Less-than-Best-Effort Service for Community Wireless Networks: Challenges at Three Layers

Michael Welzl
Department of Informatics
University of Oslo, Norway
Email: michawe@ifi.uio.no

Stein Gjessing
Department of Informatics
University of Oslo, Norway
Email: steing@ifi.uio.no

Naeem Khademi
Department of Informatics
University of Oslo, Norway
Email: naeemk@ifi.uio.no

***Abstract*—Community Wireless Networks can be a way to make "Internet access for everyone" possible, by sharing a broadband Internet connection via WLAN. The underlying idea is to freely provide network access to anybody in the vicinity of the wireless access point via a Lower-than-Best-Effort (LBE) service, such that non-paying users interfere as little as possible with the "regular" Internet usage. Such a service faces challenges at various network layers; this paper discusses some of them, focusing on layers 2, 3 and 4.**

I. Introduction

"Lowest-Cost Denominator Networking (LCD-Net)" was introduced in [1] as a new paradigm that should help make the Internet available for everyone, by ensuring that users and network operators can freely share their resources with anybody without being negatively affected. Community Wireless Networks such as the "Public Access WiFi Service (PAWS)" project in the UK make such sharing possible via wireless Access Points (AP) that have to be installed in private homes which are equipped with a broadband Internet connection – everyone in range of this access point can get free Internet access, and the users who share their resources could receive some form of compensation.

One particular technical requirement is at the core of this idea: the ability to share the resources of subscribers *without being negatively affected*. This requires using the network in such a way that the shared usage interferes as little as possible with any other traffic – this is commonly referred to as providing a "scavenger" or "Less-than-Best-Effort (LBE)" service. Doing this with Community Wireless Networks brings about various challenges at (at least) OSI layers 2, 3 and 4. In the rest of this paper we will discuss these challenges. Section 2, 3 and 4 focus on layer 2, 3 and 4 respectively, and finally in section 5 the paper concludes.

II. Layer 2

If, as in the PAWS project, a dedicated AP is set up to support LBE traffic, the channel for this access point could be chosen such that frequency overlap with other WLANs in the same area is minimized. This by itself could remove any link layer related problems due to the overlapping channels between the LBE traffic and that of subscribers. However, this is not always possible, e.g. in the vicinity of residential or commercial buildings with many private users. Then, one may want to consider letting the LBE traffic share the same access point as the other traffic due to the lack of available spectrum space, which would make it necessary to tune the wireless MAC behavior accordingly.

978-1-4799-4937-3/14/$31.00

Generally, it is not possible to guarantee that a certain way of accessing the wireless medium with Distributed Coordination Function (DCF) mode, i.e. CSMA/CA, would not interfere with other users. For instance, since any station that wants to transmit data applies Listen Before Talk (LBT), any ongoing transmission, LBE or not, will keep that station from data transmission. However, certain parameters could be tuned – e.g. Request-to-Send/Clear-to-Send (RTS/CTS) could be enabled, reducing the chance for frame collision, a large Contention Window (CW) could be used [2] for transmitting LBE traffic to reduce the channel access probability for LBE user when subscriber's device has some data to transmit, or even complete sets of functions such as 802.11e (QoS) or Point Coordination Function (PCF) could be applied. PCF generally takes precedence over DCF, and could therefore be used to serve all but the LBE traffic.

However, there are few challenges in applying such methods. For instance increase of CW can only be applied to the downlink LBE traffic since this can be adjusted at the AP which is in control of subscriber. On the other hand, for the uplink traffic, this is out of subscriber's control. However, such CW increase on the downlink path can still be partially useful for uplink TCP traffic (or any other transport mechanism that uses receiver's feedback) since it can mitigate the dynamism of TCP's control loop by treating the uplink TCP's ACK packets in LBE fashion. The degree of such impact is a subject for further research.

