

entire duration of the experiment. The links that $f_{1 \rightarrow 18}$ traverses do not share airtime with the links that $f_{20 \rightarrow 9}$ traverses, so they do not contend for bandwidth. This can be seen in Figure 8(a): the links $20 \rightarrow 6$ and $6 \rightarrow 9$ get an airtime limit of 60%, and the reverse links get (not shown) an airtime limit of 20% each, totaling 100%.

We start another flow $f_{11 \rightarrow 18}$ at 30 seconds, and its links are in the same neighborhood as the links of $f_{1 \rightarrow 18}$. Figure 8(a) shows the airtime-limit allocation that CNA assigns. Before $f_{11 \rightarrow 18}$ starts, links $1 \rightarrow 2$ and $2 \rightarrow 18$ have the same airtime limit. Once $f_{11 \rightarrow 18}$ starts, however, the airtime-limits on both its links are reduced, and but they each get an airtime-limit twice that of link $11 \rightarrow 1$. That is because these links carry two flows ($f_{1 \rightarrow 18}$ and $f_{11 \rightarrow 18}$) while link $11 \rightarrow 1$ carries only one flow ($f_{11 \rightarrow 18}$). This illustrates that CNA assigns airtime-limits proportional to the number of flows traversing a link, resulting in fair throughput allocation to all the flows (Figure 8(b)). A closer look at the figure reveals that $f_{11 \rightarrow 18}$ gets a slightly lower throughput because it has a slightly higher RTT.

Auto-rate adaptation. Now consider $f_{20 \rightarrow 6}$ and observe that auto-rate adaptation results in about 11 PHY-layer rate changes before 60 seconds (Figure 10(a)). Interestingly, this does not affect the throughput of this flow. In 802.11, the preamble is transmitted at 2Mbps. Even if each packet waits for the average initial backoff period, it turns out that, with our 512-byte payloads and a TCP ACK, the airtime difference between a packet sent at 11Mbps and at 5.5 Mbps is a little over 20%. This difference is comparable to the throughput fluctuations shown in Figure 9(b).

We introduce, a little after 60 seconds, $f_{9 \rightarrow 17}$. This flow passes through neighborhoods that overlap with $f_{20 \rightarrow 9}$, but not with the other flows. Although one would expect these two flows to get equal throughput because CNA allocates approximately the same airtime-limits to all four links (Figure 9(a)), they don't (Figure 9(b)). Flow $f_{20 \rightarrow 9}$ gets much lower goodput. Some of this throughput loss can be explained by the brief choice of a 2Mbps rate by the auto-rate mechanism (Figure 10(a)).

More interesting, however is the observation that the introduction of the new flow increases channel losses and the average number of retries on $20 \rightarrow 6$ (Figure 10(b)), causing it to expend airtime on retransmissions with a consequent loss of throughput. We believe there is a very subtle reason for this. In Section 4, we explained how, using device-level registers, CNA is able to correctly *detect* internal interference: i.e., interference from nodes whose packets it cannot decode. When all such interferers are within two hops, CNA correctly allocates airtime to them. However, in the relatively infrequent event that an interferer is three hops away, CNA's signaling is unable to appropriately assign airtime. In this example, we conjecture that transmissions on $10 \rightarrow 17$, which is three hops away from $20 \rightarrow 6$, causes interference at node 6, resulting in reduced goodput on $f_{20 \rightarrow 9}$. Regardless, $f_{9 \rightarrow 17}$ gets its fair share of airtime and goodput,

and is unaffected.

