Distributed Split-Path Routing Strategy for Multi-hop Mesh Networks

Paulo Alexandre Regis, Shamik Sengupta
Department of Computer Science and Engineering
University of Nevada, Reno, USA, 89557
pregis@nevada.unr.edu, ssengupta@unr.edu

***Abstract*—The concept of mesh networks brings many challenges to the research community when designing such system. In this paper, we propose a new routing model to account for the link utilization between nodes, as well as the remaining energy left in each node. The framework, called Distributed Split-path, is implemented on top of the Optimized Link-State Routing (OLSR) protocol. We demonstrate the improvement our proposed scheme offers to the longevity of the network while maintaining similar performance as the original OLSR protocol. We devise a performance indicator to show in which cases our approach benefits the network system designer, depending on the priorities of the network application. Simulation results show that by allowing a trade-off in the performance, the new model outperforms current protocol.**

***Keywords*—Wireless mesh network, routing, metric, energy efficient, multi-path**

I. Introduction

Due to the rapid evolution of handheld electronic devices, a massive surge in wireless communication demand was seen in the last decade. Together with the development of small unmanned aerial vehicles (UAVs), this advancement brought a new age of mobile ad-hoc networks (MANET) applications. The new UAV technologies, coupled with developments in unmanned ground vehicles (UGVs) added a new dimension to service provision of wireless networking systems. The new trend of 3D wireless mesh networks enables mission-centric operations like military tactical network and first responders in disaster areas network in situations where infrastructure is not available.

Throughput hungry applications, such as real-time video streaming and large files transfer, push forward the development of higher capacity wireless networks. One alternative is to distribute the load in the network among all of its nodes. This way, if well designed, the network reduces the traffic bottleneck while maintaining the performance at an acceptable level. When supporting a specific mission, the priority of the network is to provide uninterrupted communication, as well as offer enough capacity to execute the task. Mesh networks can provide a redundant connection by forming multiple paths between two nodes, this way if one path is disrupted, communication can be rerouted through a different set of nodes. Simple examples are border control and surveillance networks. Sensors are deployed to collect certain types of data and send back to a control base, where the data can be analyzed. These sensors are typically small and only perform a single specific task (measure, record, film, etc.). They can be deployed in areas with difficult access, which makes maintenance a challenge. Thus, creating mechanisms to maximize the longevity of the devices is very important. Devices capable of forwarding packets can autonomously create a multi-hop network to provide wireless coverage to tactical or disaster teams over a wide area by relaying the data through multiple nodes and frequencies. We see an example of such an application in Figure 1. The UAV network provides connectivity to the soldiers when there is no infrastructure available. It enables the team to communicate among themselves or send information to a remote location, such as the headquarters, through a gateway in the network. Having more than a single path of communication makes the system more robust in case of a node malfunction.

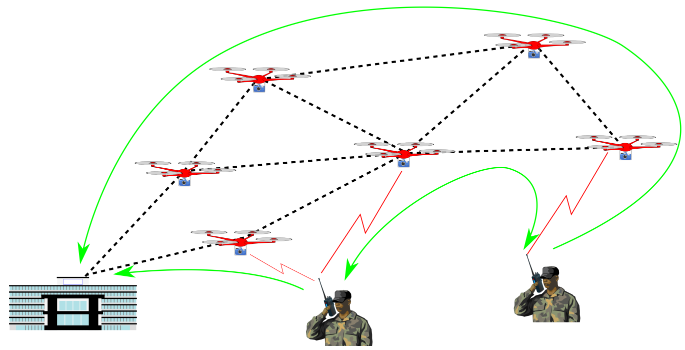

Figure 1: Example application scenario: arrows represent a possible path of communication between end-nodes. The different colors (red and black) distinguishes different communication channels

The multi-hop nature of the network poses some serious challenges for resource (time, bandwidth, power) allocation among the devices. Ideally, the network should not have any bottleneck to avoid interruption of the service, and nodes should be able to function for as long as possible before running out of energy. The longevity problem becomes very complicated if no centralized controller defines the positioning of nodes, and the paths the end-to-end communication should traverse, based on the information of the entire network. However, having a complete information of the network might not even be feasible, due to the dynamic and autonomous characteristics of a mesh network. In a network with energy constrained devices, increasing the throughput without concerning about the individual power consumption of the nodes can interrupt the connectivity of the entire system if there is a bottleneck (i.e. most of the communications have to pass through a small set of nodes).