Regarding the use of PCF, while it provides a better centralized control over channel access using polling by the AP in contention-free period, there are practical limitations in place. While PCF is an optional part of 802.11 standard, it hasn't been adopted by Wi-Fi alliance and therefore most Wi-Fi devices don't support PCF. It has also traditionally been argued that PCF is less efficient than DCF in terms of channel access since it requires the wireless nodes to transmit a null frame in response to Contention-Free-Poll (CF-Poll) frames when those nodes have no data to transmit. Apart than this, there has also been other reasons for the lack of PCF deployment in Wi-Fi devices among them the need for a scalable solution (e.g. DCF) that can be used in wireless mesh and ad-hoc networks with various number of active and associated 802.11 nodes positioned in different topological patterns as the polling of all stations in PCF mode incurs overhead when the number of nodes are large and it may also not be possible where a centralized coordination solution is not favored (e.g. in wireless mesh networks).

Similar QoS mechanisms based on PCF such as HCCA in 802.11e share the same fate with PCF in hardware deployment. Applying PCF for all non-LBE traffic requires the change in supporting hardware in all subscriber's wireless devices including the clients and the AP which in turn limits the deployability of such solutions.

However, As explained above some but not all of a WLANs parameters could be tuned at the AP – e.g. the AP can only use a different CW for a certain user on downlink traffic, i.e. traffic that the AP sends. Given that the LBE user's equipment is not under control of the person offering the service, there will therefore be limitations in how much "LBEness" can be achieved in a practical setting, in particular for uplink traffic.

Several modulation and coding schemes which translate into different physical layer bit-rates are provided by the IEEE 802.11 standard [3]. This flexibility can enable robust frame transmission in the presence of a given SINR (Signal to Interference and Noise Ratio). Choosing the right PHY behavior is automatized using a so-called "Rate Adaptation (RA)" mechanism; common examples of such mechanisms include AMRR [4], SampleRate [5] and Minstrel [6] (these schemes are a part of the well-known *madwifi* RA suite). In unfavorable conditions, an RA algorithm typically picks a low rate [7] – but the frame transmission time increases at lower rates, thereby reducing the chances for other nodes to access the channel. This is an inherent fairness problem of WLAN RA: a user who gets the least utility from the network (lowest physical rate) creates the greatest problem for others. If a user is supposed to get an LBE service, then this user should, in principle, send as much as data as it can within the shortest time possible and quickly clear the channel. It is up to the AP to decide which rate to use on downlink traffic, and the AP also advertises the rates that can be used on the uplink; the AP could therefore impose an extra limitation on the rates that LBE users can use and thus avoiding the LBE users experiencing poor channel conditions to starve the subscriber's traffic by using low bit-rates.

In the envisioned LBE scenario, the LBE user is likely to be located in a neighboring building or on the outside, while the AP could be installed in a house or flat. The signal would then have to pass through walls and obstacles, rendering the capability to lower the physical rate in presence of poor channel conditions especially important. When the LBE user is alone, e.g. when the paying user is not at home or asleep, lowering the physical rate does not do any harm. The right solution therefore seems to be not to simply limit the rates that can be used, but to use a RA mechanism that is very cautious at picking a low rate; as shown in Figure 1, there is a significant variety in the behavior of RA mechanisms, even under ideal conditions. Here we define ideal condition as the environment where all nodes stationed within the proximity of the AP and the noise level is minimal and hence the frame loss events are dominated by frame collisions as a result of channel access contention rather than poor channel conditions.

Alternatively, a new RA mechanism could be designed for LBE, to quickly limit the rates to the higher bit-rates when

