Supporting Session and Access Point Mobility in a Large Multi-provider Multi-vendor Municipal WiFi Network

Timo Ojala, Toni Hakanen, Ossi Salmi, Mikko Kenttälä, and Juha Tiensyrjä

MediaTeam Oulu, University of Oulu
P.O. Box 4500, 90014 University of Oulu, Finland
{firstname.lastname}@ee.oulu.fi

Abstract. We present a simple network design for a large multi-provider municipal WiFi network where the WiFi zones of multiple organizations are merged into a single IP subnet at a central layer 2 switch. The design provides built-in session mobility support for the WiFi clients using the network, without any additional software at the client. We also provide a simple design for mobile access points so that they belong to the same IP subnet. We analyze the session mobility of the over 12000 devices using our network in January 2008. We also provide a comparison of the usage of multi-mode mobile devices and WiFi laptops, and characterize the roaming performance of a standard WiFi laptop in our network.

Keywords: municipal, wireless, network design, mobility, roaming, handover.

1 Introduction

This paper addresses session and access point (AP) mobility in a large multi-provider multi-vendor municipal WiFi network provisioned primarily for the purpose of providing open and free wireless Internet access to the general public. With session mobility we refer to a WiFi user device using multiple APs during a continuous application session so that connections are maintained during transitions between APs. Thus, we wish to make a clear distinction to user mobility, which has been used to characterize usage of multiple access points by an individual user over many sessions [2]. With AP mobility we refer to an AP, installed for example in a bus, moving around, possibly among stationary APs. With multi-provider we refer to a municipal wireless network comprising of subnets provided by multiple independent organizations. With multi-vendor we refer to the network comprising of devices produced by multiple manufacturers.

Session mobility in a WiFi network is important from both the application point of view and the network design point of view. This is particularly true for a multi-provider network comprising of subnets provided by multiple independent organizations. If the user device roams between APs in different IP subnets so that the subnet prefix changes, then the IP address of the user device changes. Since the IP address determines the communication endpoints at the application layer, e.g. a 4-tuple TCP connection < source IP address, source port number, destination IP address, destination port number>, roaming and the subsequent change of the IP address

C. Wang (Ed.): AccessNets, LNICST 6, pp. 89–101, 2009.

results in the interruption of all existing TCP connections. So, in terms of session mobility, we want to make sure that the IP address of the user device does not change, even though the device roams from an AP of one organization to an AP of another organization during the session.

AP mobility is equally important, particularly if mobile AP's move around among fixed APs belonging to the same network. If a stationary user device can hear many APs of a network, the device can "ping pong" between different APs. In other words, a user device may associate and reassociate with several APs many times in succession, without moving an inch. This does destroy session mobility, if the APs in question reside in the same LAN for example. But the "ping pong" effect may become a problem if the user device is associated to an AP placed in a bus driving through a cluster of fixed APs, for typically the fixed and mobile APs have different backhaul networks (fixed vs wireless) with different IP address spaces. Again, we want to make sure that the IP address of the user device does not change, even though it roams from the mobile AP to a fixed AP or vice versa.

We present a simple network design for supporting session and AP mobility in a large multi-provider multi-vendor municipal wireless network. The WiFi zones of multiple organizations are merged into a single IP subnet at a central layer 2 switch, i.e. we effectively reduce the multi-provider network into a single-provider case in terms of IP addressing. Session mobility is based on the self-learning property of layer 2 switches, i.e. upon receiving frames from a mobile node the switches automatically learn the current location of the mobile node and build their forwarding tables accordingly. The design is user-friendly in the sense that session mobility is provided without any additional client application at the user device. This is partially motivated by the fact that a large proportion of the users of our municipal wireless network are one time users, visiting our city for business or leisure and happily using our network for Internet access. The design is suitable for a multi-vendor network, where we cannot and do not want to rely on proprietary vendor-specific solutions. However, as the experimental evaluation of the roaming performance will show, the design does not provide good support for highly mobile devices.