In this paper, we expand our previous work where we designed a strategy to pair traffic flows and paths in an ad-hoc network [1]. We develop a new routing mechanism that encompasses information from other layers of the network stack. The new approach, called Distributed Split-path, chooses the next hop for a packet based on the previous route requests, combined with the local energy information and transmission characteristics of the link. The traffic should be split into multiple paths to improve the performance regarding energy consumption. We put both together in a link-state routing protocol and test over a reliable network simulator, ns-3 [2].

Resource allocation in wireless ad hoc networks has been studied in many aspects such as joint power allocation and routing, joint routing and dynamic spectrum access, resource allocation based on the quality of service requirements, energy efficiency, etc [3], [4], [5]. The work in [6] proposes different strategies to allocate power and minimize the bit error rate in multi-hop networks. While authors in [7] propose a mechanism tries to maximize the throughput in cognitive radio ad-hoc networks with a joint routing and dynamic spectrum access strategy. The possibility of having multiple radio interfaces, enabling access to multiple frequencies, makes the problem more complex. In [8] the authors tackle resource allocation based on the quality of service requirements considering only homogeneous networks, where a base station serves an arbitrary number of users. Similarly, the work in [9] investigates the importance of energy efficiency in LTE and WiFi systems. [10] investigates the channel assignment problem in wireless mesh networks. The problem is elaborated as a joint routing and channel assignment task, not considering power balance among the nodes. Regarding heterogeneous networks resource allocation, [11] tries to maximize the energy efficiency of individual nodes, but do not account for interference. [12] proposes a prioritized resource allocation method for heterogeneous networks. The model splits traffic flows into multiple paths to achieve better throughput, but energy efficiency and network longevity are not considered. Studying novel mechanisms to increase the capacity of the system is important. However, without careful design of flow assignment and novel distributed mechanisms to increase scalability, performance, and longevity jointly, the very features of mission-centric mesh network can turn into disadvantages.

In contrast to the above works, this work is focused on route assignment based on the mission-centric objective as well as maintaining scalability, longevity, and quality of services. We propose a method to split the generated traffic into multiple paths in a mission-centric ad-hoc mesh network. We tackle this problem by considering power source, medium access control, and history of local route requests and incorporate them in a new metric used by the routing protocol. Since the cost of a link is based on different variables, the local algorithm balances the traffic among its peers based on the updated value of this cost. The objective is to increase the longevity of the network, even though it may not yield the optimum result since nodes do not have complete information of the system at any given point in time. Simulation results show that the proposed mechanism may perform better when a trade-off between delivery ratio and remaining power is allowed when designing the network. We can maintain similar performance metrics while reducing the throughput, and increasing the energy saved at the end of the simulations.

The remainder of the paper is organized as follows. Section II presents a literature review of the field. In Section III we describe the system model, and the proposed Distributed Split-path model. Section IV describes the simulation experiment setup, the parameters, and indicators used to evaluate, and discusses the obtained results. Finally, Section V gives the final remarks and concludes the paper.

II. Related Work

The routing problem in wireless ad-hoc networks has been extensively studied in the past [13], [14], and cross-layer optimization is also a well known problem [15], [16]. In this section, we divide the literature into two distinct groups: routing protocols, and network metrics. The former targets the implementation of proposed mechanisms and its applicability, and the later focuses on new link utility metrics and estimations.

A. Routing Protocols

There are many routing protocols proposed in the literature; we focus on the most common ones, which have requests for comments, and their variations.