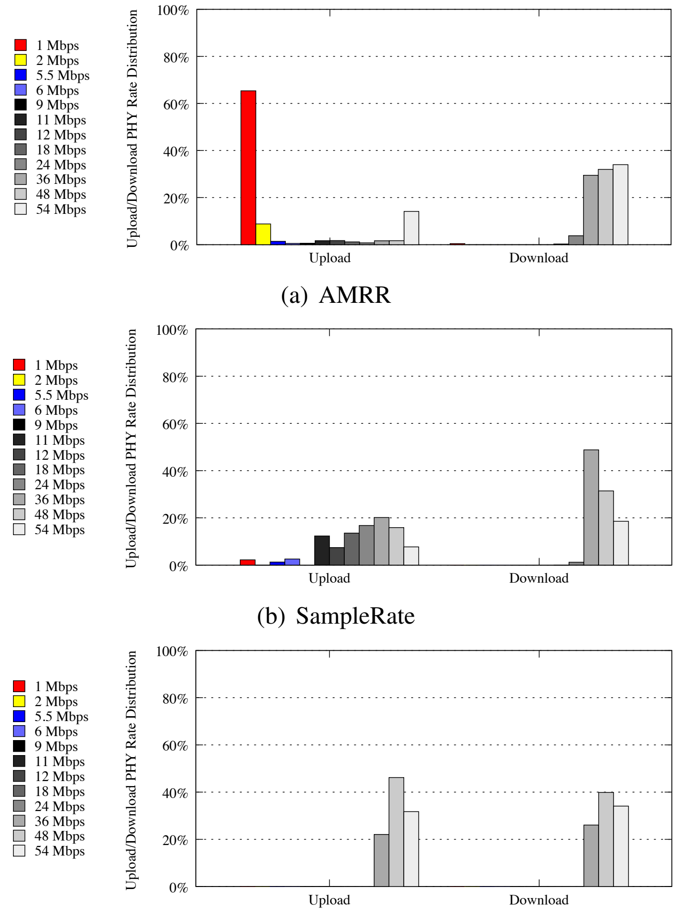

(a) AMRR

(b) SampleRate

(c) Minstrel

Fig. 1. Rate distribution of TCP data packets in a real-life test, using a local testbed under noise-free conditions with different rate adaptation mechanisms. Figure taken from [7].

other users are present (i.e., whenever a collision occurs). At first sight, this may seem like a pretty unusual approach for an RA mechanism, as e.g. AMRR uses the lack of acknowledgment frames (ACKs) – which is normally taken as an indication of a collision – to *lower* the rate. However, the underlying logic in AMRR is that the *consecutive* occurrence of collisions is improbable, and hence missing ACK frames in a certain number of consecutive retries could be taken as an indication of a poor signal instead. Similarly, if an ACK frame is missing only once, and communication works fine immediately afterwards, chances are higher that the missing ACK frame was indeed caused by a collision. Such an *isolated* occurrence of a missing ACK frame could therefore become a reason to increase the physical rate for LBE users. Other than AMRR, Minstrel measures the throughput per bit-rate and uses a table of these measurements to pick the best rates, consecutively lowering the rates according to a "retry chain" in case of failure; the required change for LBE would therefore look quite different for Minstrel (e.g. the retry chain could be limited in some fashion), if it is even necessary at all. For instance a utility function can be taken into account by choosing the bit-rates that deliver a certain amount of data in shortest time possible and if the transmission fails at higher rates terminating the retry chain to avoid the transmission with low bit-rates and hence freeing the channel for the subscriber's devices to access.

As shown in Figure 1, while rates chosen by RA mechanisms can be quite different for the TCP uplink traffic, this is not always the case for TCP downlink traffic. This confirms the earlier finding from [8] that consecutive collisions are especially rare on TCP downlink traffic, making it unlikely for an RA mechanism to erroneously choose a very low physical rate. Since the LBE node is not under control in the assumed Community Wireless Network scenario, the upload behavior cannot be controlled, and hence it might seem unnecessary to try to construct a more suitable RA mechanism. However, when channel conditions are worse, RA mechanisms do choose a lower rate [6], and then the suggested improvement could be applied even on downloads. As for uploads, it could be a possibility to at least recommend that LBE users upgrade their system accordingly, and an AP could try to determine whether the upgrade is in place on the other side of the connection by tracking the rates that were chosen.