This paper is organized as follows. After briefly discussing related work we first provide an overview of our panOULU network [5]. Then we describe the network design and its simple built-in support for session and AP mobility. Our municipal network has been in everyday use for several years, hence we have a plenty of real life traces to work on. After presenting some general usage statistics of the network, we focus on estimating the amount of session mobility in our network. If there is very little session mobility, then there is very little motivation to provide specific support for it and vice versa. We also present quantitative performance characterization of our design in different roaming scenarios. We conclude with a discussion on the strengths and weaknesses of our design, which should prove useful in designing similar networks.

2 Related Work

2.1 Municipal WiFi Networks

In terms of the design of the wireless user access and backhaul, municipal WiFi networks have evolved from the basic ESS (Extended Service Set) topology towards

mesh topology. In the former multiple single-radio APs connected to a wired distribution (backhaul) system provide wireless user access. In leading-edge mesh topologies multiple-radio APs typically provide user access in the 2.4 Ghz band using the 802.11b/g technology, while inter-AP backhaul links are implemented in the 5 GHz band using the 802.11a technology. One motivation is to reduce operational expenses in the long run, as the mesh topology allows cutting down the number of wired backhaul links, possibly down to just 10% of the traditional design. Our wireless network is a mixture of the two topologies.

In terms of combining multiple network providers' WiFi zones into a single access network, in our case each organization first aggregates its APs into a VLAN. These VLANs are then merged into a single aerially large layer 2 network and a single IP subnet. This is a very rare if not a unique design. Boingo acting as the service provider for multiple network providers is a good example of more traditional WiFi roaming. The opposite are commercial open access networks, where the goal is to connect multiple service providers into a single WiFi access network provided by some network operator. The best-known example of commercial open access networks is probably StockholmOpen.net [6]. Its core component is a DHCP relay, which first establishes the initial connection with the user device and then provides a list of ISPs to choose from. Once the user has selected the ISP, the DHCP relay redirects the connection to the selected ISP, which assigns IP address and other network settings to the user device.

Another type of multiple-provider networks include community networks such as SeattleWireless [7] or FreiFunk [8], where the members of the community agree on granting each other access to APs and forwarding traffic. If a mesh topology is used, then only a fraction of the APs need to have backhaul access.

2.2 Session Mobility

We can identify two principal levels of session mobility. On the first level we just maintain active sessions (connections) during transitions between APs, which is effectively a mobility management problem. On the second, more advanced level we provide session mobility with some level of QoS. It could for example refer to communication delays or packet loss, which are very relevant to real-time interactive applications such as VoIP. If no specific support is provided, even roaming between APs of the same technology (horizontal handover) and belonging to the same subnet can lead to packet loss and degraded quality of service. This is mainly due to the long latency in the network learning of the new location of a highly mobile node to route packets accordingly.

Interestingly, the original IEEE 802.11 specifications did not specify the intra AP communication protocol to inform the network about a client's roaming. This led to many proprietary vendor-specific implementations and interoperability concerns. Later, in 2003 the 802.11f IAPP (Inter-Access Point Protocol) protocol was published as a trial standard for this purpose. However, it did not gain popularity in the industry and consequently was withdrawn in 2006. The current standardization effort is 802.11r ("Fast Roaming/Fast BSS Transition"), which is scheduled to be published in 2008. In 802.11r the client can use the current AP as a conduit to other APs, which allows the client to establish security and QoS state information at a new AP before making a transition.

A number of mobility management technologies have been presented for higher layers of the protocol stack, for example Mobile IP [9] at the network layer, TCP-Migrate [10] at the transport layer, and SIP [11] at the application layer. Further, WiFi vendors have developed proprietary mobility managers for maintaining seamless connections for mobile nodes traversing between subnets. The common goal of the mobility management technologies and mobility managers is to provide location transparency for a mobile node. The main challenge in their deployment is that they require possibly costly changes to the network infrastructure and/or mobile node.