- **Ad-hoc On-demand Distance Vector (AODV)**: this protocol finds the route from source to destination based on route requests and route reply packets. When a node has data to transfer but has no entry in the routing table, it asks its neighbors for a route to the destination. Once found, it sends the data to the neighbor closest to the destination in terms of hop-count. It relies on the neighbor to decide the subsequent hop to send the packet.
- **Better Approach To Mobile Adhoc Networking (B.A.T.M.A.N.)**: uses a broadcast of originator messages to advertise the existence of a node to the entire network. Based on the sequence number, the time-to-live field in the packet (hop count), and the receiving nodes decide which neighbor is best suited to send data when the destination is the originator node.
- **Dynamic Source Routing (DSR)**: similar in nature to the AODV, this protocol discovers the route to the destination with route requests and replies but differentiates because if forces the packet to traverse a predefined path. It inserts in each packet the entire path that packet has to travel until it reaches the destination.
- **Optimized Link-State Routing (OLSR)**: is a link-state protocol, meaning that all nodes have complete information of the current state of the network. It uses *Hello* messages to advertise and discover new nodes on the network, and *Topology Control (TC)* to broadcast neighboring information to distant nodes. It floods the network using a *Multi Point Relay* mechanism, where only a subset of immediate neighbors relay the packets forward.
- **Multi-Path Optimized Link-State Routing (MP-OLSR)**: similar to OLSR, but instead of keeping an updated routing table, it calculates the next hop whenever requested, due to the mechanism used to calculate multiple routes. Packets can be sent using source routing (the protocol inserts the entire path in the header), this way avoiding loops. Alternatively, it can rely on the decision of the next hop neighbor, but not guarantee a loop-free network.

All protocol designs mitigate certain issues a network can suffer from, they all act in the best effort manner; a node forwards packets to whichever node apparent to be closer to the final destination. Although it may yield better performance, the longevity of the network is not considered. If a node runs out of battery, the network simply redirects the flow to the next best route. Next, we describe alternative routing metrics that reflect different characteristics of the links.

B. Routing Metrics

Routing metric does not necessarily mean it is only used by a certain protocol. The value assigned to a link may reflect other parameters of the network, and help neighboring nodes decide the best route based on these values. Hop count, for example, is the most widely known link cost. It simply reflects the number of nodes a packet has to traverse to reach the destination. Due to the ease of implementation, it is present in most routing protocols.

However, in cases such as multihop wireless networks, hop count may not reflect the actual conditions of each link. Take a wireless link in an urban environment for example. The wireless signal may reflect into the buildings, and the same signal can interfere with itself on the other end of the link. It is still a single hop, but this interference can significantly lower the quality of the transmission.

The *expected transmission count (etx)* metric was created to overcome the problem of the unreliable wireless medium. In practice, it is based on the sequence number of packets transmitted between neighbors. In perfect conditions each packet received will have an incremental sequence number. However, in practice packets may be lost. *Etx* is simply the number of times a packet is expected to be transmitted, and its acknowledgment received. The details of implementation can be seen in [17]. *Etx* incorporates the conditions of the medium into the link cost, values such as delay, throughput, are indirectly reflected in the value of the link. However, *etx* does not account for the conditions of the hardware itself, such as computing power of the neighbor, queue load, or power consumption of the node on the other side of the link.

III. Distributed Split-path

To enhance the longevity of the network in a distributed manner, we devise a new metric that incorporates the local conditions of the nodes. In this section, we explain the split-path components and how they work together. We make assumptions on our system model that allows us to implement the new approach in a reliable environment. While in this paper we consider a single-radio single-channel network, the model can be used in heterogeneous multi-radio multi-channel models as well. We first describe the system model and then move to the distributed split-path description.

A. System model

In this section, we describe the assumptions that enable the proposed mechanism and explain the reasoning behind it. As mentioned before, our approach aims to tackle the longevity problem of the network. In other words, we want to make the network last for as long as possible, and we focus on the network aspect of it. We do it by balancing the network load among the nodes in the network. Although the proposed mechanism allows for multi-radio multi-channel nodes, in our system, we use single-radio nodes to gather insights on how the mechanism behaves in such case.

