How is your event Wi-Fi doing? Performance measurements of large-scale and dense IEEE 802.11n/ac networks

Patrick Bosch*, Jeroen Wyffels†, Bart Braem* and Steven Latré*
*University of Antwerp - imec
Middelheimlaan 1, B-2020 Antwerp
firstname.lastname@uantwerpen.be
†Telenet
Liersesteenweg 4, B-2800 Mechelen
jeroen.wyffels@staff.telenet.be

***Abstract*—The popularity of the IEEE 802.11 (Wi-Fi) standard has resulted in a plethora of hotspot deployments. While home hotspots offer high and consistent performance, at large-scale events such as conferences and festivals, Wi-Fi performance is often poor and highly fluctuating. There are several factors explaining this increased difficulty: the required scale, the complexity of backhaul network topologies, the density of devices connected to the access point (AP), and the interference caused by both people and radio frequency (RF) equipment. While these factors are all known to degrade the performance of public hotspots, little is known about the actual performance of IEEE 802.11 at large-scale events. In this paper, we present the results of quantitative Wi-Fi performance measurement study undertaken at a music festival with 80,000 visitors over a geographical area of 0.3 square kilometres. Two separate networks were constructed for this study. The first was an IEEE 802.11n-based wireless mesh consisting of 37 devices in 15 nodes working as a network backhaul and the second an IEEE 802.11n/ac-based public hotspot that was accessible by the festival goers. We characterise the performance of the wireless spectrum and illustrate the impact of interference factors such as crowds and RF equipment. Finally, we report on the results of the deployment of a public hotspot at the festival, focusing more on application-layer metrics parameters such as user experience and session statistics. The results show that the interference at such events is so high that adaptations to the protocol configuration are needed to improve performance.**

I. Introduction

With the everyday use of wireless devices such as smartphones and tablets, users expect to be always connected to the Internet. This leads to an increased deployment of Wi-Fi APs, which in turn results in a higher density of wireless devices. Dense wireless networks (DenseNets) will be more common and bring with them the challenge of huge interference. One particular and large-scale example of this is Wi-Fi at events. Wi-Fi is commonly deployed for two reasons at large-scale events: as a public hotspot for visitors or as a wireless backhaul network that provides connectivity to harder to reach regions. The expectation of the quality of wireless broadband connectivity is the same as the one that people have at home. To deliver that Quality of Service (QoS) in DenseNets is still an open research challenge. This makes it increasingly difficult to deploy hotspots at events as there are three main challenges.

First, wireless devices interfere with each other when they are competing for air-time. The higher the density of the devices, the higher the interference will be, resulting in significant performance degradation. Not only active devices, but also inactive devices can have an impact, as there are thousands of them probing or keeping the connection alive [1]. This is especially the case for a very dense environment, as the devices are in close proximity to each other.

A second challenge is that, the backhaul networks themselves are evolving towards wireless networks. They are easier to deploy, as organisers do not need the installation of cables towards every network device. But the networks need to be reliable, have a low latency and sufficient throughput.

The third challenge is the large presence of interference by RF equipment. This is especially relevant for events where RF equipment such as audio, video and the control of fireworks lead to interference in the wireless spectrum. This is a completely new challenge, as usually this is not considered during the design or implementation phase of wireless deployments and hardware is not designed to cope with it.

In this paper, we investigate the performance of Wi-Fi in large-scale and dense environments, as a public network and a wireless backhaul. For this purpose, we performed a large-scale field trial for dense Wi-Fi hotspots. This was composed of a backhaul network and a public hotspot at a large festival hosting 80,000 visitors over an area of 0.3 square kilometres. To the best of our knowledge, this is the first paper that presents a comprehensive study of two deployments in such a particular large-scale dense environment. The contributions of the paper are extensive measurements from low level measurements up to the application layer of a field trial.

The remainder of the paper is structured as follows. We discuss previous work in Section II and continue with the setup overview in Section III. We follow up with an explanation of the extensive results in Section IV and a guideline for large deployments Section V. Finally, in Section VI we conclude.

978-3-901882-89-0 @2017 IFIP

II. Related work

In this section, we provide an overview of the state of the art in the performance characterisation of large-scale deployments of wireless mesh networks and public Wi-Fi hotspots.

