Performance of Dual Wi-Fi Radios in Infrastructure-Supported Multi-Hop Networks

Fabian Dreier
Disney Research Zurich
8092 Zurich, Switzerland
fdreier@disneyresearch.com

Vladimir Vukadinovic
Disney Research Zurich
8092 Zurich, Switzerland
vvuk@disneyresearch.com

Stefan Mangold
Disney Research Zurich
8092 Zurich, Switzerland
stefan@disneyresearch.com

***Abstract*—Wireless multi-hop networks are a promising approach to provide connectivity in areas with sparse wireless infrastructure. In multi-hop networks, data is relayed from one mobile device to another until it reaches the destination (e.g. a Wi-Fi access points). Different hops preferably use different radio channels to avoid interfering with each other. It is not practical to use multiple channels with single-radio devices because the network would suffer from partitioning. The addition of the second radio interface to the devices may simplify the multi-channel management and improve the overall network performance. In this paper, we evaluate the potential benefits of the dual Wi-Fi radio devices in terms of throughput and energy consumption. For the evaluation, we use our Wi-Fi network simulator, Jemula802, extended with the models of multi-channel radio management algorithms.**

Keywords-802.11;multi-hop;multi-radio;multi-channel

I. INTRODUCTION

We consider a network architecture where mobile stations (e.g. handheld devices) use ad hoc multi-hop routes to connect to sparsely deployed Wi-Fi access points, which all have wired backhauls. Such a hybrid architecture of ad hoc and infrastructure-based components can be used to reduce the number of access points needed to cover an area. In entertainment theme parks, for example, it is often not feasible to densely deploy access points to provide full coverage because they might be too visible to guests and interfere with creative/artistic intentions. Besides, large theme parks are comparable in size with cities: One of the entertainment parks in Florida spans over 100 km^2, an area as large as San Francisco. Therefore, it could be prohibitively expensive to cover such an area with a traditional infrastructure-based Wi-Fi network.

Multi-hop networks provide range extension through relaying. If all stations are equipped with single Wi-Fi radios tuned to the same channel, capacity problems arise because nearby stations interfere with each other when they transmit simultaneously. This problem is severe enough to prevent the network from handling a large number of users and covering large areas. One way to avoid this problem is to equip the stations with two Wi-Fi radios. By tuning the radios to different (e.g. orthogonal) channels, the stations can communicate simultaneously with minimal interference. This approach is commercially attractive because it is relatively inexpensive nowadays to incorporate two off-the-shelf radios in each station.

The contribution of this paper is the detailed evaluation study of benefits and drawbacks of dual Wi-Fi radios for multi-hop communication in the described scenarios and architecture, given constraints and metrics such as quality-of-throughput, reliability, scalability, fairness, and power consumption. The evaluation is based on detailed simulation. For this purpose, we extended our Jemula802 network simulator [1] with the current state-of-art multi-channel radio resource management algorithms (channel allocation and routing). For two simple network topologies (line and grid), we illustrate throughput gains archived with dual radio devices relative to single-channel single radio devices. Throughput fairness among competing flows is also studied. Based on a Wi-Fi energy consumption model that we carefully calibrated with real-life measurements, we evaluate the energy efficiency (Joules per successfully transmitted byte) of the single and dual radio devices. All tests are performed both with UDP and TCP data flows and differences in performance gains are noted.

The studied architecture relies on protocols and mechanisms developed for mobile ad hoc networks (MANETs) and wireless mesh networks (WMNs). The authors in [2] identify some of the main challenges in multi-channel mesh networks including channel allocation and packet routing. In [3], [4], the authors propose channel allocation algorithms that exploit multiple channels with single radio devices. The solutions either require changes to the current commodity hardware [3] or rely on a pre-defined channel hopping sequence, which does not allow for optimization based on current traffic conditions [4]. Therefore, multi-radio devices have been considered and various channel allocation protocols for such devices proposed in [5]-[10]. Most routing protocols for mobile ad hoc and mesh networks are based on the dynamic source routing (DSR) [11] or ad hoc on-demand distance vector (AODV) [12] routing protocols. The performance of the AODV protocol in multi-channel multi-radio networks has been studied in [13]-[15]. Approaches that aim to solve the channel allocation and routing problems jointly through a complex cross-optimization are described in [16], [17]. A practical approach is to keep the channel allocation and routing decisions separated, as proposed in [18] and described in Sections III.C and III.D of this paper.