Flow departure. We have illustrated CNA's adaptation to flow arrival above. We now explore flow departure, from 90 to 180 seconds, in Figure 7. Around 90 seconds, we introduce one more flow $f_{2 \rightarrow 17}$, and it runs for 60 seconds. This flow traverses the middle of topology, causing all 5 active flows to traverse at least one link in the neighborhood of link $3 \rightarrow 4$. Figure 11(a) shows how the airtime-limits change on some links in $N_{3 \rightarrow 4}$. Adding $f_{2 \rightarrow 17}$ results in a reduction of airtime-limits for all links, but before $f_{2 \rightarrow 17}$ starts and after it departs, the airtime-limit allocations are roughly same. The factor of two difference in the airtime-limit for links $1 \rightarrow 2$ and $2 \rightarrow 18$ exists because these links carry two flows. Figure 11(b) shows the throughput each flow achieves in this interval. As explained above, $f_{20 \rightarrow 9}$ gets lower throughput because of increased retransmissions. The throughput differences between the other flows are due to differences in RTT (the 3 hop flows $f_{2 \rightarrow 17}$ and $f_{11 \rightarrow 18}$ get lower throughput) and in packet loss rates.

Short Flows. CNA uses per-packet signaling to achieve responsiveness. The interval from 180 and 270 seconds in Figure 7 demonstrates this. In this interval, two long flows $f_{1 \rightarrow 18}$ and $f_{9 \rightarrow 17}$ are present, and two short flows $f_{3 \rightarrow 4}$ and $f_{3 \rightarrow 20}$ start around at 210 and 240 respectively. Each runs for 10 seconds, and they transmit approximately 611 KB and 370 KB respectively. As Figures 12 and 7 show, CNA quickly responds to each new flow, reducing the airtime of the contending links and, as a result, the throughput of the long flows. When the short flows leave, the absence of their packets is quickly noticed, and the long flows regain their airtime-limits and throughputs.

Unused airtime redistribution. We use the interval between 300 and 350 seconds in Figure 7 to demonstrate the efficacy of CNA's unused airtime re-distribution mechanism. Around at 310 seconds, the existing flow $f_{9 \rightarrow 17}$ ends, and another low-rate flow $f_{9 \rightarrow 17}$, which sends only at 8 Kbps starts. Links $9 \rightarrow 10$ and $10 \rightarrow 17$ have unused airtime, and this is redistributed to other links in their neighborhood. Thus, $7 \rightarrow 8$'s airtime-limit is now increased to 30% (Figure 13), and flow $f_{7 \rightarrow 8}$ achieves greater throughput while $f_{9 \rightarrow 17}$ runs. Notice that $f_{7 \rightarrow 8}$ incurs a TCP timeout in the middle, which causes a 1-second long downward spike in its instantaneous throughput; CNA cannot, of course, completely eliminate TCP timeouts, but greatly reduces their occurrence.

Responsiveness to route change. The time interval between 340 and 380 seconds in Figure 7 demonstrates CNA's responsiveness to route changes. During this interval, two flows, $f_{7 \rightarrow 8}$ and $f_{4 \rightarrow 20}$, are present. Initially, $f_{4 \rightarrow 20}$ uses the route $4 \rightarrow 6 \rightarrow 20$. Around 360 seconds, we change it to $4 \rightarrow 3 \rightarrow 5 \rightarrow 20$. As the Figure 14 shows, CNA re-assigns airtime-limits, and both flows achieve relatively fair throughput. There is a short transient where air-time limits are lower than the steady-state value: during this interval, CNA has not detected the departure of the flow along the old route,

but has detected the arrival of the flow on the new route.

Adapting to external interference. CNA explicitly measures external interference, and sets airtime-limits based on the available channel airtime. To demonstrate this, consider the time interval between 410 and 470 seconds in Figure 7, where there are three flows running: $f_{7 \rightarrow 8}$, $f_{3 \rightarrow 4}$ and $f_{2 \rightarrow 17}$. At about 420 seconds, we turned on a microwave oven (placed near nodes 1 and 2, Figure 6) for 30 seconds.

Figure 16(a) and 16(b) shows the computed available channel airtime and airtime-limits. Link $7 \rightarrow 8$ does not detect external interference, but its airtime-limit is reduced, because it is in the neighborhood of $2 \rightarrow 3$ and $3 \rightarrow 4$. Therefore, throughputs for all three flows are decreased. Note that in this period one three-hop flow $f_{2 \rightarrow 17}$ is competing with two one-hop flows $f_{7 \rightarrow 8}$ and $f_{3 \rightarrow 4}$. Link $3 \rightarrow 4$ gets proportionally higher airtime because it carries two flows.