A number of studies on session or user mobility in WiFi networks have been conducted in the past. For example, in their analysis of the usage of a large campus network Henderson *et al.* [3] reported that 95.1% of users had a home location, and 50% of those users spent 98.7% of their time there. Balazinska and Castro [2] found in their study of a corporate wireless network that 50-80% of users fell into the occasionally and somewhat mobile categories. In our earlier study [4] based on a 14-months long monitoring of our network covering both the city and university campus, we found that on average 9% of the sessions in the city were mobile, while the corresponding figure on the campus was 5%. These studies and reported findings are not directly comparable as different things were measured in different contexts.

3 panOULU Network

3.1 General Description

panOULU is a municipal wireless network in the City of Oulu in Northern Finland. With about 130000 citizens Oulu is the sixth largest city in Finland. The Oulu region is well-known for its ICT competence and with about 14000 ICT jobs has by far the largest regional R&D expenditure per capita in Finland. Back in 1998 the Wired Magazine ranked Oulu as #3 'silicon valley' in the world, after the original Silicon Valley in California and Austin in Texas. The city's central administration is very much pro ICT and pro R&D, which has proven very valuable over the years, also for the panOULU network.

panOULU comprises of two basic types of subnets: campus networks of five public organizations and so-called panOULU subscriptions sold by three ISPs. The City of Oulu, University of Oulu, Oulu University of Applied Sciences, VTT Technical Research Centre of Finland, and Pulmonary Association Heli each have large premises around Oulu covered by WiFi networks. Each organization is responsible for the expenses and the management of its own campus network. panOULU subscription is an ISP product, and by purchasing it any organization can acquire panOULU coverage into its premises.

As of now the campus networks and panOULU subscriptions total 1050 APs. The APs come from multiple vendors (Cisco, Linksys, Strix, Buffalo) and have different types of backhaul links (Ethernet, xDSL, IEEE 802.11a, Flash-OFDM). From the user's point of view the APs appear as one large uniform network with SSID 'panoulu'. The APs provide both indoor and outdoor coverage in places deemed relevant for public access. The city center and its immediate surroundings are blanketed with a WiFi mesh network, but otherwise the coverage is provided in a

hotspot manner. In some locations such as the city hospital providing public access is secondary and supplementary to actual production use of wireless access to patient databases.

In its coverage area the panOULU network provides open (no login, authentication or registration) and free (no payment) wireless internet access to the general public with a WiFi-equipped device. After associating to the network the device is granted a private IPv4 address created with NAT and allocated with DHCP, and a public IPv6 address. Public IPv4 addresses are also available upon request for R&D use. The first HTTP request of a particular device on a particular day is redirected to a splash page providing basic information about the network. Excluding the blocking of outgoing port 25 (SMTP), which is required by the Finnish legislation, there are no limitations or restrictions on the use of the network. Other services include SMTP server for sending email, help desk during office hours and panOULU Luotsi [12]. It is a mash up combining real-time AP-ID positioning of the user device and various information feeds into a location-based web service with map-based guidance.

3.2 Network Topology

A simplified version of the panOULU network topology is shown in Figure 1. Each organization contributing a WiFi zone depicted as a green cloud groups its APs into a VLAN. The VLANs are then aggregated at the central layer 2 switch #1 residing at the DNA server room. The WiFi clouds labeled DNA, Elisa and Netplaza correspond to the APs of the panOULU subscriptions (see Section 3.3) sold by the three respective telcos. OukaMESH is a 60 AP WiFi mesh network built with Strix System's OWS-2400 series APs. Other WiFi clouds are built with Cisco 1100/1200 series, Linksys WRT54GL and Buffalo WHR-G54S APs using the 'traditional' approach of hooking up each AP into the backbone with different types of wired backhaul (xDSL, Ethernet). The Mobile APs are explained in detail in Section 3.4.

The server farm providing the core services (e.g. DHCP, DNS, HTTP, SMTP) is located at the City of Oulu's server room. The server farm residing in the DNA server room includes a high performance probe for collecting packet headers from the central switch for monitoring purposes. The 100 Mbps Internet uplink is provided by Oulu University of Applied Sciences. The management of the server farms is currently sponsored and conducted by the City of Oulu, but will soon be purchased as a service from a company.