III. LAYER 3

At the network layer, the goal of a routing protocol has usually been to find the shortest path from a source to a destination. However, this shortest route may be one with low capacity and it may share links, sub-paths and routers with high priority traffic. The state of the art in the current Internet is that there may be several (totally or partially) disjoint paths between sender and receiver, but only a fraction of these paths are used. Multiple paths are especially common when hosts are multi-homed, which is more and more common, in particular in a Wireless Mesh Network setting that may well be a basis of a Wireless Community Network. However, such multiple paths may have very few disjoint links, because each interface finds its shortest path to the (may be common) destination. MPLS and other circuit switched protocols can of course create paths that are different from the shortest path, but the MPLS paths will be more static than what is needed in an LBE network. There are other ways to utilize more than one path, e.g. Equal cost multipath [9], but here, still the main strategy is to use the shortest paths. Other forms of multipath routing exist, but how the different paths are created and used are open research questions [10].

An LBE network should probably be built as a Software defined network (SDN), where rules for the different types of traffic can be inserted in the routers [11]. In an LBE network one could use an SDN based control plane and program it such that it does not always look for the shortest and fastest path, but rather finds paths that cause no harm to high priority traffic or the most energy efficient or least used paths. Instead of giving weights to links, the routing protocol should use an algorithm that gives weights to routers based on these other parameters. Packets that have very low priority (LBE traffic) may be routed a longer way in the Internet; but as long as the packets do not interfere with regular traffic or use less energy, such paths would be ideal for LBE routing.

However, it is not easy to create and maintain such "non-shortest" paths. An obvious obstacle is that they must not contain loops. This topic has been researched for IP fast reroute, where a loop free alternative is needed in order to route packets around components (switches, routers, links) that have failed. This is being immediately, i.e. when a router sees a failed component on the next hop, it instantly switches traffic to another path that avoids the failed component. If information about a failed network element is transmitted to an ingress router, this ingress router can choose a new path for the data it sends, and in this way avoid the faulty network element. A mechanism that may achieve this uses a set of preconfigured topologies, where a set of routers (and links) that may fail can be pre-deleted from the topology [12]. When forwarding tables are built for these topologies, these routers and links will not be used for forwarding traffic. The IETF standard for multi topology (MT) routing has been identified as a mechanism that can be used for this, although MT routing was originally meant for transporting different kinds of traffic over different paths [13]. MT routing is ideal to route LBE traffic over paths that are not used for regular higher priority traffic, e.g. long delay paths, cheap paths, unreliable paths and energy efficient paths. Using a lightweight tunneling approach, it is possible to always find a route that avoids just one link or one router that is dedicated to high priority traffic, and should not be used at all for LBE traffic [14] [15].

Traffic Engineering (TE) methods are used in order to optimize and configure the use of network resources, e.g. ensuring that urgent data are not routed over long lasting bottlenecks in the network, and if it is, alleviate the bottleneck by routing traffic around them. With SDN many of the TE tasks are integrated in the software defined control plane so that forwarding tables are set according to the wishes of the network owner. Since the time aspect in TE traditionally is long (hours, days, weeks), parameters like traffic load (the traffic matrix), QoS and resilience (e.g. creating backup paths) are considered. Traffic engineering is often classified as on-line or off-line [16]. Usually off-line TE is performed with manual help, while online TE means that automated methods are in use. With SDN one might say that online TE is now much easier to perform as an integral part of the regular control plane.

Traditionally there has been a sharp distinction between routing and Traffic Engineering. Within a network that also is to accommodate LBE traffic, the control plane must be programmed to undertake tasks that have previously been considered TE tasks. The network owner must get the ability to implement the routing protocol of choice, i.e. the owner must be able to build network layer components that offer exactly the services that the high priority customer needs, and at the same time be able to accommodate LBE traffic when resources are available. The routers do not need to have only one forwarding table, or use only one forwarding rule [13]. Forwarding can be based on any set of parameters that the network owner wants. Using the shortest path might be a choice taken only for small high priority packets. For less urgent traffic, a longer path with more delay might be the better choice.