For the time interval between 500 and 530 seconds, we turned on the microwave oven for 20 seconds. During this time interval, flows $f_{1 \rightarrow 18}$, $f_{11 \rightarrow 18}$ and $f_{20 \rightarrow 9}$ are present. As Figure 16(c) and 16(d) show, links $20 \rightarrow 6$ and $6 \rightarrow 9$ do not get reduced airtime-limits, since they are outside the neighborhood affected by the external interference. Notice though, in Figure 7, that there is throughput gap between $f_{1 \rightarrow 18}$ and $f_{11 \rightarrow 18}$. The external interference causes increased retransmissions on $f_{11 \rightarrow 1}$ shown in Figure 15, so the throughput of $f_{11 \rightarrow 18}$ is reduced. This results in unused airtime on $1 \rightarrow 18$, which is taken up by $f_{1 \rightarrow 18}$.

6. OPTIMALITY

How far off from the optimal is CNA? The most comprehensive answer to this question can be given if one computes the achievable rate region of representative, real-world mesh topologies and identifies the point corresponding to CNA on this regions.³ Prior work [14] has developed a theoretical framework which, given information about the interference graph and the link loss rates in a topology, can compute the achievable rate region for an 802.11 mesh network. Unfortunately, this framework cannot be directly applied to real-world wireless networks, such as the ones we have used in Section 5, because in the real world, both the interference graph and the loss rate of each link change over time (as they did during our experiments). We could have attempted to use the “most likely interference graph” and the “average loss rate of a link” to compute the optimal rates. However, it is unclear whether, and to what extent, this calculation would under-estimate or over-estimate the actual achievable-rate region. So, we leave such efforts for future work, and instead use simulation to study the optimality of CNA in a number of topologies for which the achievable rate region can be computed using [14]. To keep the exposition simple, we find the max-min rate allocation on the boundary of

³The achievable rate region is the set of all flow-rates which do not blow up any queues. Its boundary corresponds to optimal rate allocations, e.g. the max-min rate allocation is one of these optimal points.

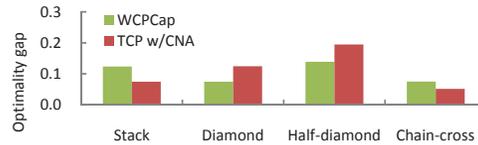


Figure 17: The Optimality Gap

the achievable rate regions and compare the max-min rates, referred to as optimal from now on, to those of CNA. We emphasize that the max-min rates under an optimal scheduler are not the same with the max-min rates under 802.11, and that our focus is on the latter. For a comparison of these max-min rates the interested reader is referred to [15].

We implemented CNA on Qualnet 3.9.5 [3]. In addition to comparing CNA and optimal, we also compare the rate allocation obtained by WCPCap [24], the state-of-the-art rate-control design for mesh network transport which closely approximates a max-min rate controller while being fully distributed (We have obtained the code in [24] from the authors to compute WCPCap rates.) In particular, under WCPCap, nodes estimate the available capacity in a congested region (using [14]), compute the max-min rate allocation in the region, and send these rates explicitly to the sources.

We use the four topologies discussed in [24] to compare the rates achieved by TCP without CNA, WCPCap, TCP with CNA, and the optimal allocation. These topologies have been used in a number of prior works, see, for example, [24, 30, 29], and are representative of most interference scenarios occurring in mesh network topologies.

Figure 17 shows how far off CNA and WCPCap are from their respective optimal max-min rates by comparing aggregate rates for each topology. Additional simulation results can be found in [13]. Note that although CNA allocates airtime, not rates, we compare rates achieved because in our simulations all radios run at a fixed rate, and we set the channel to be perfect.

For all the four topologies, CNA is *no further away from the optimal than WCPCap*. More surprisingly, CNA enables TCP to achieve throughputs between 5-20% of the optimal rate allocation across the topologies, and is about 12% off the optimal on average across the topologies. Although we made several conservative design choices to preserve stability (Section 3), these have not impacted performance significantly.