The remainder of this paper is organized as follows: Section 2 gives a brief overview of multiple channel access mechanisms in 802.11 networks and potential benefits of dual radios. Section 3 describes the extensions made to the network simulator Jemula802. Section 4 presents the results of the dual radio performance evaluation. Finally, Section 5 concludes this paper.

II. CHANNEL ACCESS IN 802.11

The performance of the 802.11 Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) protocol can be substantially improved with the use of dual radio devices. Consider a network of four stations equipped with single Wi-Fi radio interfaces shown in Fig. 1. Each station has a unique station address and a network ID. A traffic stream originates from Station 1's traffic generator (TG) and ends at Station 4's back-off entity (BE). The maximum radius at which Station 1's transmission can be sensed is marked; we assume this to be the isotropic *transmission range* of Station 1. Each station in Fig. 1 can sense the transmissions of its immediate neighbor(s). If Station 1 sends a packet to Station 4, then Stations 2 and 3 have to relay the packets. If all stations use the same radio channel, a problem arises when two stations, which are not in the sensing range of each other (e.g. Stations 1 and 3), try to send packets at the same time. Since the stations do not hear each other's transmissions, the listen-before-talk policy of CSMA/CA becomes ineffective. Station 3, for example, may try to transmit on the channel while Station 1 is already transmitting. This results in a collision at Station 2, which is then unable to decode the packet sent by Station 1. To alleviate this problem, 802.11 standard defines an optional RTS/CTS (request to send / clear to send) handshake mechanism. However, RTS/CTS is rarely used since shown to be ineffective in most practical scenarios [19].

Consider now a solution where each of the stations has two Wi-Fi interfaces: we call such stations *dual radio devices*. Dual radios allow two channels to be used at the same time at each station. In the network shown in Fig 1, the hidden station problem can be avoided if channels used on different hops are selected carefully. For example, if Stations 1 and 3 are sending on orthogonal channels, there is no hidden station problem at Station 2. However, the problem of channel assignment to the dual radio devices is not trivial. Static assignment cannot be optimal if stations are mobile. Even when stations are not moving, the authors in [20] show that the problem of optimal channel assignment is np-complete. Therefore, heuristics that approximate the optimal solution should be used. In our evaluation, we use a channel allocation algorithm proposed in [18]. The algorithm assumes that every station has its first radio (semi)-statically assigned to one channel, which is then used only for receiving data. The second radio may switch between channels at any time and it is used for sending data: It selects a channel based on the intended receiver. The first radio is referred to as *fixed* while the second is *switchable*. This distinction greatly simplifies the channel allocation problem. When neighboring stations select different channels for their fixed (receiving) radios, the probability of the hidden station problem is lowered. Therefore, one may expect that the maximum achievable throughput with dual radio devices is *more* than twice the throughput achievable with single radio devices: The gain is due to the simultaneous use of two channels and due to decreased collision probability.

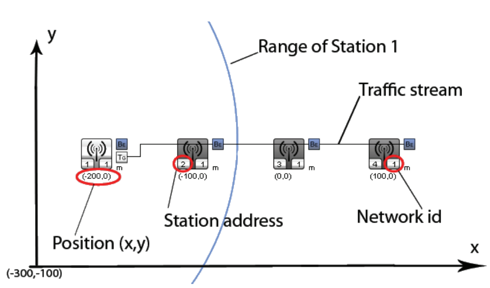

Figure 1. This simple example illustrates the notation used in scenarios.

III. WI-FI SIMULATION TOOL AND EVALUATION MODELS

In order to evaluate the performance of dual radios, we extended our network simulator Jemula802 [1] with multi-channel radio resource management and link adaptation algorithms, as well as radio propagation and interference models. Each of these models is described in more details in the following.

A. Propagation and Interference Models

Originally, the connectivity model in Jemula802 was a simple unit disk model. In the model, a node can receive/transmit from/to all other nodes that are within a fixed radius. Link quality (hence, achievable data rate) did not depend on the distance, as long as two nodes are within the radius. To account for signal attenuation, we added a free space propagation model to Jemula802. In this model, received power P_{RX} decreases with the distance d to the transmitting station as

$$P_{RX} = P_{TX} - P_0 - L(d), \qquad (1)$$

where P_{TX} is the transmit power, P_0 is the received power at the reference distance $d_0 = 1$ m, which is computed as

$$P_0 = 20 \cdot \log_{10}(4\pi/\lambda), \qquad (2)$$

and $L(d)$ is the path loss given by $10 \cdot n \cdot \log_{10}(d/d_0)$ if $d \geq d_0$ and zero otherwise. The path loss exponent for free space propagation is $n = 2$. For simplicity, shadowing and multipath fading are not included in the propagation modeled.