We assume that nodes have access to its local information, such as remaining energy, voltage level, MAC layer configuration. This is possible with current technologies and software libraries in open-source communities. Since we are dealing with mesh networks, we also assume nodes can act as relays; any ad-hoc network routing protocol enables this feature. In this paper, we use OLSR protocol as a base to implement the proposed mechanism. As we consider the network to perform an arbitrary task, each node in the network is able to generate traffic. The amount of traffic generated can vary with the type of application the node/network is executing. All the traffic generated by the nodes are assumed to have a common destination. In a mesh network, this can be seen as the gateway to the Internet or a remote location, such as a command and control center.

B. Relative link quality: $q^{i,j}$

The relative link quality ($q^{i,j} \in (0,1]$) is the probability of successful packet transmission from a neighbor to a node, based on the number of received and transmitted packets. This is achieved by keeping track of the received Hello packets used in neighbor discovery for the OLSR protocol. Every Hello packet has an incremental sequence number, and on the receiver side, if two consecutive Hello packets are not sequential, it means something was lost in between two consecutive receptions.

Earlier we mentioned the *expected transmission count (etx)* metric for mesh networks. *Etx* is calculated using the link quality from both neighbors: $etx = \frac{1}{q^{i,j} \times q^{j,i}}$. Both peering nodes need to know the value of their neighbor's perspective of the link, this is done by including the q information in the Hello packets.

However, the split-path uses the Hello packet to let peering nodes know about the combined metric (described later). Thus, it can only use the locally measured link quality. As a result, neighboring nodes may have different costs for the same link. This may be a disadvantage, as nodes perceive links as symmetrical when in practice $q^{i,j}$ can be different from $q^{j,i}$.

C. Normalized link data-rate: δ_i

We use the link transmission data rate as part of the metric. We dothis as a method to give preference to higher speed links. We normalize the value of the transmission data rate because the other components are not in the same measurement unit. For example, combining bits per second with remaining energy fraction makes the energy fraction irrelevant as bps can be a value of the thousands or millions. Thus, in our approach, we normalize the data-rate by the highest data-rate interface available in the node. If a node has multiple interfaces, then the normalized link data rate will look like this:

$$\delta_i = \frac{d_i}{max[D]}$$

Where δ_i is the normalized data rate, $d_i \in D$ is the transmission data rate mode configured to device i, and D is the set of all interfaces on the node.

D. Remaining energy fraction: e

As the objective of our new method is to save energy, we incorporate the energy fraction in the metric:

$$e = \frac{remaining\_energy}{maximum\_battery\_capacity}$$

Where $e \in (0,1]$. This is done in a way to favour the links in which the composing nodes have a higher energy capacity, and diverge the traffic from nodes with lower energy.

E. History percentage: H_l

In a realistic application scenario, it is tough to predict how much traffic an application will generate and for how long it will run. To try and estimate these parameters, we create the history stack. We use a finite stack data structure to store the results of route request by the upper and lower layers of the network stack. The History percentage H_l is the number of entries in the stack destined to link l:

$$H_l = \frac{\sum^H h_l}{|H|}$$

Where $H_l \in [0, 1]$ is the estimated link usage, $h_l = 1$ is an entry in the stack representing a packet sent through link l, and $|H|$ is the size of the history stack.

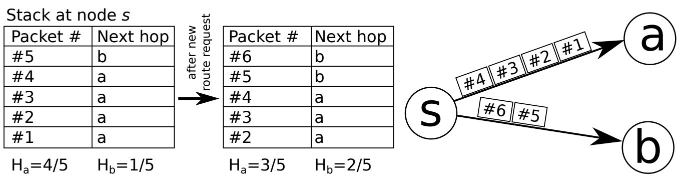

Figure 2: Node s is sending packets to its neighbors a and b.

For example, in Figure 2 we see packets in the pipe sent from s to its neighbors a and b. In this example, we assume that the history stack size is 5. As new packets arrive and need to be forwarded, old entries are discarded, changing the calculated H_l value.