A. Large-scale wireless mesh networks

While the Wi-Fi infrastructure mode is still dominant, Wi-Fi-based mesh setups are also used in production networks with latency sensitive applications. In this setup, layer 2 or layer 3 routing protocols such as AODV, B.A.T.M.A.N. or B.A.T.M.A.N. advanced are being used [2]. As such, ample examples exist of such deployments (e.g., in office environments to handle backhaul traffic [3], as a city-wide community Wi-Fi [4]). Vural et al. provide a survey of experimental evaluations of such wireless city-based mesh networks [5]. The impact of external interference effects is there identified as one of the most important challenges in setting up a wireless mesh network. Moreover, Vural et al. present guidelines in terms of node location, use of directional antennas, etc. Many of these guidelines were taken into account in the setup of our large-scale wireless mesh. Our work therefore differs from the above studies in three ways. First, we focus on much denser deployments with up to 80,000 people on a small geographic area. Second, our environments are much more challenged from interference due to a high amount of RF equipment. Third, our environment is more dynamic: static deployment guidelines (e.g., in terms of node location) are often not possible due to the short time frame in which these networks are set up and the lack of power in several locations.

The study of very high density deployments (e.g., those featured at a festival) have mainly been limited to analytical and/or simulation models. With this respect, Michaloliakos et al. provide a model for characterising the performance through simulation of a medium density deployment such as a conference [6]. Abinader et al. show decreased throughput in simulations for high density environments [7].

These simulations show that (i) the density introduces important performance costs and (ii) that the performance is reverse proportional to the hop count in the wireless mesh. However, it does not accurately reflect external interference effects which is crucial in a real deployment. To the best of our knowledge, our work is the first that characterises the performance of a deployed wireless mesh backhaul in such a high-density and dynamic setting.

B. Public Wi-Fi hotspots

With Wi-Fi being one of the most important consumer oriented wireless technologies, the Wi-Fi infrastructure mode is more and more used for large-scale deployments as well, offering public Internet connectivity to the users. Characterising the performance of these deployments is key in understanding their challenges. Again, most studies focus on analytical models to estimate the load and corresponding performance in large-scale hotspots. Ghosh et al. present a model to estimate traffic in large-scale deployments [8]. Collected data is used to create a model, which is able to accurately model the traffic and session distribution. Paul and Ogunfunmi's work is more focused on link performance [9]. Their focus is on a complete analytical model for the IEEE 802.11n standard to model the behaviour for different parameters. While Zhang et al.'s main focus is a signal to noise ratio (SNR) based rate adaption, they provide a good overview over the impact of SNR and interference on the packet delivery ratio (PDR) [10]. Focusing on the packet error rate (PER), Ramachandran et al. present measurements for saturated and non-saturated channels [11]. Their setup is relatively dense and it gives a good overview over the impact of density in wireless networks. Gummadi et al.s work focuses on interference outside of IEEE 802.11, while Rayanchu et al. present a system to detect it with commodity Wi-Fi hardware [12] [13]. They show significant vulnerability regarding latency and throughput with only a very low power output on the interfering device. In contrast to the above described analytical models, we focus on an experimental characterisation of such public hotspots, relying on an actual deployment.

Experimental performance characterisations of public hotspot deployments are rare. McHenry et al.. present spectrum analysis measurements for Chicago [14]. Their Wi-Fi measurements show that the band is already well used with received signal strength indicator (RSSI) up to around -65dB. Their work is however limited to a medium-sized hotspot and focuses only on lower layer measurements. Biswas et al. present a large number of network measurements from a multitude of deployments, which includes data about channel utilisation, delivery ratio and spectrum analysis, but also usage of operating systems and devices [15]. This is one of the first large-scale experimental performance studies that investigate a real deployment. Our work complements this approach in the sense that we focus on the performance of such hotspots in large-scale events with thousands of potential users, while Biswas et al. rather focus on the performance of medium-sized hotspots targeting hundreds of potential users.

III. Experimental setup overview

In this section we will explain the deployed setup consisting of both a wireless backhaul and public accessible hotspots.

