with the ripple effect, stabilizing channel assignments that are providing useful links yet remaining equally capable of reacting to topology changes and link breaks. To provide this function, we decided that every node in the network must have a consistent identical asymetric view of the cost of every link. Equation 1 does not provide a consistent asymmetric metric. SNRs are not symmetric and fluctuate over time. Furthermore, the addition of a new node on a given channel may dramatically change the EstBW, creating a ripple of channel reassignments throughout the network.

Fortunately, link asymmetry and a consistent global view is one of the key features of link-state routing protocols. Subsequently, our design uses the routing table to decide whether an interface must be reassigned a new channel. In RDCS, if an interface is providing a unique route in the routing table, it will remain unchanged. However, if an interface is no longer providing a route in the routing table then it will be reassigned through the afore mentioned processes. Although initial channel selections are done using the channel masking and channel costing mentioned above, the final decision is based on the routing protocol. This cross layer functionality is beneficial for a number of reasons. Firstly, the channel assignment mechanism can be sure that its view of the network is identical to all other nodes. Secondly, it will optimize topologies for whatever metric the routing protocol values. There is a negligible benefit to creating a topology that is not valued equally by the routing protocol. Finally, if an interface is not providing a link in the routing table, no harm or network instability can be caused through reassignment.

4.4 Algorithm Overview

With the major features of RDCS already described, this section provides the sequence in which these operations are performed. Each interface is evaluated individually. RDCS begins by performing a check of the OLSR routing table to determine whether an interface is providing a route. If the interface is not providing a unique route in the routing table then it will undergo channel selection which includes masking, costing and then assignment. Alternatively, if the interface is providing a route, RDCS will perform two additional checks. The first check will determine how many other nodes are sharing the same channel. This variable is user defined. If this test fails, the interface will undergo masking, costing and assignment. The other check is the Internet gateway check. Consecutive fails of this test will also cause channel reassignment. Overall, RDCS tends to form stable channel diverse mesh networks by maintaining the channel assignment for interfaces providing stable routes and reassigning the channel on interfaces that are not providing routes.

Initial channel assignment schemes were highly theoretical. Since these contributions, experimental work uncovered a number of highly restricting factors. For example, experimental work discovered that co-located radios must utilize completely different bands to enable nodes to simultaneously transmit [10, 3, 4]. This is in strong contrast to theoretical work advocating the use of five overlapping channels in the 2.4GHz band. Also, the majority of studies, perform an initial channel assignment while the network is forming then terminate; surrendering the ability to adapt to future topology changes [12, 8, 6] whereas RDCS retains the ability to adapt to future changes. While some channel assignment mechanisms are coupled with relatively primitive routing protocols [9] that may increase congestion through inefficient routing, RDCS is based on a link-state routing protocol. Finally, some studies rely on one common control channel [6], while this can simplify a number of issues and ensure connectedness, it may increase contention compared to other approaches that do not use a control channel. RDCS was designed around experimental observations and prior work in wireless mesh networks. While it addresses many problems in prior work, RDCS is also not immune from issues and drawbacks.

5. IMPLEMENTATION ISSUES AND DISCUSSION

Before discussing the results, we will review the problems encountered when implementing and testing RDCS. Currently, OLSR uses a reliability metric called ETX, however, this metric is unable to incorporate bandwidth. In RDCS, channel selection is routing driven and hence determined by the routing metric. Therefore, the reliability metric, ETX, prevents the routing protocol and therefore RDCS from differentiating between high and low bandwidth links. Using the current metric, OLSR will never prefer a high quality two hop path over a poor single hop path. When the effects of bandwidth can be incorporated into the OLSR metric, improvements in both routing and RDCS channel selection performance should also be seen.

During testing, we also discovered a number of fairness issues. Currently, there is no working layer 2 QoS mechanism

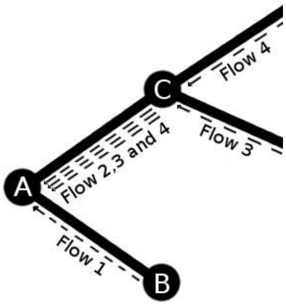

Figure 1: Flow Unfairness

for ad-hoc mode wireless interfaces. With network congestion, OLSR hello messages can be delayed causing the ETX metric to rise and occasionally lead to route changes. In Figure 1 there are two nodes, B and C which are directly connected to the gateway. Ideally, the flows 1, 2, 3 and 4 should share the bandwidth equally. However, if node C and Node B both share the same channel, the 802.11 MAC layer will provide contention resolution between the two nodes. Subsequently, flow 1 will consume half of the channel time leaving the other half to be shared between flow 2, 3 and 4; resulting in unfair bandwidth allocation.