Small packets may incur overhead and extra energy usage in a network. A router in an area with poor access to power – which is not unlikely in a LCD-Net-like setting, e.g. in economically challenged parts of the world – may only be operational for a few minutes or a few hours or a day. Then packets destined for the same egress should be aggregated in the ingress and when power and a path become available, sent immediately as one large burst. Forwarding of such aggregated burst could also be performed as tunnels from ingress to egress through an Aggregation Domain (AD). This could also simplify the forwarding tables within the network: all the information needed is the next hop towards the egress. A burst of packets (an aggregate) is sent from one ingress edge router to an egress edge router in the same AD when the interior network within the domain is connected and has available capacity for LBE traffic. An aggregate in an egress router is transmitted to the ingress router of the next AD when this link is up and there is room for LBE traffic. The architecture may be such that sometimes the egress router of one domain is the same as the ingress router of the next domain. For example, in the scenario shown in Figure. 2, E3 and E5 may be one router.

In this example (Figure 2), several (three) messages are sent to E1, and they are all to be routed to E9. When the interior network from E1 to E3 are available for LBE traffic, the aggregate is tunneled to E3. The routing protocol is such that the chosen path may be either via C2 or via C3, depending of which routers and links have most spare capacity and incur least overhead and cost. In the same way the aggregate is tunneled further via E5, E6 and E7, until it ends up in E9. Here the aggregate is unpacked and three messages are sent on to their final destinations.

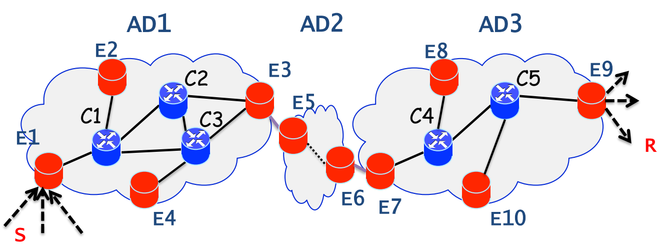

Fig. 2. The figure shows three aggregation domains (ADi). The regular interior routers (Ci) do not have aggregate or caching properties. Packets may only be aggregated in the edge routers (Ei).

IV. Layer 4

Transport connections interfere with each other when their traffic ends up in the same bottleneck queue, via the delay or loss that this queue produces. The most immediate method to achieve LBE behavior is therefore to differentiate between the packet types and accordingly treat them in a router that contains a congested queue, e.g. using packet scheduling mechanisms or Active Queue Management (AQM) functions [17], or inform the senders to back off immediately upon queue growth using ECN [18].

While being more than two decades old in context with the introduction of Random Early Detection (RED) [19], new AQM schemes such as CoDel [20] and PIE [21] have recently been proposed with the aim to keep the latency on the Internet low, and hence tackling the commonly known problem of bufferbloat on the Internet [22]. Having that in mind, AQM schemes can be customized to treat LBE traffic differently than normal traffic generated by the subscriber e.g. using a different set of marking/dropping thresholds for LBE traffic and/or with different priority queues. Packets scheduling mechanisms such as weighted fair queuing (WFQ) and weighted RED (WRED) [23] can also be combined with AQM mechanisms to fulfill this goal by giving precedence to the subscriber's traffic over LBE traffic.

In the envisioned LCD-Net LBE scenario, it is desirable for the LBE traffic to "step out of the way" of other traffic at any queue in the Internet (e.g., if both the LBE user and the paying user download a file from the same server and that server's access link is the bottleneck). While it might be possible to opportunistically mark packets for LBE, akin to what is currently being proposed for WebRTC [24], one cannot rely on routers to support such a marking. Hence, LCD-Net LBE requires a different approach.

When a queue begins to grow, delay is the most immediate effect that is typically seen by end systems. Hence, if an end-to-end congestion control mechanism immediately reduces its rate when delay increases, it can achieve an LBE behavior, as most other senders (at least standard TCP and most experimental TCP variants) keep increasing their rates until a packet is lost. A well-known mechanism that does this is LEDBAT [25], a variant of which is deployed in BitTorrent. LEDBAT illustrates the difficulties involved in developing such a mechanism correctly: it requires to build a stable queue at a certain *target delay*, which in itself is a trade-off that requires careful tuning [26]. LEDBAT has been shown to be prone to various sorts of problems, among them the well-known "late-comer advantage" [27]. Figure 3 shows another problem (described in detail in [28]): because it uses a minimum ("base") delay in its calculation and repeatedly measures this minimum delay while continuing to send traffic, it considers its self-caused delay as the new minimum, making the overall delay of LEDBAT grow larger and larger when it is used long enough.