Each entry in the history stack is simply the address of the resulting route request. A new entry is added every time the routing protocol decides to forward a packet to one of its neighbors, the neighbor's address is then pushed down the stack.

The size of the stack will influence the resulting metric. The reason to have a finite stack (apart from physical memory limits), is to account for intermittent data flows. As a new flow is created and transmitted, new entries are added to the history stack. When the same flow does not exist anymore (a video transmission has stopped, for example), those entries will be removed from the stack as new traffic is forwarded.

F. Combined metric

We combine the aforementioned components into a single split-path metric as follows:

$$f_l^{i,j}(\delta_l, q^{i,j}, H_l, e) = \frac{\delta_l}{q^{i,j}} + \frac{H_l}{e}$$

Where $f^{i,j}$ is the link cost function. We divide the normalized data rate with the relative link quality to regulate the maximum achievable data rate in that link (l). If the link quality is high (close to 1), the link cost tends to be lower. Otherwise, it gives a penalty to that link by increasing its cost. The Same principle is applied on the second part of the function. If the estimated link utilization H_l is high, then the link cost will increase. The H_l value is regulated by the remaining energy fraction e of the node. Thus, the more energy left, the less H_l will affect the link cost.

G. Split-path routing

The protocol implementation is built on top of the OLSR protocol. It uses Hello packets to estimate the relative link quality. Then the value is used to calculate the link cost. Advertised Hello packets contain the calculated link cost of all immediate neighbors. Neighbors do not need to know their peers link cost to themselves, but they use the information in the Hello packet to maintain the two-hop neighbor list updated.

Nodes advertise the link states to the network through Topology Control (TC) messages. In the message, there is a list of all immediate neighbors. Nodes use TC message to keep the link state information from nodes beyond two-hop distance. Due to the usage of relative link quality information, TC messages received from two distinct nodes that are neighbors among themselves may contain different link cost values representing the same link. In this case, while processing the TC message, the node only uses the information of the most recent one.

IV. Simulation

We implement the proposed mechanism in the ns3 [2] environment. The simulations measure network performance in terms of average throughput, delay, hop count, packet delivery ratio, and total remaining power. We use OLSR as a reference benchmark. We simulate each variation 10 times and average the results to obtain the final measurement.

The main parameters used in the simulation are listed in Table I. Each node is equipped with a single radio interface. We used the simple wireless model described in [18]. The destination of all traffic generated in all simulations is the first node, located in the left bottom position of the node grid, as seen in Figure 3.

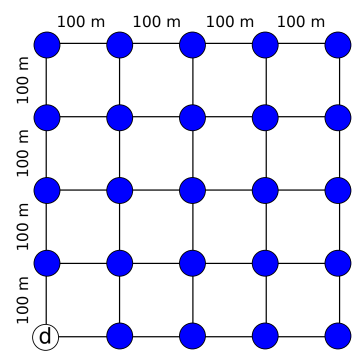

Figure 3: Topology grid, d is the destination of all traffic

We vary the number of source nodes generating traffic from 1 to 24 (the maximum possible number of source nodes). At each experiment run, the source nodes are randomly selected between the possible 24 nodes (excluding the destination). Next, we show and discuss the obtained results. We run the simulation for $400s$, which is enough to drain the energy of all the nodes.

A. Performance indicators

We evaluate the mechanism regarding common network performance measurements, namely: throughput, hop count, delay, and average lifetime of the nodes. To allow for a trade-off between throughput and lifetime, we also introduce a new indicator: a weighted aggregation between both performance metrics. The indicator is defined as $t = w_1|tp| + w_2|lt|$, where w_1 and w_2 are the weights ($w_1 + w_2 = 1$) of the normalized throughput ($|tp| \in [0, 1]$) and the normalized average lifetime ($|lt| \in [0, 1]$), respectively. The normalization is necessary since both metrics are of different units. The goal

TABLE I: Simulation parameters

Parameter	Value
Number of nodes	25
Communication range	100 m
Transmission bitrate	1 Mbps
Initial energy per node	100 J
Tx energy drain	0.5 A
Default energy drain	0.1 A
Simulation duration	400 s
Positioning	5 x 5 grid
Spacing	100 m
Source CBR	50 kbps

is to see when it is worth using the proposed model, as well as comparing the trade-off of different sizes of the history stack.