A. QoS performance of a Wi-Fi backhaul network

Each event requires a backhaul communication infrastructure to allow connecting different locations across the event area. Each location requires communication for subsystems such as a cash register, personnel management, etc. Currently, events deploy kilometres of wired connections to form this backhaul infrastructure. This is obviously both a costly, time consuming and non-trivial task. We evaluated the feasibility of a Wi-Fi-based approach to replace the typical wired backhaul. To do this, we deployed a Wi-Fi-based mesh network of 15 nodes across the festival area, including six stages with, depending on the location, hundreds to several thousand people attending. We chose directional antennas to improve performance, as our links were point-to-point and the maximum distance was close to one kilometre. We used the Ubiquiti

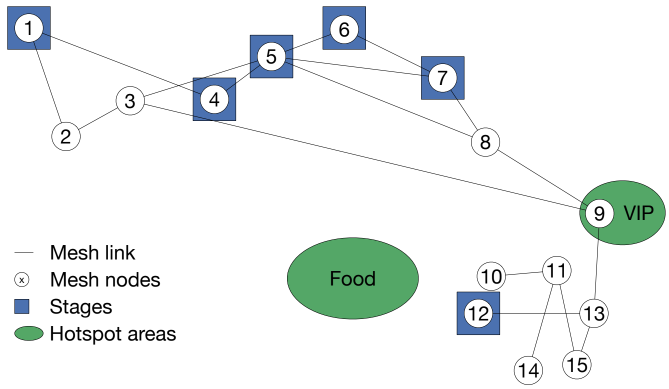

Fig. 1. Initial setup for the Wi-Fi mesh, showing all the nodes and the areas for the hotspot deployment.

NanoStation M5 [16]. Each device ran OpenWrt and used the IEEE 802.11n standard [17]. To be able to perform mesh-based routing, the B.A.T.M.A.N. advanced protocol was used [18].

The topology of the deployed Wi-Fi mesh network is illustrated in Figure 1. The nodes themselves were placed in areas with stages close to them, or other places that attract a large number of people like bars or food stalls. This means that they were close to other electronic equipment (which were often also transmitters of - possibly interfering - RF signals), but also to the stages, bars and food stalls. The height of the nodes did vary, as some could be placed very high, up to 20 metres high, and others were close to ground level. Most of the nodes we deployed were on ground level at about one to three metres height. This means that for most links, both people and user devices such as smartphones potentially cause interference as they were moving between the nodes.

B. Application-level performance for a public hotspot

Except for replacing the wired backhaul network, Wi-Fi is obviously also used to provide public Internet access to visitors at the event. In this section, we explain the setup of the publicly accessible hotspot for all festival visitors. To this end, we deployed 54 APs, spread across three different zones: the VIP and food area in Figure 1 on the festival area and one on the camping area with limited RF interference.

Frequency-wise, both 2.4GHz and 5GHz were used. A maximum of 200 users could be served by each AP, leading to a theoretical maximum of 10,800 simultaneous users. In order to simultaneously serve that many clients, a traffic shaping profile was used on all Wi-Fi sessions. For 2.4GHz, IEEE 802.11n was supported, and for 5GHz, IEEE 802.11ac was enabled. Channel bonding was disabled. Legacy IEEE 802.11b rates were disabled as well to improve usability and reliability. Directional antennas were used and in the camping area, only a part was covered by the deployment. The hotspots were not publicly announced on every AP, but only at few places.

IV. Performance results description

In this section, we will present the results of the different deployments.

A. QoS performance of a Wi-Fi backhaul network

In this section, we discuss the QoS performance of the Wi-Fi mesh throughout the field trial. We define QoS using the following metrics:

- The RSSI
- Availability of a link describing the percentage how often that link was seen during our measurements.
- Usage of a link describing the percentage how often that link was actively used in routing.
- Latency of ping messages periodically send every second
- Loss ratio of pings

In Figure 2 we can see the availability of links (a) and the usage of links (b) by the B.A.T.M.A.N. advanced protocol as a fraction of the 3-day time zone in which the nodes were online. When compared with the original topology assumed at design phase and illustrated in Figure 1, we can see a lot more available links than planned (Figure 2a). Especially in the denser clustered area with nodes 10 to 15, most of the nodes can see each other during the entire festival. Although directional antennas are by nature focused in one direction, their angle of radiation usually is wide enough to allow for not completely aligned links. While the availability for our chosen links is rather high, we have a high number of links that have an availability from 20% to 70%. But the availability itself is not sufficient to explain the quality of the link. If we have a look at the usage of the links and compare it to the availability, we can see links with a high availability, for example between node 3 and node 8, but a low usage (Figure 2b). The link was simply not reliable enough to be chosen by the B.A.T.M.A.N. advanced protocol, although it was available most of the time. The overall medium usage of specific links, indicates frequent changes in the topology. Other links were chosen because the quality of the link decreased. This unreliability is mainly due to the fact of too much interference and therefor too much packet loss on links.