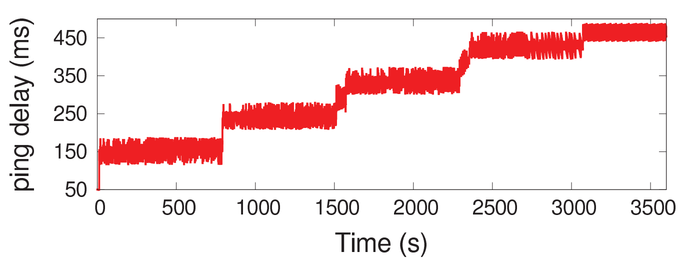

Fig. 3. Ping delay across an Ethernet link that was set to 500 kbit/s / 50 ms with netem, with the libutp utp test program (BitTorrent's LEDBAT implementation) running. Figure taken from [28].

LEDBAT is by no means the only end-to-end approach for LBE at the transport layer; a survey is given in [29]. Not all LBE mechanisms might have the same problems as LEDBAT – e.g. Delay-Gradient TCP [30] only considers the gradient of the delay signal and hence does not produce a standing queue. On the other hand, Delay-Gradient TCP bases its decision on the Round-Trip Time (RTT), which can be affected by noise on the return link. To get a more precise idea of the state of the forward queue, LEDBAT uses changes in One-Way Delay (OWD). This requires the sender to timestamp packets, giving the receiver the necessary OWD information, which it would then have to feed back to the sender where the rate is calculated. Hence, other than Delay-Gradient TCP which is available in FreeBSD since version 9.0 as a pluggable congestion control mechanism for TCP, LEDBAT can probably not be correctly implemented as a one-sided change to TCP.

Such deployment considerations are important for the LCD-Net LBE scenario because the LBE user's stack is not under control of the system. Here, we basically have two choices for realizing a LBE service at the transport layer:

1) Split TCP connections [31] at the LBE AP such that the AP acts like a receiver towards the remote Internet node and a sender towards the LBE node for downstream traffic, and like a receiver towards the LBE node and a sender towards the remote Internet node for upstream traffic. Then, whenever the AP acts like a sender, sender-based LBE schemes could be used, and whenever it acts like a receiver, receiver-based LBE schemes could be used.
2) Do not split TCP connections but monitor their delay and change the TCP Receiver Window (rwnd) in packets when needed to force the sender to slow down. The Receiver Window is used by several of the receiver-based schemes in [29].

This is only considering "normal" TCP connections between the remote Internet node and the LBE node, not UDP or other relatively common transports such as Multipath TCP [32] or SCTP [33]. As with TCP, it should also be possible to intercept such other traffic at the LBE router and affect it accordingly.

Note that it may be too cautious to *always* impose LBE-oriented congestion control onto the LBE user: it is really only necessary when the paying user is active, and even then, it is actually only necessary when the two users share the same bottleneck queue. Shared bottlenecks can be detected using active or passive methods (cf. [34] and references therein), but such schemes are not yet in widespread use, as they have historically been regarded as either too hard to use (too computationally intensive) or not reliable enough. This is perhaps going to change now, as the RTP Media Congestion Avoidance Techniques (RMCAT) IETF group is planning on standardizing a shared bottleneck detection method for the sake of WebRTC communication.

V. Conclusions

The LCD-Net paradigm – making LBE-style Internet access possible for everyone – has a noble goal. However, in practice, it seems that very little well-functioning LBE technology is in place in today's networks. In particular, having *no* interference whatsoever with the paying customer's traffic may be hard to achieve; instead, it may be worth trying to quantify the degree of such interference. As we have seen in this paper, problems exist at several layers – quite possibly also the physical and application layer in addition to the three layers that we have focused on. These challenges are not necessarily obstacles that cannot be overcome, but it seems clear that a significant amount of research is needed before all the right LBE mechanisms can be put in place.