To capture the effect of cross-channel interference, we implemented the model proposed in [21]. In this model, the cross-channel interference $I(a, b)$ caused to the receiver on channel a by neighbouring transmissions on channel b is

$$I(a, b) = f(a, b) \cdot P_{RX,b}, \qquad (3)$$

where $P_{RX,b}$ is the total received signal power on the channel b and $f(a, b)$ is a cost function defined in [21]. The signal to noise plus interference ratio (SNIR) at the receiver is then equal to the received power from sender i on channel a divided by the noise plus interference of other all other transmissions:

$$SNIR_a(i) = \frac{P_{RX,a}(i)}{N + \sum_{j \neq i} P_{RX,a}(j) + \sum_{b \neq a} I(a, b)}, \qquad (4)$$

where $\sum_{j \neq i} P_{RX,a}(j)$ is the intra-channel interference and

$\sum_{b \neq a} I(a, b)$ is the cross-channel interference on channel a. The SNIR may change during a reception of a packet. The simulator keeps track of the lowest SNIR during a packet reception. Whether this lowest SNIR is sufficient for a correct packet reception is decided based on the bit error rate (BER) model described in [22]. The model specifies BER as a function of SNIR for different modulation and coding schemes (MCS) used in 802.11 (Fig. 3). The more complex the coding scheme, the higher is the required SNIR to receive bits correctly. The packet error rate (PER), which is the probability that at least one bit of the packet is erroneous, is then computed as $PER = 1 - (1 - BER)^L$, where L is the packet length in bits.

B. Link Adaptation

An algorithm is needed by which a transmitter selects a more or less robust modulation and coding scheme (MCS) depending on the current link quality. This is known as *link adaptation* in 802.11. The least robust MCS (64QAM 3/4 in Fig. 2) provides the highest link rate (54 Mb/s in 802.11g), but suffers from high packet error rates at low SNIRs. In the simulator, we assume that the sender can estimate the current SNIR at the receiving station before every packet transmission. For this SNIR, the packet error rate (PER) is calculated for each MCS based on the BER in Fig. 2. The algorithm then chooses an MSC that provides the highest data rate with the PER lower than 1%. Of course, such an ideal closed-loop algorithm is only possible in a simulator, since the SNIR is looked up directly at the receiving station. The design and performance of practical link adaption algorithms is not addressed in our study. For example, the authors in [23] propose a link adaptation algorithm with a SNIR feedback from the receiver based on modified RTS/CTS mechanism.

C. Channel Allocation

We implemented the channel allocation algorithm proposed in [18]. The algorithm allows to use an arbitrary number of channels with dual radio devices. The main idea is to use one of the radios for receiving and the other one for transmitting. The receiving radio is semi-statically assigned a channel, which is may only change infrequently. This receiving radio is called *fixed* radio. The channel number of the fixed radio is periodically advertised via hello packets to all neighbors on all available channels. A station initially selects a fixed channel at random when it joins a network. It keeps track of the fixed channels used by neighbors by listening to the broadcasted hello packets. Then it selects a channel that is used by fewest of its neighbors and assigns it to its fixed radio. This balances the number of stations on each channel. A station only changes its fixed channel with a certain probability < 1 when it recognizes that there is a channel used by fewer stations than the current one. The transmitting (*switchable*) radio has to switch to the receiving (*fixed*) channel of the intended receiver of the transmitted packet. Switching channels introduces some delay (~1 ms) until the radio is tuned to the new channel. Due to the delay, the switchable radio does not switch at every single packet, but it stays on the same channel at least for a few tens of milliseconds. Packets for other channels/receivers are queued until the transmitting radio switches to the correct channel. Although this may introduce some queuing delay at every hop, it is still more time efficient than per-packet switching.

Fig. 3 shows an example of a channel assignment for three stations. Stations 1 and 3 are both sending packets to Station 2. Since Station's 2 fixed radio is receiving on Channel 1, both stations have to tune their switchable radios to the Channel 1. Station 2 uses its switchable radio to transmit packets to Station 1. If Station 2 needs to send a packet to Station 3, it has to switch its switchable radio to Channel 3.