The straightforward "KISS" layer 2 design merges the organizations' WiFi zones into a single IP subnet. Among other things the design provides built-in session mobility support for the WiFi clients, which was one of the original design decisions of the panOULU network. Session mobility is based on the self-learning property of the layer 2 switches used to connect the APs into the backbone (not necessarily switches #1 and #2 shown in Figure 1). When a mobile node moves between APs in two different BSSs (Basic Service Sets), the layer 2 switch connecting the two BSSs will eventually receive a frame from the mobile node, thus automatically learning the new location of the mobile node and updating its forwarding table accordingly. However, as we will see in the experiments, this approach is not really optimal for highly mobile users.

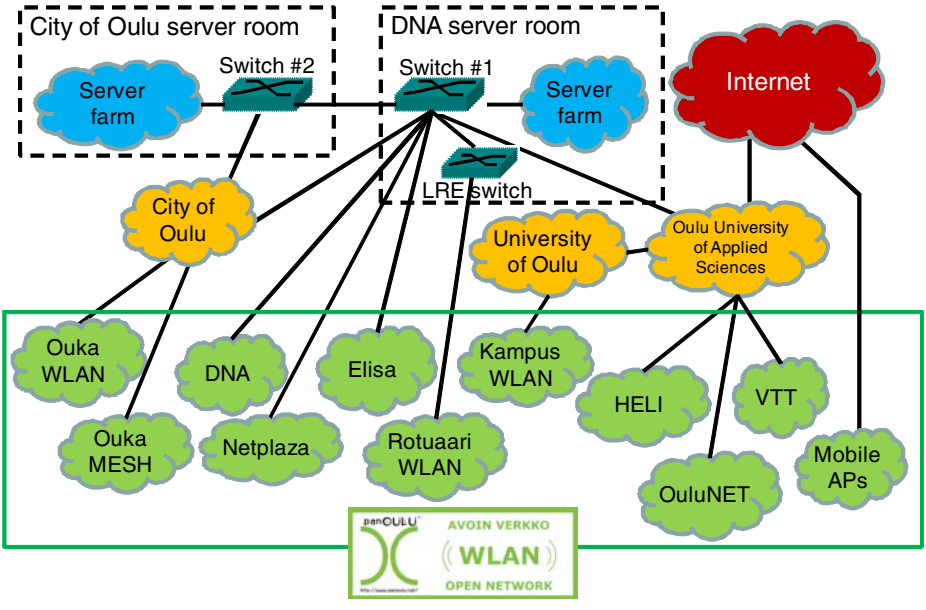

Fig. 1. Simplified topology of the panOULU network

3.3 panOULU Subscription

panOULU subscription is a panOULU compatible product offered by three ISPs. It includes both a regular business subscription with whatever services the customer wishes to purchase, and a panOULU hotspot providing open and free wireless Internet access at the customer's premises. The idea behind this bundling is that the customer does not have to purchase two separate services in order to provide a panOULU hotspot. panOULU subscription is a leasing package in the sense that the ISP own all hardware. Hence, the subscriber does not have to invest in WiFi APs or other devices, but pays for them in the monthly fee. The subscription also includes installation and maintenance.

Figure 2 illustrates the network topology of a panOULU subscription. It is implemented with one physical link (the red solid line from operator's point-of-presence to the customer's premises), typically an xDSL, atop which two VLANs are set up. The first VLAN is for the regular business subscription of the restaurant manager and its traffic is routed via the ISP's Internet uplink (green dashed line below). The other VLAN connects the WiFi AP(s) providing the panOULU hotspot in the restaurant into the panOULU's central switch. The traffic of the WiFi AP(s) is routed via the panOULU's Internet uplink (blue/yellow dashed line). We see that the AP installed inside the restaurant resides in the same layer 2 network as other panOULU APs.