B. Results and discussion

Since the number of entries stored in the history of requests is a fixed number, we compare the proposed model with different size of history entries. We varied the history stack size in memory from 10 to 10^5. In Figure 4 we observe that OLSR outperforms in terms of throughput. However, our proposed model achieves superior lifetime than OLSR, while maintaining a comparable performance regarding delay and hop count. At first, it seems that the new model does not present any advantage. The apparent decline in lifetime looks proportional to the decreased throughput. However, analyzing the other metrics, the version of the mechanism with a history size of 10^5 entries stands out from the other versions. It is capable of providing similar delay and hop count values as OLSR while reducing the consumption to a similar level as the other variations. These reflect in a better throughput for the 10^5 version when compared to the smaller history sizes.

C. Effect of trade-off indicator

The objective of the trade-off indicator $t(w_1, w_2) = w_1|tp| + w_2|lt|$ is to examine when it becomes advantageous to keep more or less entries in the history stack. From Figure 5 we observe the transition from equally weighted throughput and lifetime, to when the lifetime of the network is more important. While the throughput depends on various network conditions, it also influences the amount of energy consumed. The higher the throughput, the higher the consumption. However, the average lifetime of the nodes in network depends not only on throughput, but it also depends on the energy provided by the energy source of each device.

We can see another advantage of the proposed model when the number of source nodes generating traffic is increased, and energy is of much more importance. While OLSR has a better indicator with fewer sources, the split metric architecture works better when there is more traffic, by splitting the flows into multiple paths. The indicator quickly transitions from OLSR to split-path once the weight for the later is slightly increased. Furthermore, we see that having a larger history size might not be more advantageous. Lower history stack means quicker computation of the history percentage. From the plots in Figure 5 we see that is the lifetime of the network is much more important than the throughput (i.e. the application does not require such high performance), having lower stack size may be a better choice.

V. Conclusion

In this paper, we proposed a new cross-layer metric model for ad-hoc wireless networks. The proposed mechanism utilizes the data rate mode of the link and battery level information to help compute the weight of all links in a local manner (i.e. there is no centralized controller that makes routing decisions). It also attempts to estimate the application needs by keeping track of the recent route requests. This way we try to measure the link utilization, and reroute the traffic if a link is being overwhelmed with traffic. The new metric combines all this information, plus the probability of successful transmission, and balances the load of the network. Simulation results showed that by allowing a trade-off in the network performance, the system is able to improve the performance indicators. We also demonstrated that the amount of past request kept in the history stack also plays a role in the overall performance.