D. Packet Routing

We extended Jemula 802 with a multi-channel version of the popular AODV [12] routing protocol. Original version of AODV aims to discover routes with a minimum number of hops. When multiple channels and radios are used, the route with fewest hops is not necessarily the best route. For example, a three-hop route that uses different channels on each hop may be better than a two-hop route that uses the same channel on both hops because fewer collisions will occur. Hence, AODV needs a new routing metric for multi-channel scenarios. Such metric, called multi-channel routing (MCR), has been proposed in [18]. MCR replaces the default hop number metric of the standard AODV with a metric that also takes into account channel diversity and channel switching cost. We implemented the MCR metric in Jemula802 (see [18] for details about the metric).

E. Energy Consumption Model

We implemented a Wi-Fi power consumption model in the simulator to estimate the battery drain of single and dual radios. Each state of a radio interface is assigned a different power consumption level. The model distinguishes between three states: sending (TX), receiving (RX), and idle (IDLE),

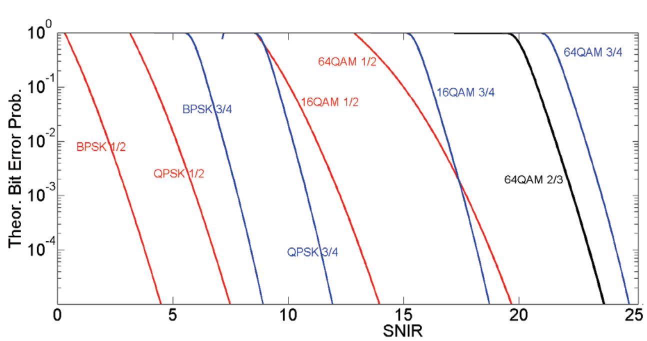

Figure 2. BER for different modulation and coding schemes in 802.11 according to [22].

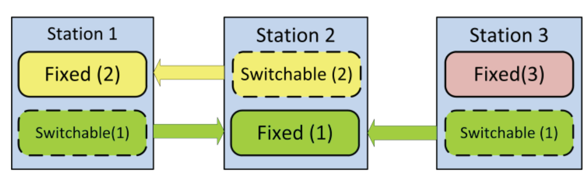

Figure 3. An example of channel assignment to dual-radio devices: channel numbers are in brackets.

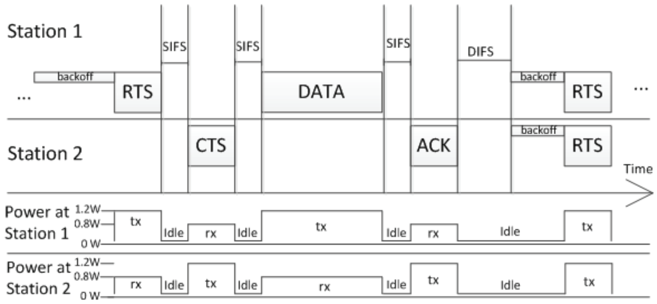

Figure 4. Power levels during a packet transmission/reception in 802.11.

which are associated with power levels P_{TX}, P_{RX}, and P_{IDLE}, respectively. The energy consumption in each state is then computed as a product of the state's power level and the time that the interface spends in that state. The total energy consumption E_{TOT} is then given by:

$$E_{TOT} = P_{TX} \cdot t_{TX} + P_{RX} \cdot t_{RX} + P_{IDLE} \cdot t_{IDLE} \quad (5)$$

The power levels are assumed to be $P_{TX} = 1.28$ W, $P_{RX} = 0.82$ W, and $P_{IDLE} = 0.045$ W, according to the measurements in [24]. The time spent in the TX/RX state is the sum of the transmission/reception times of all sent/received packets. If it is not sending or receiving, a radio is in the IDLE state. Fig. 4 shows the power levels at two communicating stations during a packet exchange.

IV. EVALUATION

We present dual radio performance results for two scenarios: line and grid scenarios. In all setups with dual radios, we assume that three orthogonal channels in the 2.4 GHz spectrum are used. In setups with single radios, we assume that only one channel is used. The performance is measured in terms of throughput, throughput fairness (given by the Jain's fairness index [25]), and energy consumption. By default, the scenarios were run for ten minutes of network traffic. Each scenario was run five times with different seeds. The coefficient of variation (CoV) of throughputs and energy consumptions in different runs is indicated.