While these results are very promising, they do not preclude the existence of topologies where CNA is further away from the optimal. With this in mind, we briefly comment on CNA’s worst case performance. Consider the stack topology again (figure 2(a)). As already discussed, an optimal scheduler would support an airtime limit of $\frac{1}{8}$ for all links, 802.11 would support an airtime limit between $\frac{1}{12}$ and $\frac{1}{8}$, and CNA assigns an airtime limit of $\frac{1}{12}$. The difference comes from the ability of an optimal scheduler to always schedule the two outer flows concurrently, the inability of 802.11 to always do this, and our CNA design choice which conservatively assumes that 802.11 will never be able to do

this. If there are more than two outer flows, the difference among the airtime limits will increase. In general, the worst case scenario that maximizes the difference consists of many edges which interfere with a common edge (the edge in the “middle”) but do not interfere with each other, and the optimal scheduler is the only one which can always schedule concurrently the flows that traverse them. (See [15] for a complete discussion.) It is worth pointing out that a large imbalance in the number of flows per edge exacerbates this problem. If, for example, there are 10 flows in the upper outer branch of the stack and one flow in the lower outer branch, the lower flow suffers because CNA conservatively assumes that 802.11 will never schedule it concurrently with any of the 10 upper flows. Despite the suboptimal performance of CNA in such corner cases, we stand by our design choice on how to allocate airtime limits because it is very simple to use and implement, and it yields near-optimal performance in many real-world scenarios.

7. CONCLUSIONS

We have discussed the design and evaluation of cooperative neighborhood airtime-limiting (CNA), a mechanism that achieves efficient explicit airtime allocation in a manner transparent to TCP/IP and the 802.11 MAC. CNA’s design handles external interference, MAC-layer rate adaptation, and permits mesh TCP connections that also traverse wired links. It achieves performance on average within 12% of the optimal on the topologies we have explored. CNA may be applicable to other wireless technologies including 802.11n, and we have left this to future work.