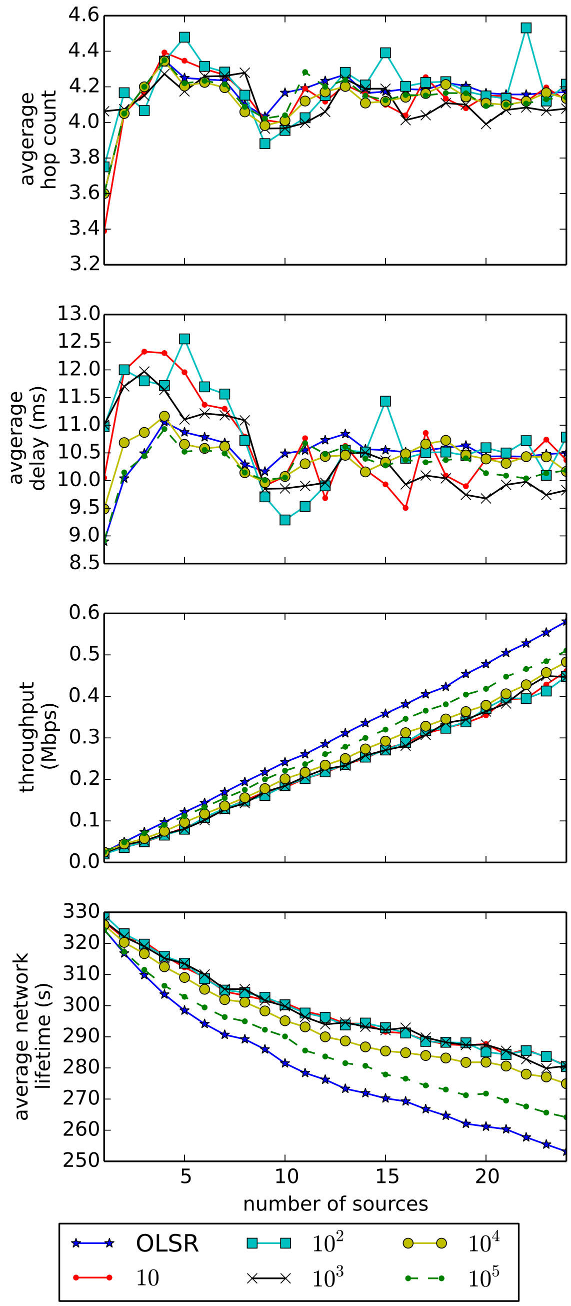

Figure 4: Network performance

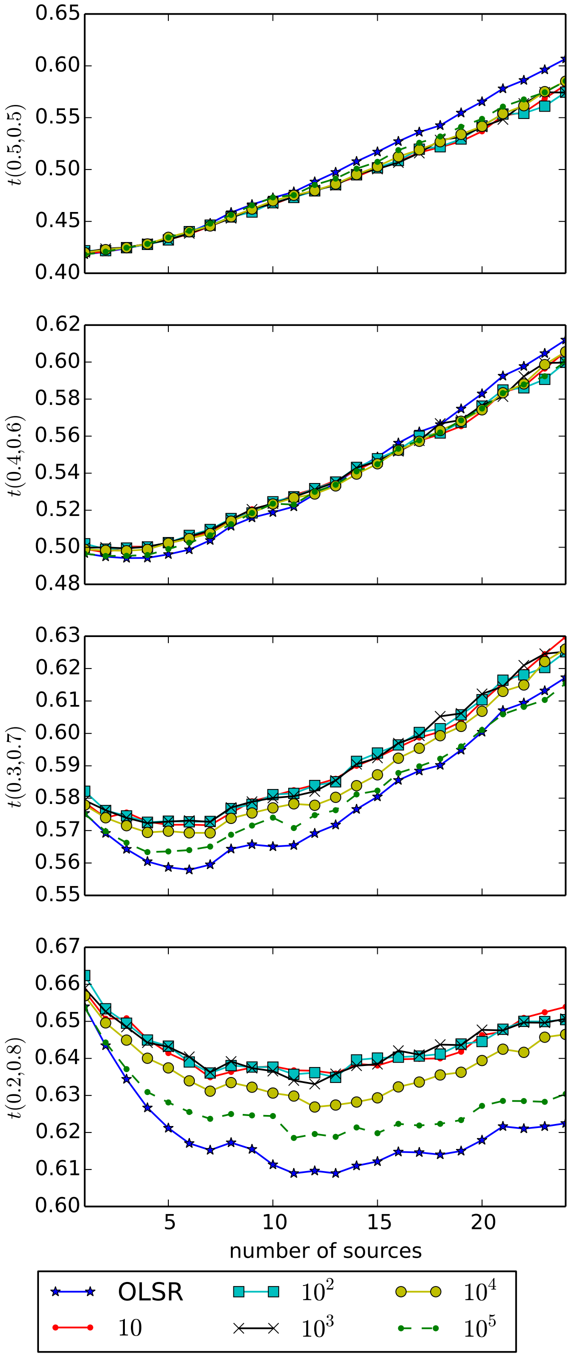

Figure 5: Trade-off comparison, varying w_1 from 0.5 to 0.2

ACKNOWLEDGMENTS

This research was partially supported by the NSF CAREER Award #1346600, NSF Partnership For Innovation Award #1430328, and CAPES Brazil 13184-13-0.