In the 15 nodes deployed in the wireless backhaul, we could distinguish three different types, depending on where they were placed (i) at the edge, (ii) indoor or (iii) outdoor. Edge nodes are nodes that located at the edge of the festival ground and had therefore less interference and less visitors around them. With outdoor, we denote nodes that were positioned outside, mainly on top of stages or other high places and therefore had sufficient line of sight with other nodes and did not suffer from the Farady effect. Indoor reflects the nodes that are either inside stages or other buildings. This means that these do not have line of sight, but send through structures. Moreover, as a stage is typically a large metal construct it acts as a Faraday shield.

Figure 3 displays the average RSSI for the three node types (Edge, Outdoor and Indoor) with the standard deviation. There is a clear distinction in the average values, while the standard deviation is very similar for all groups at around 10.5dBm. The nodes of type edge have an average of around -66dBm, which is acceptable. Home setups can achieve better values, but the performance of Wi-Fi is still fine with -66dBm. For the nodes

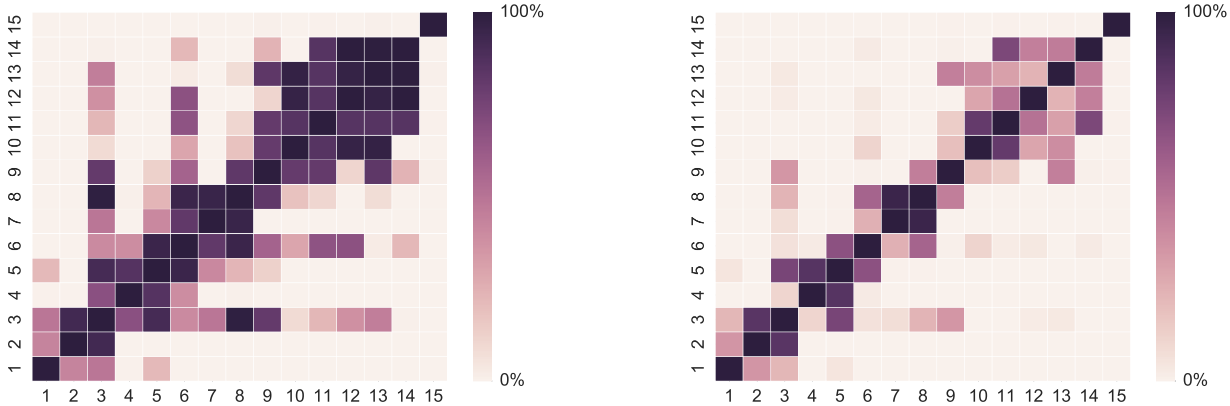

(a) Available links.

(b) Used link.

Fig. 2. A lot more links are available than were planned and link usage is not consistent over time.

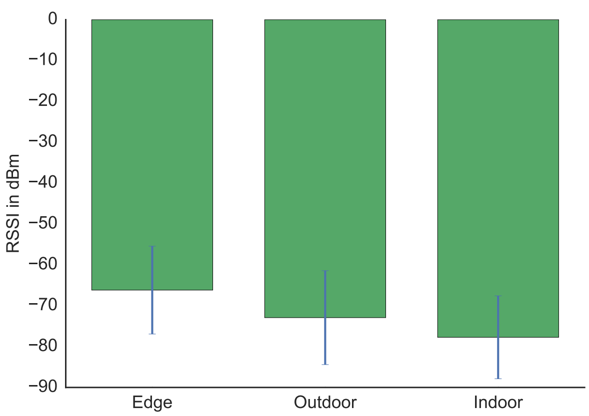

Fig. 3. The average RSSI for the three node types. It is important to place nodes correctly to avoid both external interference and Faraday effects.