Another problem, which may exacerbate the previous issue of flow based unfairness is RTT (Round Trip Time) and link quality unfairness. TCP congestion windows are highly dependent on RTTs. Subsequently, flows originating close to the gateway will have lower RTT's and better TCP performance. This is similar to link quality unfairness which is equally problematic. Frames transmitted by nodes further from the gateway must cross a larger number of links, lowering reliability. Although 802.11 uses acknowledgments to ensure reliability, retransmitted frames may still result in unexpected RTTs and TCP timeouts. Put simply, flows originating closer to the gateway will experience lower RTTs and fewer retransmissions; resulting in better performance.

Generally, we found that nodes one hop away connected to the gateway quickly capitalized on the available bandwidth, starving more distant nodes of bandwidth. To improve the consistency and fairness of our results we used rate limiting at the Internet gateway, limiting each node to 5 Mb/s and allowing bursting to 10Mb/s. Such rate limiting techniques are congruent with the methods used in many community mesh networks.

6. TESTBED AND RESULTS

The testing platform consisted of eight ALIX wireless nodes, each with two 802.11a/b/g network cards. The ALIX platform, is an embedded x86 PC with dual miniPCI slots. The operating system used was Voyage Linux and the wireless cards were Atheros CM9's running a modified version of MadWiFi v9.3.3. OLSR version 5.6 was used with the txtinfo plug-in installed. We have compared RDCS with a single radio static channel assignment and a dual radio static channel assignment, where each nodes two radios were assigned the same set of channels.

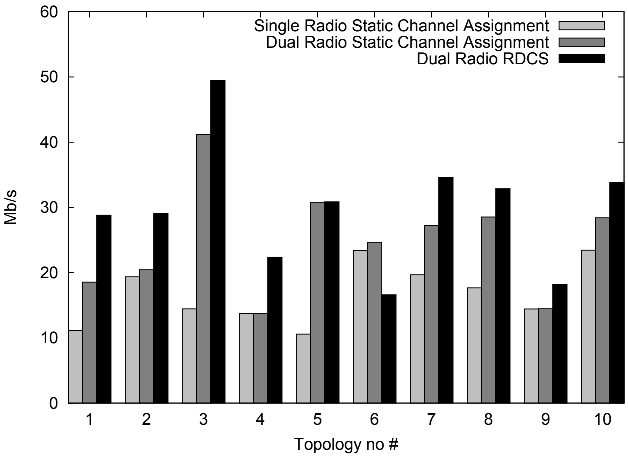

Figure 2: Gateway/Internet Bound Traffic

In our first test, a wired server performed simultaneous iperf tests to every wireless node via the mesh network gateway. A range of dense and sparse topologies were examined. The results of each topology (or physical placement of nodes), is numbered 1 to 10 in Figure 2. A cursory glance reveals variations in single radio, dual radio and RDCS results. While differences were anticipated, the difference between single radio, dual radio and RDCS were expected to be more consistent. The unexpected variation is a result of the routing protocol being channel and bandwidth unaware.

Given a multi channel assignment, the routing protocol (OLSR) sometimes preferred a single channel throughout the network resulting a large amount of contention and significantly lower bandwidth. Equally, sometimes OLSR balanced the multiple channels resulting in significantly higher throughput. Also, due to OLSR's and hence, RDCS's inability to differentiate between links with disproportionate bandwidths, topologies that are reliable yet can only provide a few Mb/s of actual throughput sometimes formed. This is evident in topology 6 where the RDCS topology performed worse than single radio and multi radio configurations. Our conclusion from these results is that variations are caused by ETX which neither captures bandwidth or channel information.

Our second test attempted to emulate a file sharing or P2P traffic model. A Ruby program was built whereby each node randomly starts iperf sessions with other nodes. The bandwidth recorded by each random iperf session was averaged. This test was performed over a longer period of time, with each node simultaneously performing iperf tests to other nodes. A number of factors such as a longer testing duration, and non-gateway centric bandwidth tests provided more consistent results; shown in Figure 3.

In both tests RDCS outperforms single radio and dual radio static channel assignments and; given a larger testbed, we believe that the difference between RDCS and the static channel assignments would be even more pronounced. With static channel assignments, larger numbers of nodes, will further increase contention and reduce the bandwidth available to each individual node. Comparatively, mechanisms such as RDCS that produce channel diverse topologies, will suffer from less contention as the number of nodes increase. Additionally, when OLSR incorporates a bandwidth and channel metric, RDCS will be able to make more intelligent channel assignments. Therefore, while RDCS currently outperforms static channel assignments, we project that the difference will also grow with more intelligent routing metrics.