A. Basic Line Scenario

Here we describe the results for the line scenario shown in Fig. 5, which was discussed in the Introduction. In the scenario, a single stream of application data packets originates from the traffic generator of Station 1 and ends at the back-off entity of Station 4. The data stream saturates the connection, which means that whenever the MAC layer queue of Station 1 is empty, a new packet is generated and added to the queue. We first consider a case where UDP is used to transport the data. The stream is unidirectional since no acknowledgements are sent back to Station 1. The packets are relayed at Stations 2 and 3. Every station can only sense transmissions from its immediate neighbors. Stations 1 and 3, for example, are not able to hear each other and may attempt to transmit at the same time. In a single radio single-channel configuration, this leads to a hidden station problem at Station 2.

Figure 5. The line scenarios with a single data stream from Station 1 to Station 4. The distance between two neighboring stations is 100m.

In Table I, we compare the throughputs and energy consumptions with single and dual radios. The throughput of dual radio configuration is 240% higher and it uses only 56% more energy compared to the single radio configuration. One may think that with dual radio devices, the throughput can only be at most twice as high as in with the single radio devices. However, when using two channels instead of one, the collision probability also decreases. The decreased collision rate leads to fewer packet retransmissions and, therefore, the dual radio throughput can be more than twice as high as with single radio. For the same reason, the energy consumption per successfully transmitted packet is lower with dual radios. This is true for high traffic loads (e.g. when the network is saturated). When the traffic is light, the second radio can be switched off to save energy.

We also considered the case where TCP is used to transport the packets between Stations 1 and 4. Now there is an additional stream of acknowledgement packets flowing from Station 4 back to Station 1. Because of that, there are now hidden station problems both at Stations 2 and 3 and not only with the single radio configuration, but also with the dual radio multi-channel configuration. This could be avoided if different channels were used for each hop and each direction. However, the channel allocation protocol (see Section III.C) only allows stations to receive data on one of their radios. Therefore, Stations 1 and 3 have to send data to Station 2 using Station 2's fixed/receiving channel. Since Stations 1 and 3 are not in sensing range of each other, there is a hidden station problem at Station 2 even with dual radios. With single radios the problem is more severe than with dual radios since collisions occure not only when Stations 1 and 3 transmit simultaniously to Station 2, but also when Station 1 transmits to Station 2 and Station 3 transmits to Station 4. Similarly, a hidden station problem is caused at Station 3 by Stations 2 and 4. The throughputs and energy

TABLE I. COMPARISON OF THROUGHPUT AND POWER CONSUMPTION FOR THE BASIC LINE SCENARIO.

	Single Radio				Dual Radio				Comparison	
	T [Mb/s]		E [J]		T [Mb/s]		E [J]		ΔT	ΔE
	Mean	CoV	Mean	CoV	Mean	CoV	Mean	CoV		
UDP	3.30	0.139	1359	0.0007	11.25	0.145	2175	0.1220	+240 %	+56 %
TCP	2.12	0.138	1068	0.0100	7.95	0.138	1969	0.0005	+275 %	+84 %

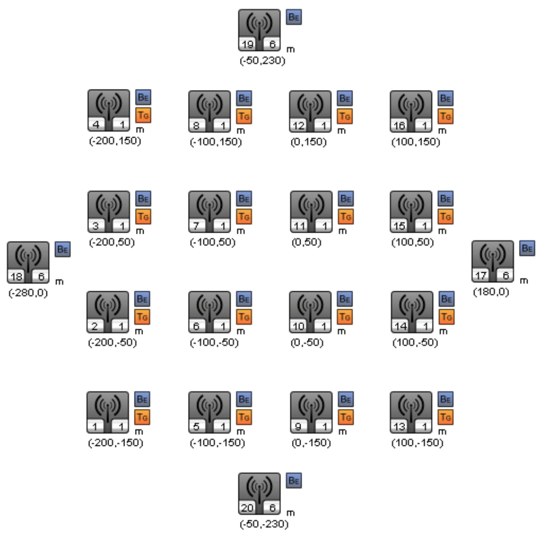

Figure 6. Layout of the grid scenario with 16 stations and 4 access points.

consumptions with single and dual radios are compared in the bottom row of Table I. Due to the hidden station problems, the TCP throughput is lower than the UDP throughput, both with single and dual radios. The dual radio configuration still outperform the single radio configuration: the throughput is 275% higher at the expense of an 84% increase in energy consumption. This makes dual radios still more energy efficient than single radios.