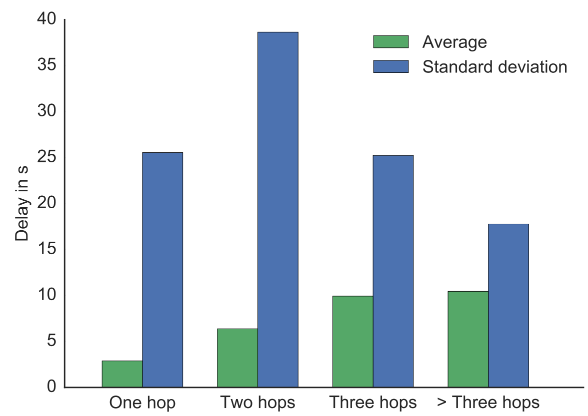

Fig. 4. The average latency for the ping tests with standard deviation show that latency sensitive communication is not feasible.

of type outdoor it already decreases about 7dBm compared to the nodes of type edge. While -73dBm is still a value where Wi-Fi performance is sufficient, the large standard deviation of 10.5dBm also clearly shows the significant instability of the RSSI values, potentially leading to frequent losses as well. The nodes of type indoor are worse again with an additional 5dBm of loss, resulting in an average of -78dBm. The again high standard deviation of 10.5dBm shows two things. On the one hand, this means that a large number of packets can not be received correctly and need to be retransmitted, on the other hand the high deviation means that there was no continuous interference, but an ever changing one.

A major QoS factor is latency, in Figure 4 we can see the average latency, for an increasing hop count in the wireless mesh. The results clearly show the deteriorating effect of the external interference on the overall latency. For only a single hop, we already see an average latency of 2.5 seconds. This so high that any type of bidirectional communication is impossible. For larger hop counts this increases up to 10 seconds and more. On the other hand, the standard deviation is extremely large as well. For one hop, around 50% of the pings were below 100ms, which is still acceptable for most applications. This percentage decreased down to 10% for more than two hops and was for two hops around 30%.

This large interference does not only have an impact on latency. Figure 5 illustrates the observed application loss rate as a function of the mesh hop count. Even to direct neighbours the loss rate is already above 50%. As soon as there are two hops, the loss rate gets close to 80%. With more than two hops the rate is even above 90%.

These abnormal results are mainly due to the large amount of RF interference that can be observed in the specific environment of a large-scale event. IEEE 802.11 uses a carrier sense multiple access with collision avoidance (CSMA/CA) scheme to control channel access. In CSMA/CA, the backoff timer associated to this window is increased every time interference

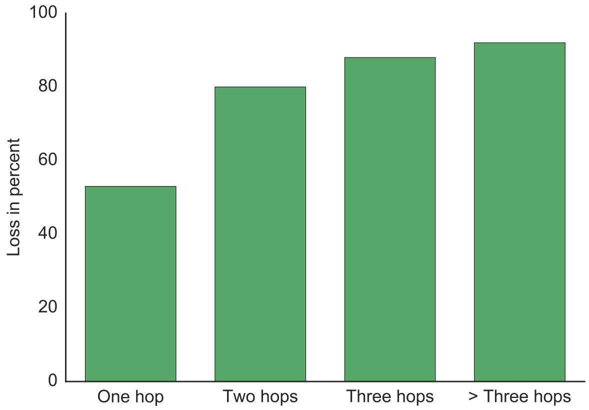

Fig. 5. The loss is above 40% for one hop connections and rises to nearly 100% for multi hop connections.

is detected. This means that the latency can easily increase in high interfering environments.

This latency is partly modelled by Chatzimisios et al. [19], which provides an analytical model for the expected latency in an IEEE 802.11b network. Their study mainly shows the increasing latency when the number of stations increases. They however do this for an ideal channel, which was clearly not the case for our deployment. But nonetheless, the model estimates up to 1 second of latency if there are 70 stations present. This can explain around 50% of the latency to direct neighbours. For neighbours further away, it is around the same number, as we need to add 1 second for each hop. As even probe requests can create significant interference [1], a high latency is less surprising if there are thousands of client devices present. The other half of the latency is harder to explain, here we need to take into account that the model from Chatzimisios et al. is only for ideal channels. The time the channel is sensed busy will be much higher than the assumption in the model, which would increase the latency significantly. This is due to the fact that not only our devices are present, but also a lot of external Wi-Fi (e.g., devices on other networks) and non-Wi-Fi interfering (e.g., other RF equipment) sources.