References

[1] A. Sathiaseelan and J. Crowcroft, "LCD-Net: Lowest Cost Denominator Networking," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 2, pp. 52–57, Apr. 2013. [Online]. Available: http://doi.acm.org/10.1145/2479957.2479966

[2] A. Khalaj, N. Yazdani, and M. Rahgozar, "Effect of the contention window size on performance and fairness of the IEEE 802.11 standard," *Wireless Personal Communications*, vol. 43, no. 4, pp. 1267–1278, 2007. [Online]. Available: http://dx.doi.org/10.1007/s11277-007-9300-5

[3] IEEE, "Standard for Information Technology- Telecommunications and Information Exchange Between Systems-Local and Metropolitan Area Networks-Specific Requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE Std 802.11-1997*, pp. i –445, 2007.

[4] M. Lacage, M. H. Manshaei, and T. Turletti, "IEEE 802.11 Rate Adaptation: A Practical Approach," in *Proceedings of the 7th ACM International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM '04. New York, NY, USA: ACM, 2004, pp. 126–134. [Online]. Available: http://doi.acm.org/10.1145/1023663.1023687

[5] R. T. Morris, J. C. Bicket, and J. C. Bicket, "Bit-rate selection in wireless networks," Masters thesis, MIT, Tech. Rep., 2005.

[6] D. Xia, J. Hart, and Q. Fu, "Evaluation of the Minstrel rate adaptation algorithm in IEEE 802.11g WLANs," in *Communications (ICC), 2013 IEEE International Conference on*, June 2013, pp. 2223–2228.

[7] N. Khademi, M. Welzl, and S. Gjessing, "Experimental evaluation of TCP performance in multi-rate 802.11 WLANs," in *13th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM) (IEEE WoWMoM 2012)*, San Francisco, California, USA, Jun. 2012.

[8] S. Choi, K. Park, and C.-k. Kim, "On the Performance Characteristics of WLANs: Revisited," in *Proceedings of the 2005 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '05. New York, NY, USA: ACM, 2005, pp. 97–108. [Online]. Available: http://doi.acm.org/10.1145/1064212.1064225

[9] D. Thaler and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection," RFC 2991 (Informational), Internet Engineering Task Force, Nov. 2000. [Online]. Available: http://www.ietf.org/rfc/rfc2991.txt

[10] A. Kvalbein, C. Dovrolis, and C. Muthu, "Multipath load-adaptive routing: putting the emphasis on robustness and simplicity," in *17th IEEE International Conference on Network Protocols, ICNP 2009*, 2009, pp. 203–2012.

[11] A. Sathiaseelan, C. Rotsos, C. S. Sriram, D. Trossen, P. Papadimitriou, and J. Crowcroft, "Virtual Public Networks," *2013 Second European Workshop on Software Defined Networks*, vol. 0, pp. 1–6, 2013.

[12] A. Kvalbein, A. Hansen, T. Čičić, S. Gjessing, and O. Lysne, "Multiple Routing Configurations for Fast IP Network Recovery," *ToN*, vol. 17, no. 2, pp. 476–486, 2009.

[13] P. Psenak, S. Mirtorabi, A. Roy, L. Nguyen, and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF," RFC 4915 (Proposed Standard), Internet Engineering Task Force, Jun. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4915.txt

[14] S. Bryant, S. Previdi, and M. Shand, "A Framework for IP and MPLS Fast Reroute Using Not-Via Addresses," RFC 6981 (Informational), Internet Engineering Task Force, Aug. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6981.txt

[15] M. Menth, M. Hartmann, R. Martin, T. Čičić, and A. Kvalbein, "Loop-free Alternates and Not-via Addresses: A Proper Combination for IP Fast Reroute?" *Comput. Netw.*, vol. 54, no. 8, pp. 1300–1315, Jun. 2010. [Online]. Available: http://dx.doi.org/10.1016/j.comnet.2009.10.020

[16] N. Wang, K. Ho, P. G., and H. M., "An overview of routing optimization for Internet traffic engineering," *IEEE Communications Surveys & Tutorials*, vol. 10, no. 1, pp. 36–56, 2008.