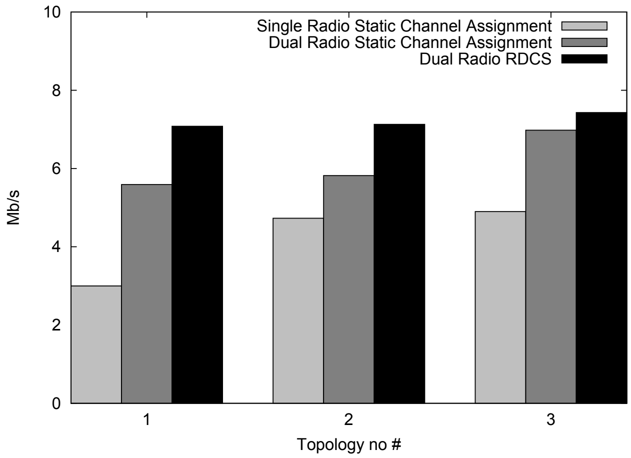

Figure 3: Random Traffic Paths

7. CONCLUSION

This paper proposes a scheme that attempts to overcome a fundamental performance limitation in 802.11 wireless mesh networks. Theoretical work has shown that in single radio networks, as the number of nodes increase, the bandwidth available to individual nodes approaches zero. The reason for this sharp degradation in bandwidth in a multi-hop environment is caused by a combination of problems including half-duplex transmission, contention, and 802.11's large carrier sensing range. RDCS uses multiple radios to build a channel diverse topology. Channel selection in multi-radio ad-hoc networks is non-trivial because the competing variables of connectedness and interference must be balanced. Furthermore, channel assignments must be performed in a distributed manner based only on locally available information. RDCS builds channel diverse topologies that minimize interference while simultaneously ensuring that the network is fully connected. The performance of RDCS confirms the benefits of a channel diverse topology. However, before multi-radio mesh networks can realize their full potential, many issues remain at all layers of the OSI networking model. Routing metrics must incorporate bandwidth and perhaps eventually channel information and transport layer unfairness including flow based unfairness and RTT unfairness also require attention. RDCS is capable of increasing spatial diversity by creating multi channel topologies, however, much future work is required to fully utilize this additional capacity.

8. REFERENCES

[1] G. Anastasi, E. Borgia, M. Conti, and E. Gregori. Wi-fi in ad hoc mode: A measurement study. In *PERCOM '04: Proceedings of the Second IEEE International Conference on Pervasive Computing and Communications (PerCom'04)*, page 145, Washington, DC, USA, 2004. IEEE Computer Society.

[2] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 134–146, New York, NY, USA, 2003. ACM Press.

[3] 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*, pages 114–128, New York, NY, USA, 2004. ACM Press.

[4] P. Fuxjager, D. Valerio, and F. Ricciato. The myth of non-overlapping channels: Interference measurements in ieee 802.11. In *WONS: Wireless on Demand Network Systems and Services*, 2007.

[5] P. Guptak and P. R. Kumar. The capacity of wireless networks. In *IEEE Transactions on Information Theory*, 2000.

[6] B.-J. Ko, V. Misra, J. Padhye, and D. Rubenstein. Distributed channel assignment in multi-radio 802.11 mesh networks. In *WCNC: Wireless Communications and Networking Conference*, 2007.

[7] 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*, pages 61–69, Rome, Italy, July 2001.

[8] K. N. Ramachandran, E. M. Belding, K. C. Almeroth, and M. M. Buddhikot. Interference-aware channel assignement in multi-radio wireless mesh networks. In *IEEE INFOCOM*, 2006.

[9] A. Raniwala and T. Chiueh. Architecture and algorithms for an ieee 802.11-based multi-channel wireless mesh network. In *INFOCOM - 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, 2005.

[10] A. Raniwala, R. Krishnan, and T. Chiueh. *Wireless Mesh Networking: Architectures, Protocols and Standards*, chapter IEEE 802.11-Based Wireless Mesh Networks, pages 79 – 109. Auerbach, 2007.

[11] J. So and N. H. Vaidya. Multi-channel mac for ad hoc networks: handling multi-channel hidden terminals using a single transceiver. In *MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*, pages 222–233, New York, NY, USA, 2004. ACM.

[12] J. Tang, G. Xue, and W. Zhang. Interference-aware topology control and qos routing in multi-channel wireless mesh networks. In *MobiHoc '05: Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, pages 68–77, New York, NY, USA, 2005. ACM.

[13] S.-L. Wu, C.-Y. Lin, Y.-C. Tseng, and J.-P. Sheu. A new multi-channel mac protocol with on-demand channel assignment for multi-hop mobile ad hoc networks. In *ISPAN '00: Proceedings of the 2000 International Symposium on Parallel Architectures, Algorithms and Networks (ISPAN '00)*, page 232, Washington, DC, USA, 2000. IEEE Computer Society.

[14] Y. Yang, J. Wang, and R. Kravets. Designing routing metrics for mesh networks. In *IEEE WiMesh 2005*, Santa Clara, Ca, September 2005.