B. Application-level performance for a public hotspot

While the previous section discussed the performance of the Wi-Fi backhaul, we now discuss application-level performance of the publicly accessible hotspot during the festival. In this section, we are evaluating the results in regard of sessions per AP, traffic per AP and traffic to session distribution. A session is the time a client was associated to an AP. During the festival, there were about 1,900 sessions on average with a peak of up to 19,000 sessions and a total of around 1,000,000 over the period of three days. The number of uniquely connected clients was on average 700, with a peak of 2,800 and a total of around 22,000. Figure 6 shows the distribution of the sessions over the APs in the areas where Wi-Fi was provided with the average and the 95th percentile. Especially striking for all APs is the large deviation in terms of connected users across them. This means, those had at some point a lot more users connected to them, as the APs around them. This brings up an important point: although the APs were geographically fairly balanced between the audience, the resulting AP association is not.

We can see a similar pattern when we look at Figure 7, where the traffic for the APs is depicted. In this case, we only focus on the traffic in the food area, but the other areas have similar results. As can be seen, the unbalance in session per AP also results in an unbalance in traffic per AP. These results clearly show the need for more advanced load balancing mechanisms, which are able to evenly distribute the traffic amongst APs. In the IEEE 802.11 standard, the decision to associate to a particular AP is done by the end device. However, as can be seen from the above results, for high density environments, this is far from optimal.

Figure 8 illustrates the number of sessions compared to the average traffic per session in the food area. We can see that few sessions generate high-bandwidth traffic. On the other hand, nearly half of the sessions are around a few kilobytes of traffic. This means, that a lot of the sessions were only used for automatic services like fetching emails or similar things. Manual services like browsing or even streaming were only used by a limited number of users. This is in line with the session duration which was between 10 and 100 seconds.

V. Lessons learned

The results show that it is crucial to have a well-planned deployment for large-scale events. For this purpose, we present guidelines in this section.

First, of utmost importance is the interference from Wi-Fi, but also RF sources. With this in mind positioning of devices becomes crucial. We have seen that placing them outside, if possible in a high position, with line of sight is important. The interference itself adds around 7dBm and placing them not intelligently adds another 5dBm. Placing them not intelligently makes them practically unusable.

Second, one should also not underestimate the impact of the interference on the performance and especially the latency as this is a crucial metric for QoS. We have seen that the impact is incredibly high. So high, that the network becomes unusable. To better cope with interference, a more sophisticated analytical model is necessary. With this, prediction and accurate planning becomes possible

Third, the Medium Access Control (MAC) protocol needs to more predictive instead of reactive. While being able to send while interference is present works in some cases, it might be more efficient to predict when no interference is present and send at this moment. Currently there is no such protocol present. If all of this fails, switching frequencies is another solution. The E-band in 70GHz and 80GHz offers an interesting prospect for long wireless links. [20]

Final, specifically for hotspot deployments another need is present. We have seen a large difference in distribution of traffic and sessions among APs. This decreases the experience for the users on the highly loaded APs, while there is capacity

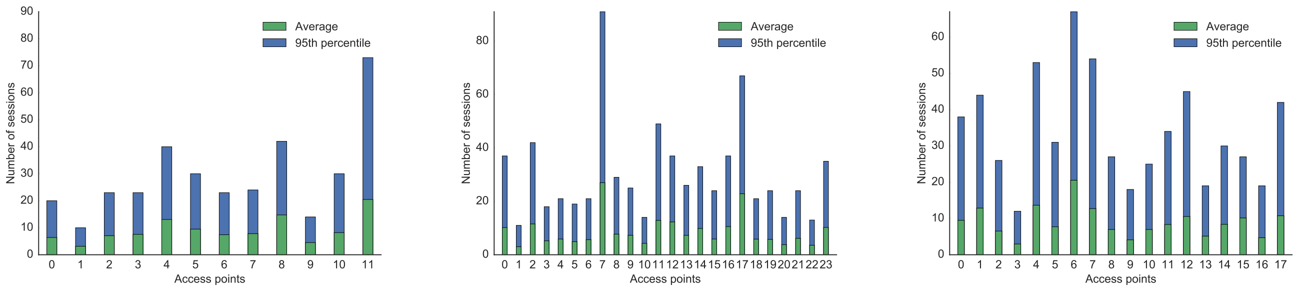

(a) VIP area. (b) Food area. (c) Camping area.