The organizations having purchased a panOULU hotspot into their premises for the purpose of improving their customer service and image include Oulu Airport, Oulu Cathedral (!), Technopolis (one of Europe's largest technology center operators), a

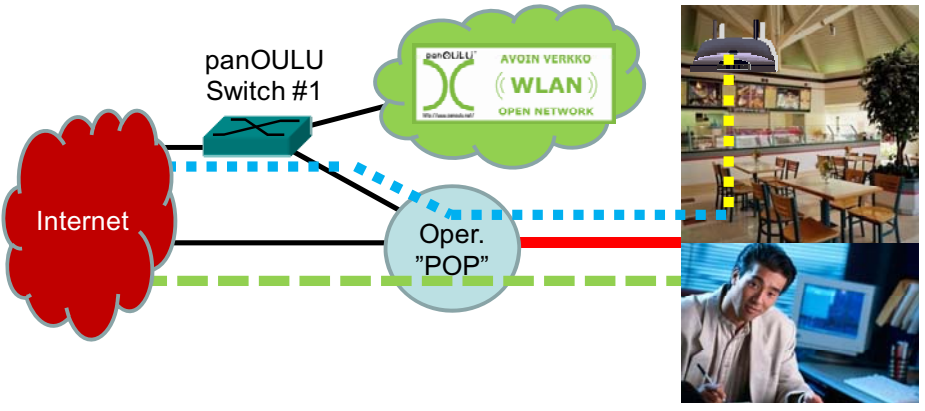

Fig. 2. Network topology of a panOULU subscription

large training and management institute, a large sports complex, a private hospital, the largest department store in the city, a large bank, many media and IT companies, and many cafes, pubs and restaurants. With the inception of the panOULU subscription we have managed to incorporate telco incentive in municipal wireless offering.

3.4 Mobile APs

The panOULU network has about 15 mobile APs, residing in buses, a ferry and a mobile library. The simple network topology of a mobile AP is illustrated in Figure 3. The Linksys WRT-54 GL AP placed in front of the bus is connected to a Flash-OFDM modem, which provides wireless backhaul over the @450 Wireless Broadband network of Digita [1].

The @450 network uses the 450 MHz band and the Flash-OFDM technology originally developed by Flarion. It provides theoretical maximum 5.3 Mbps downlink and 1.8 Mbps uplink data rates per sector. The @450 service can be purchased either with 1 Mbps or 512 kbps downlink data rate. The nominal radius of a @450 cell is 30 km, but with a high-gain directional antenna the range can be easily doubled outdoors. The @450 MHz network provides very good mobility support and about 50 ms RTT.

The @450 network has naturally its own IP address space different from that of the panOULU network. The WiFi clients of the mobile AP are provided with a regular panOULU IP address as follows. Upon startup the mobile AP creates a VPN tunnel through the @450 network to the panOULU core. A WiFi client associating with the mobile AP obtains an IP address from the panOULU DHCP server, just like with a regular stationary AP connected ‘directly’ to the panOULU network.

Hence, from the clients’ point of view, the mobile AP is in the same IP subnet as stationary APs. Consequently, a client’s possible “ping ponging” between a mobile AP in a bus and a nearby stationary AP would not have any effect on the IP address and on-going connections.

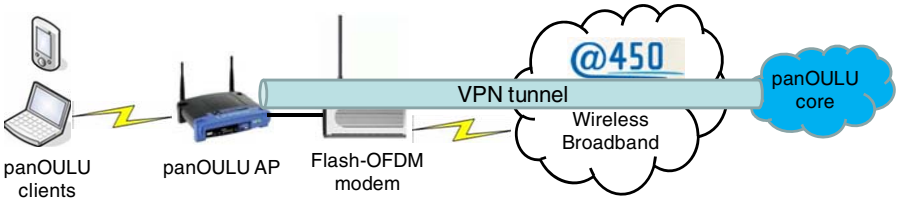

Fig. 3. Network topology of a mobile AP

We assessed the behavior of WiFi clients aboard panOULU buses by riding back and forth through the city center, which is the most probable place for the "ping pong" effect to take place. However, once the WiFi client has established a connection with the AP inside the bus, it never strayed to an outside stationary AP nearby the bus.