B. Grid Scenario

In this scenario, 16 stations are aligned in a grid as shown in Fig. 6. The distance between two neighboring stations in the grid is 100 m, allowing them to transmit to each other at the highest possible data rate (54 Mb/s) at the chosen transmit power level, as described in Section III.B. Stations that are not immediate neighbors in the grid can still communicate directly, but at a lower data rate. There are four access points (APs) at the edges of the grid (in some some setups only one of them is available). All APs have wired conections to a server, which is not shown in Fig. 6. Therefore, data destined to or originating from the server can be routed through any of the APs.

We consider several different setups depending on the number of active stations and available access points. First we consider a setup where only 4 out of 16 stations (Stations 4, 7, 10, 13, which are on the diagonal of the grid) have data to transmit. Each of the stations generates 16 Mb/s of uplink UDP traffic destined to the server. Than we consider a setup where 16 stations are active and each of the stations generates 4 Mb/s. The total traffic load of 64 Mb/s saturates the network in both setups in order to gauge the maximum achivable throughput. The generated traffic flows are routed to the server either through any of the four APs, if they are all available, or through a single AP, if only one of them is available. Therefore, we distinguish between four setups: four flows served by four APs, four flows served by one AP, 16 flows served by four APs, and 16 flows served by one AP. Table II shows the comparison of the single and dual radio configurations in terms of the average throughput per flow T, throughput fairness J, and total energy E consumed by all stations in the grid.

The dual radio performed significantly better than the single radio both in terms of average throughput and energy consumed per succesufuly transmitted packet. Jain's fairness index also indicates that the throughput allocation to the flows is more fair with dual radios. The performance improvement are due to hops being able to use different, non-interfering channels with dual radios. The throughput gain of dual radios ΔT is higher with 16 flows than with four flows, which confirms that major benefits of dual radios are for congested scenarios with many stations attempting to transmit at the same time. Dual radio configuration scales better when the number of flows increases. For example, with four APs, the average throughput per flow with single radios decreases by 90 % (from 0.75 Mb/s to 0.076 Mb/s) when the number of flows increases from 4 to 16. With dual radios, the throughput decrease is still significant, but less severe (from 2.16 Mb/s to 0.74 Mb/s). For setups with the same number of flows, the throughput gains with dual radios are higher with four APs compared to one AP. With one AP, the gains are limited by the capacity of the receiving radio at that single AP. It is interesting to notice that, with single radios, the average throughput per flow is higher with one AP than with four APs. However, with one AP, the throughput is very unevenly distributed among the flows: The stations that are close to the AP achieve much higher throughputs than the stations that are further away. Fig. 7 shows the throughput per stations with 16 active stations and one AP for single and dual radios. Although the not perfectly fair, the dual radio configuration provides certain throughputs to the stations that are not immediate neighbours to the AP, which are otherwise starved with single radios. Dual radio configuration is more power efficient in all setups. For example, with 16 flows and one AP, the throughput increased by 414 % while the energy consumption increased only by 68 % This increased power efficency is due to lower number of packet retransmissions with dual radios.

TABLE II. COMPARISON OF UDP THROUGHPUT AND POWER CONSUMPTION FOR THE GRID SCENARIO.

#AP	#Tx	Single Radio					Dual Radio					Comparison	
		T [Mb/s]		JF	E[J]		T [Mb/s]		JF	E [J]		ΔT	ΔE
		Mean	CoV		Mean	CoV	Mean	CoV		Mean	CoV		
4	4	0.750	0.137	0.986	7514	0.001	2.165	0.138	0.998	9308	0.082	+188 %	+24 %
	16	0.076	0.153	0.540	9086	0.003	0.745	0.148	0.623	14705	0.012	+880 %	+61 %
1	4	0.777	0.186	0.806	6195	0.048	1.487	0.138	0.921	8607	0.071	+91 %	+38 %
	16	0.092	0.150	0.220	6594	0.025	0.473	0.179	0.371	11088	0.013	+414 %	+68 %