Fig. 6. The distribution of sessions is unbalanced among the APs.

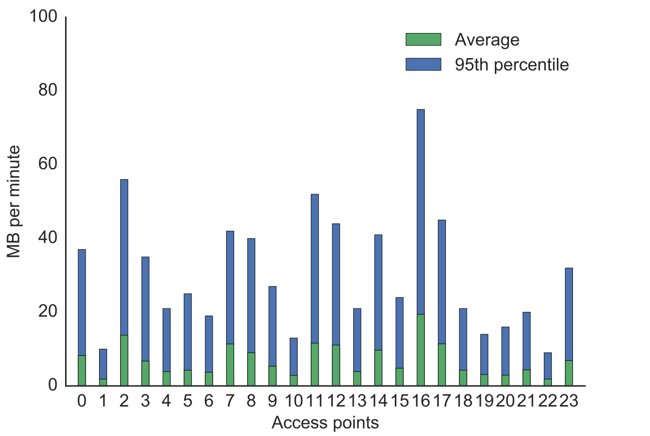

Fig. 7. We can see the average traffic among the APs in the food area and the 95th percentile. It shows a large difference in distribution among the APs.

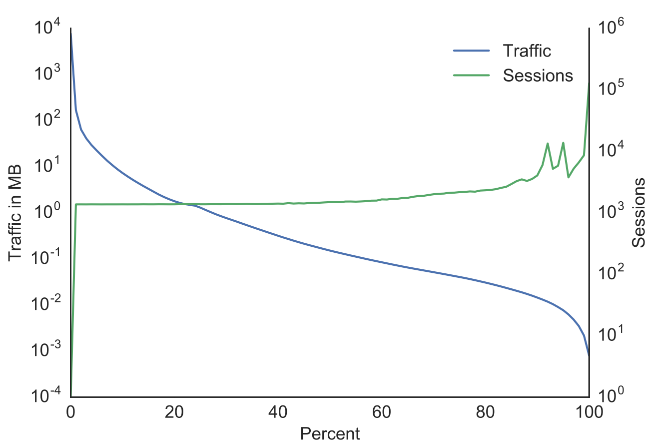

Fig. 8. We can see the amount of traffic per session and the number of sessions with a logarithmic scale in the food area. A low number of sessions had immense traffic, while the majority had a moderate amount of traffic.

available on nearby APs. Advanced load-balancing algorithms are needed that can cope with dynamically changing clients and especially mobility of clients.

VI. Conclusion

In this article, we presented the results and lessons learned of a large-scale measurement study of dense Wi-Fi deployments. We focused both on the performance of a wireless backhaul and a public hotspot at a large-scale music festival. The results show how complicated it is to deploy a wireless backhaul in such a challenging environment, containing multiple RF interfering sources. Both RSSI (below -70dB) and latency values (hundreds of milliseconds) are several factors worse than compared to small scale deployments. This results in strong fluctuations in the topology and the routing behaviour. Based on these results, we presented some guidelines for future deployment that can minimise the impact of interference. However, the results also show that there is a need for a more specific MAC protocol for wireless backhauls in such challenging environments, in order to optimise performance as a response to external interference effects.

The performance of public hotspots in the same geographic area is better since the backhaul itself is using a wired backhaul instead of a wireless backhaul. In this experimental performance study, we therefore focused more on application level performance, characterising the user behaviour during the festival. The results show that it is possible to provide public internet access to users with public hotspots on such a large scale. In terms of user behaviour, the duration of sessions is rather short with the majority between 10 and 100 seconds. This indicates a high mobility, which is expected as festival goers move around the area. The usage per traffic with below 10MB for 90% of the sessions is also rather low. Overall, this article provides several guidelines and lessons learned when deploying and managing Wi-Fi networks on large-scale events.

Acknowledgment

Part of this work has been funded by the iFEST project, co-funded by imec and VLAIO, and partially by FWO.

References

[1] X. Hu, L. Song, D. Van Bruggen, and A. Striegel, "Is There WiFi Yet? How Aggressive WiFi Probe Requests Deteriorate Energy and Throughput," *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*, pp. 317–323, 2015.