REFERENCES

[1] P. A. Regis, S. Bhunia, and S. Sengupta, "Enhancing performance and longevity of multi-radio multi-channel hetnets through dynamic path-assignment,"

[2] ns-3 Consortium, "Network simulator, ns-3." https://www.nsnam.org/.

[3] Y. Zhang, C. Lee, D. Niyato, and P. Wang, "Auction approaches for resource allocation in wireless systems: A survey," *IEEE Communications surveys & tutorials*, vol. 15, no. 3, pp. 1020–1041, 2013.

[4] D. Feng, C. Jiang, G. Lim, L. J. Cimini, G. Feng, and G. Y. Li, "A survey of energy-efficient wireless communications," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 1, pp. 167–178, 2013.

[5] N. Chakchouk, "A survey on opportunistic routing in wireless communication networks," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, pp. 2214–2241, 2015.

[6] S. Gupta and R. Bose, "Joint power allocation and routing optimization in ber constrained multihop wireless networks," in *Communications (NCC), 2013 National Conference on*, pp. 1–5, Feb 2013.

[7] L. Ding, T. Melodia, S. N. Batalama, J. D. Matyjas, and M. J. Medley, "Cross-layer routing and dynamic spectrum allocation in cognitive radio ad hoc networks," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 4, pp. 1969–1979, 2010.

[8] M. Fadel, A. S. Ibrahim, and H. Elgebaly, "Qos-aware multi-rat resource allocation with minimum transmit power in multiuser ofdm system," in *2012 IEEE Globecom Workshops*, pp. 670–675, IEEE, 2012.

[9] S. Navaratnarajah, A. Saeed, M. Dianati, and M. A. Imran, "Energy efficiency in heterogeneous wireless access networks," *IEEE Wireless Communications*, vol. 20, pp. 37–43, October 2013.

[10] J. J. Gálvez and P. M. Ruiz, "Efficient rate allocation, routing and channel assignment in wireless mesh networks supporting dynamic traffic flows," *Ad Hoc Networks*, vol. 11, no. 6, pp. 1765–1781, 2013.

[11] G. Yu, Y. Jiang, L. Xu, and G. Y. Li, "Multi-objective energy-efficient resource allocation for multi-rat heterogeneous networks," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 10, pp. 2118–2127, 2015.

[12] M. Gerasimenko, D. Moltchanov, R. Florea, N. Himayat, S. Andreev, and Y. Koucheryavy, "Prioritized centrally-controlled resource allocation in integrated multi-rat hetnets," in *2015 IEEE 81st Vehicular Technology Conference (VTC Spring)*, pp. 1–7, IEEE, 2015.

[13] J. N. Al-Karaki and A. E. Kamal, "Routing techniques in wireless sensor networks: a survey," *IEEE wireless communications*, vol. 11, no. 6, pp. 6–28, 2004.

[14] S. Adibi and S. Erfani, "A multipath routing survey for mobile ad-hoc networks," in *Proceedings of the IEEE Consumer Communications and Networking Conference (CCNC)*, vol. 2, pp. 984–988, 2006.

[15] S. Shakkottai, T. S. Rappaport, and P. C. Karlsson, "Cross-layer design for wireless networks," *IEEE Communications magazine*, vol. 41, no. 10, pp. 74–80, 2003.

[16] V. Srivastava and M. Motani, "Cross-layer design: a survey and the road ahead," *IEEE Communications Magazine*, vol. 43, no. 12, pp. 112–119, 2005.

[17] D. S. De Couto, *High-throughput routing for multi-hop wireless networks*. PhD thesis, Massachusetts Institute of Technology, 2004.

[18] P. Deutsch, L. Veytser, and B.-N. Cheng, "Ll simplewireless: A controlled mac/phy wireless model to enable network protocol research," in *Proceedings of the Workshop on Ns-3*, WNS3 '16, (New York, NY, USA), pp. 71–78, ACM, 2016.