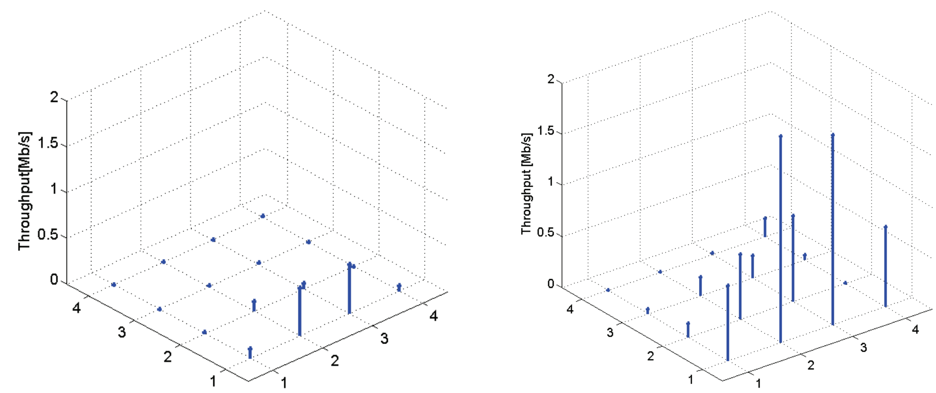

Figure 7. Throughput per flow with 16 flows and one AP for (a) single radios. (b) dual radios.

We repeated the simulations using TCP as a transport protocol instead of UDP; results are shown in Table III. With four APs, the throughput gain of dual radios was not as high as with UDP. Due to packet acknowledgements, the overall contention for the channels increases. Our dual radio setup uses three orthogonal channels, which are no longer sufficient to avoid collisions. In the setup with only one AP, the throughput with dual radios is even lower than with single radios. The reason is high overhead of multi-channel routing. With the single AP, routes are longer in average (they contain more hops). Therefore, link breaks, which trigger new route discoveries, are more frequent. With dual radios and multiple channels it takes longer time to discover new routes: RREQ messages need to be broadcasted on all available channels. Channel switching delay adds to the discovery time. This may lead to packet losses and TCP time-outs, which have negative impact on the throughput. However, the average troughput per flow does not paint the full picture. In the single radio setup with one AP and four active stations, the station closest to the AP received more than 50 % of the total throughput. With dual radios, the throughput allocation is more fair as indicated by the Jain's fairness index.

V. Conclusions

We extended the Jemula802 simulator with some of the current state-of-art multi-channel routing and channel allocation algorithms in order to study the performance of dual Wi-Fi radio devices. The comparison with the single radio devices has shown the following: i) Dual radios may reduce the hidden station problem because fewer stations are sending data on the same channel. ii) Forwarding of data packets on orthogonal channels, which is enabled by dual radios, improves the overall throughput. iii) When traffic load is high, dual radios are more energy efficient than single radios since fewer packets are re-transmitted. iv) The performance with dual radios scales better when the number of devices increases.

During the tests, we identified the following issues that need to be addressed: i) An algorithm is needed to decide when to switch off the second radio. Such algorithm may help reduce the energy consumption of dual radio devices when the traffic load is low and, hence, the second radio is not needed. ii) The problem of high multi-channel routing overhead needs to be addressed. The overhead is likely to create even more problems in highly mobile networks where links break frequently. iii) Falsely detected link breaks due to collisions are harmful since new routes have to be rediscovered. Better criteria are needed to decide whether a link is broken or just overloaded due to collisions. iv) The multi-channel routing (MCR) metric could be adapted to the QoS requirements of the data flow to further improve the performance.

References

[1] Jemula802. http://jemula802.origo.ethz.ch/. [Accessed May 26, 2011].

[2] P. Kyasanur, J. So, C. Chereddi, and N. H. Vaidya,. "Multichannel mesh networks: challenges and protocols," IEEE Wireless Communications, vol. 13, no. 2, pp. 30–36, 2006.

[3] J. So and N. Vaidya, "Multi-channel MAC for ad hoc networks: Handling multi-channel hidden terminals using a single transceiver," *Proc. ACM MobiHoc*, 2004.

[4] P. Bahl and R. Chandra, "SSCH: Slotted seeded channel hopping for

TABLE III. Comparison of TCP throughput and power consumption for the grid scenario.

#AP	#Tx	Single Radio					Dual Radio					Comparison	
		T [Mb/s]		JF	E[J]		T [Mb/s]		JF	E [J]		ΔT	ΔE
		Mean	CoV		Mean	CoV	Mean	CoV		Mean	CoV		
4	4	0.865	0.140	0.999	5641	0.003	1.100	0.138	0.999	5873	0.036	+27%	+4 %
	16	0.172	0.141	0.911	6730	0.013	0.477	0.137	0.877	8298	0.011	+177 %	+23 %
1	4	0.773	0.141	0.583	5641	0.022	0.433	0.163	0.771	3080	0.097	−44 %	−45 %
	16	0.499	0.153	0.245	4970	0.004	0.309	0.140	0.465	4629	0.029	−38 %	−7 %

capacity improvement in IEEE 802.11 ad-hoc wireless networks," *Proc ACM Mobicom*, 2004.