[2] R. H. Jhaveri, "Mobile Ad-hoc Networking with AODV : A Review," *International Journal of Next-Generation Computing*, 2015.

[3] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye, "Feasibility study of mesh networks for all-wireless offices," *Proceedings of the 4th international conference on Mobile systems, applications and services - MobiSys 2006*, p. 69, 2006.

[4] L. Cerdà-Alabern, A. Neumann, and P. Escrich, "Experimental evaluation of a wireless community mesh network," *Proceedings of the 16th ACM international conference on Modeling, analysis & simulation of wireless and mobile systems - MSWiM '13*, pp. 23–30, 2013.

[5] S. Vural, D. Wei, and K. Moessner, "Survey of experimental evaluation studies for wireless mesh network deployments in urban areas towards ubiquitous internet," *IEEE Communications Surveys and Tutorials*, vol. 15, no. 1, pp. 223–239, 2013.

[6] A. Michaloliakos, R. Rogalin, Y. Zhang, K. Psounis, and G. Caire, "Performance Modeling of Next-Generation Wireless Networks," *Computer Networks*, vol. 105, pp. 150–165, 2016.

[7] F. M. Abinader, E. P. L. Almeida, S. Choudhury, V. A. Sousa, A. M. Cavalcante, F. S. Chaves, E. Tuomaala, R. D. Vieira, and K. Doppler, "Performance evaluation of IEEE 802.11n WLAN in dense deployment scenarios," *IEEE Vehicular Technology Conference*, 2014.

[8] A. Ghosh, R. Jana, V. Ramaswami, J. Rowland, and N. K. Shankaranarayanan, "Modeling and characterization of large-scale Wi-Fi traffic in public hot-spots," *Proceedings - IEEE INFOCOM*, pp. 2921–2929, 2011.

[9] T. K. Paul and T. Ogunfunmi, "Wireless LAN comes of age: Understanding the IEEE 802.11n amendment," *IEEE Circuits and Systems Magazine*, vol. 8, no. 1, pp. 28–54, 2008.

[10] J. Zhang, K. Tan, J. Zhao, H. Wu, and Y. Zhang, "A practical SNR-guided rate adaptation," *Proceedings - IEEE INFOCOM*, pp. 146–150, 2008.

[11] K. Ramachandran, H. Kremo, M. Gruteser, P. Spasojevic, and I. Seskar, "Scalability Analysis of Rate Adaptation Techniques in Congested IEEE 802.11 Networks: An ORBIT Testbed Comparative Study," *2007 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, pp. 1–12, 2007.

[12] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of RF interference on 802.11 networks," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, p. 385, 2007.

[13] S. Rayanchu, A. Patro, and S. Banerjee, "Airshark: Detecting Non-WiFi RF Devices using Commodity WiFi Hardware," *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, 2011.

[14] M. a. McHenry, P. a. Tenhula, D. McCloskey, D. a. Roberson, and C. S. Hood, "Chicago spectrum occupancy measurements & analysis and a long-term studies proposal," *Proceedings of the first international workshop on Technology and policy for accessing spectrum - TAPAS '06*, pp. 1–12, 2006.

[15] S. Biswas, J. Bicket, T. F. Blvd, S. Francisco, E. Wong, and D. Aguayo, "Large-scale Measurements of Wireless Network Behavior," *ACM SIGCOMM Computer Communication Review*, pp. 153–165, 2015.

[16] "Ubiquiti Nanostation M Datasheet." [Online]. Available: https://dl.ubnt.com/datasheets/nanostationm/nsm_ds_web.pdf

[17] "OpenWrt." [Online]. Available: https://openwrt.org/

[18] "B.A.T.M.A.N. advanced." [Online]. Available: http://www.open-mesh.org/projects/open-mesh/wiki

[19] P. Chatzimisios, A. Boucouvalas, and V. Vitsas, "IEEE 802.11 Packet Delay - A Finite Retry Limit Analysis," *GLOBECOM '03. IEEE Global Telecommunications Conference (IEEE Cat. No.03CH37489)*, vol. 2, pp. 950–954, 2003.

[20] P. Wang, Y. Li, L. Song, and B. Vucetic, "Multi-gigabit millimeter wave wireless communications for 5G: From fixed access to cellular networks," *IEEE Communications Magazine*, vol. 53, no. 1, pp. 168–178, 2015.