8. REFERENCES

- [1] Iperf. <http://sourceforge.net/projects/iperf/>.
- [2] MadWifi. <http://madwifi-project.org/>.
- [3] Qualnet. <http://www.scalable-networks.com/>.
- [4] AKYOL, U., ANDREWS, M., GUPTA, P., HOBBY, J., SANIEE, I., AND STOLYAR, A. Joint scheduling and congestion control in mobile ad hoc networks. In *Proc. of IEEE INFOCOM* (2008).
- [5] AZIZ, A., STAROBINSKI, D., THIRAN, P., AND EL FAWAL, A. Ez-flow: removing turbulence in ieee 802.11 wireless mesh networks without message passing. In *Proc. ACM CoNEXT* (2009).
- [6] BAKRE, A., AND BADRINATH, B. I-TCP: indirect TCP for mobile hosts. In *Proc. of IEEE ICDCS* (1995).
- [7] BALAKRISHNAN, H., SESHAN, S., AND KATZ, R. H. Improving reliable transport and handoff performance in cellular wireless networks. *Wireless Networks* (1995).
- [8] FLOYD, S., AND JACOBSON, V. Random Early Detection gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking* (1993).
- [9] FU, Z., LUO, H., ZERFOS, P., LU, S., ZHANG, L., AND GERLA, M. The impact of multihop wireless channel on tcp performance. In *IEEE Transactions on Mobile Computing* (2005).
- [10] GARETTO, M., SALONIDIS, T., AND KNIGHTLY, E. Modeling Per-flow Throughput and Capturing Starvation in CSMA Multi-hop Wireless Networks. In *Proc. of IEEE INFOCOM* (2006).
- [11] HOLLAND, G., AND VAIDYA, N. Analysis of TCP performance over mobile ad hoc networks. In *Proc. of ACM MobiCom* (1999).
- [12] JAIN, K., PADHYE, J., PADMANABHAN, V. N., AND QIU, L. Impact of interference on multi-hop wireless network performance. In *Proc. of ACM MobiCom* (2003).
- [13] JANG, K., GOVINDAN, R., AND PSOUNIS, K. Simple yet efficient transparent airtime allocation in wireless mesh networks. Tech. Rep. 915, University of Southern California, July 2010.
- [14] JINDAL, A., AND PSOUNIS, K. Characterizing the Achievable Rate Region of Wireless Multi-hop Networks with 802.11 Scheduling. *IEEE/ACM Transactions on Networking* (2009).
- [15] JINDAL, A., AND PSOUNIS, K. Making the Case for Random Access Scheduling in Wireless Multi-hop Networks. In *Proc. of IEEE Infocom (Mini-Conference)* (San Diego, CA, March 2010).
- [16] KIM, D., TOH, C.-K., AND CHOI, Y. TCP-BuS: improving TCP performance in wireless ad hoc networks. *IEEE International Conference on Communications* (2000).
- [17] KUMAR, V., M.V. M., PARTHASARATHY, S., AND SRINIVASAN, A. Algorithmic Aspects of Capacity in Wireless Networks. In *Proc. of ACM SIGMETRICS* (2005).
- [18] LI, M., AGRAWAL, D., GANESAN, D., AND VENKATARAMANI, A. Block-switched Networks: A New Paradigm for Wireless Transport. In *Proc. of NSDI* (2009).
- [19] LIU, J., AND SINGH, S. Atp: Tcp for mobile ad hoc networks. *IEEE Journal on Selected Areas in Communications* (2001).
- [20] LOCHERT, C., SCHEUERMANN, B., AND MAUVE, M. A survey on congestion control for mobile ad hoc networks: Research Articles. *Wirel. Commun. Mob. Comput.* (2007).
- [21] MORRIS, R., KOHLER, E., JANNOTTI, J., AND KAASHOEK, M. F. The Click modular router. *SIGOPS Oper. Syst. Rev.* (1999).
- [22] NEELY, M., AND MODIANO, E. Capacity and Delay Tradeoffs for Ad-Hoc Mobile Networks. *IEEE Transactions on Information Theory* (2005).
- [23] RADUNOVIĆ, B., GKANTSIDIS, C., GUNAWARDENA, D., AND KEY, P. Horizon: balancing tcp over multiple paths in wireless mesh network. In *Proc. of ACM MobiCom* (2008).
- [24] RANGWALA, S., JINDAL, A., JANG, K.-Y., PSOUNIS, K., AND GOVINDAN, R. Understanding congestion control in multi-hop wireless mesh networks. In *Proc. of ACM MobiCom* (2008).
- [25] SHANNON, C., MOORE, D., AND CLAFFY, K. C. Beyond folklore: observations on fragmented traffic. *IEEE/ACM Transactions on Networking* 10, 6 (2002).
- [26] SINHA, P., NANDAGOPAL, T., VENKITARAMAN, N., SIVAKUMAR, R., AND BHARGHAVAN, V. WTCP: a reliable transport protocol for wireless wide-area networks. *Wireless Networks* (2002).
- [27] SUNDARESAN, K., ANANTHARAMAN, V., HSIEH, H.-Y., AND SIVAKUMAR, R. ATP: A Reliable Transport Protocol for Ad Hoc Networks. *IEEE Transactions on Mobile Computing* (2005).
- [28] TAN, K., JIANG, F., ZHANG, Q., AND SHEN, X. Congestion Control in Multihop Wireless Networks. *IEEE Transactions on Vehicular Technology* (2006).
- [29] WARRIER, A., JANAKIRAMAN, S., HA, S., AND RHEE, I. DiffQ: Differential Backlog Congestion Control for Multi-hop Wireless Networks. In *Proc. of IEEE INFOCOM* (2009).
- [30] XU, K., GERLA, M., QI, L., AND SHU, Y. Enhancing TCP fairness in ad hoc wireless networks using neighborhood RED. In *Proc. of ACM MobiCom* (2003).