[5] H. Skalli, S. Ghosh, S. K. Das, and L. Lenzini, "Channel assignment strategies for multi-radio wireless mesh networks: issues and solutions," *IEEE Communications Magazine*, vol. 45, no. 11, pp. 86–95, 2007.

[6] A. Dhananjay, H. Zhang, J. Li, and L. Subramanian, "Practical, distributed channel assignment and routing in dual-radio mesh networks," *Proc. ACM Sigcomm*, Barcelona, Spain, 2009.

[7] A. H. M. Rad and V. W. S. Wong, "Joint channel allocation, interface assignment and MAC design for multi-channel wireless mesh networks," *Proc. IEEE Infocom*, Anchorage, USA, 2007.

[8] A. P. Subramanian, H. Gupta, S. R. Das, and J. Cao, "Minimum interference channel assignment in multi-radio wireless mesh networks," *IEEE Trans. Mobile Computing*, vol. 7, no. 12, pp. 1459–1473, 2008.

[9] K. Xing, X. Cheng, L. Ma, and Q. Liang, "Superimposed code based channel assignment in multi-radio multi-channel wireless mesh networks," *Proc. ACM MobiCom*, Montreal, Canada, 2007.

[10] K. N. Ramachandran, E. M. Belding, K. C. Almeroth, and M. M. Buddhikot, "Interference-aware channel assignment in multi-radio wireless mesh networks," *Proc. IEEE Infocom*, Barcelona, Spain, 2006.

[11] D. B. Johnson, D. A. Maltz, and J. Broch, "DSR: The dynamic source routing protocol for multi-hop wireless ad hoc networks," *Ad Hoc Networking*, edited by C. E. Perkins, Chapter 5, pp. 139–172, Addison-Wesley, 2001.

[12] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing," *IETF RFC 3561* (experimental), July 2003.

[13] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," *Proc. ACM Mobicom*, Philadelphia, USA, 2004.

[14] T. Liu and W. Liao, "Interference-aware QoS routing for multi-rate multi-radio multi-channel IEEE 802.11 wireless mesh networks," *IEEE Trans. Wireless. Comm*., vol. 8, no. 1. pp. 166–175, 2009.

[15] A. A. Pirzada, M. Portmann, and J. Indulska, "Performance analysis of multi-radio AODV in hybrid wireless mesh networks," *Elsevier Computer Comm.*, vol. 31, no. 5, pp. 885–895, 2008.

[16] M. Alicherry, R. Bhatia, and L. E. Li, "Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks," *Proc. ACM MobiCom*, Cologne, Germany, 2005.

[17] H. S. Chiu, K. L. Yeung, and K.-S. Lui, "J-CAR: an efficient joint channel assignment and routing protocol for IEEE 802.11-based multi-channel multi-interface mobile ad hoc networks," *IEEE Trans. Wireless Comm.*, vol. 8, no. 4, pp. 1706–1715, 2009.

[18] P. Kyasanur and N. H. Vaidya, "Routing and link-layer protocols for multi-channel multi-interface ad hoc wireless networks," *ACM Mobile. Computing and Communication Review*, vol. 10, pp. 31–43, 2006.

[19] K. Xu, M. Gerla, and S. Bae, "How effective is the IEEE 802.11 RTS/CTS handshake in ad-hoc networks," *Proc. IEEE GLOBECOM*, Taipei, China, 2002.

[20] W.K. Hale, "Frequency assignment: Theory and applications," *Proceedings of the IEEE*, vol. 68, pp. 1497–1514, 1980.

[21] B.-J. Ko, V. Misra, J. Padhye, and D. Rubenstein, "Distributed channel assignment in multi-radio 802.11 mesh networks," *Proc. IEEE WCNC*, Hong Kong, China, 2007.

[22] S. Mangold, S. Choi, and N. Esseling, "An error model for radio transmissions of wireless LANs at 5 GHz," *Proc. 10th Aachen Symposium on Signal Theory*, Aachen, Germany, 2001.

[23] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," *Proc. ACM MobiCom*, Rome, Italy, 2001.

[24] D. Halperin, B. Greensteiny, A. Shethy, and D. Wetherall, "Demystifying 802.11n power consumption," *Proc. Int. Conf. Power Aware Computing and Systems* (*HotPower*), Berkeley, USA, 2010.

[25] R. Jain, D. M. Chiu, W. R. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," DEC, Tech. Report TR-301, 1984.