[17] R. Bless, K. Nichols, and K. Wehrle, "A Lower Effort Per-Domain Behavior (PDB) for Differentiated Services," RFC 3662 (Informational), Internet Engineering Task Force, Dec. 2003. [Online]. Available: http://www.ietf.org/rfc/rfc3662.txt

[18] M. Arumaithurai, X. Fu, and K. Ramakrishnan, "NF-TCP: Network Friendly TCP," in *Local and Metropolitan Area Networks (LANMAN), 2010 17th IEEE Workshop on*, May 2010, pp. 1–6.

[19] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Trans. Netw.*, vol. 1, no. 4, pp.

397–413, Aug. 1993. [Online]. Available: http://dx.doi.org/10.1109/90.251892

[20] K. Nichols and V. Jacobson, "Controlling Queue Delay," *Queue*, vol. 10, no. 5, pp. 20:20–20:34, May 2012.

[21] R. Pang, "PIE: A Lightweight Control Scheme to Address the Bufferbloat Problem," Working Draft, Internet-Draft draft-pan-tsvwg-pie.txt, Jun. 2013.

[22] Bufferbloat. http://www.bufferbloat.net/.

[23] Cisco WRED Guide. http://www.cisco.com/en/US/docs/ios/qos/configuration/guide/config_wred.pdf.

[24] S. Dhesikan, D. D. (Ed.), P. Jones, and J. Polk, "DSCP and other packet markings for RTCWeb QoS," Internet Draft draft-dhesikan-tsvwg-rtcweb-qos-04, work in progress, Jan. 2014.

[25] S. Shalunov, G. Hazel, J. Iyengar, and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)," RFC 6817 (Experimental), Internet Engineering Task Force, Dec. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6817.txt

[26] N. Kuhn, O. Mehani, A. Sathiaseelan, and E. Lochin, "Less-than-Best-Effort Capacity Sharing over High BDP Networks with LEDBAT," in *Vehicular Technology Conference (VTC Fall), 2013 IEEE 78th*, Sept 2013, pp. 1–5.

[27] G. Carofiglio, L. Muscariello, D. Rossi, and S. Valenti, "The quest for LEDBAT fairness," in *Proc. IEEE GLOBECOM*, Miami, Dec. 2010.

[28] D. Ros and M. Welzl, "Assessing LEDBAT's Delay Impact," *Communications Letters, IEEE*, vol. 17, no. 5, pp. 1044–1047, 2013.

[29] ——, "Less-than-Best-Effort Service: A Survey of End-to-End Approaches," *Communications Surveys Tutorials, IEEE*, vol. 15, no. 2, pp. 898–908, 2013.

[30] G. Armitage and N. Khademi, "Using Delay-Gradient TCP for Multimedia-Friendly 'Background' Transport in Home Networks," in *Proceedings of the 38th IEEE Conference on Local Computer Networks (LCN 2013)*, October 2013.

[31] J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations," RFC 3135 (Informational), Internet Engineering Task Force, Jun. 2001. [Online]. Available: http://www.ietf.org/rfc/rfc3135.txt

[32] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses," RFC 6824 (Experimental), Internet Engineering Task Force, Jan. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6824.txt

[33] R. Stewart, "Stream Control Transmission Protocol," RFC 4960 (Proposed Standard), Internet Engineering Task Force, Sep. 2007, updated by RFCs 6096, 6335, 7053. [Online]. Available: http://www.ietf.org/rfc/rfc4960.txt

[34] M. M. Yousaf and M. Welzl, "Less-than-Best-Effort Service: A Survey of End-to-End Approaches," *The Scientific World Journal*, vol. 2013, no. Article ID 890578, 2013.