4 Session Mobility in panOULU Network

4.1 General Network Usage Statistics

In January 2008, 12118 unique WiFi devices used the panOULU network, totaling 279850 sessions and 12 million minutes of online time. A unique device is identified by its unique MAC address. A session is defined to start when the DHCP server acknowledges a client device's DCHP REQUEST with DHCP ACK. A session is deemed to end when the device does not respond to four consecutive ***arping*** calls made at 60 second intervals. The session is clocked to have ended at the time of the first unanswered ***arping*** call. Online time corresponds to the duration of the session. The devices are mostly PCs furnished with the Windows operating system. The proportion of Nokia manufactured multi-mode mobile handsets equipped with both cellular and WiFi radios has been growing steadily so that they make up 15-20% of all devices now.

As discussed earlier, one of the motivations behind panOULU network is to provide visitors coming to Oulu free and convenient Internet access. The graph in Figure 4 provides a weak estimate on the proportion of visiting users, based on identifying the devices that had not been seen in the network before during year 2007. We see that up to 40% of the devices using the network in a given month had not used the network before during year 2007. Of course, we have no way telling if these 'new' devices actually correspond to devices owned by visiting users, or if they are just new devices acquired by local panOULU users.

A somewhat more reliable estimate can be obtained by profiling individual devices based on their network usage patterns. For this purpose we define as a "heavy user" a device that uses the network at least 50% of the days in the analysis period, i.e. in a 30-day period a "heavy user" has used the network on 15 days minimum. As a "one time user" we regard a device that has used the network on at least four days during a period of at most one week in length. Devices not categorized as "heavy users" or "one time users" are regarded as "casual users". Using this categorization the 12118

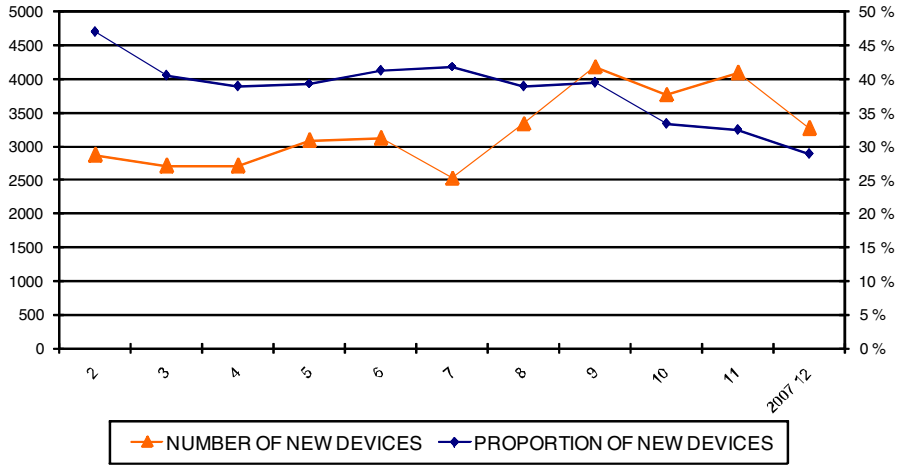

Fig. 4. Numbers and proportions of new devices in the panOULU network in 2007

devices seen in the network in January 2008 were classified as follows: 9.8% "heavy users", 70.0% "one time users" and 20.3% "casual users".

4.2 Session Mobility

We have studied session mobility, in order to obtain solid evidence on the need to provide mobility support. To quantify session mobility we define a session to be mobile, if during the session the WiFi device uses at least three APs so that at least two of them are 50 meters apart. Employing this definition only 5.1% of the 279850 sessions in January 2008 were mobile.

Another way at looking mobility is to analyze how many devices have a so-called "home AP", via which the majority of the network usage of this particular device takes place. A device is deemed to have a "home AP", if 0.7 x (homeAPsessions / totalsessions) + 0.3 x (homeAPonlinetime / totalonlinetime) $\geq$ 0.5. Using this criterion a whopping 91.5% of the 12118 devices had a "home AP".

We have studied in more detail the macro mobility of devices between administrative domains (providers), to obtain solid evidence on the need for IP mobility support such as the one provided by the layer 2 design of our multi-provider network. The most prominent place to study macro mobility is the city center where five of the eight providers in the panOULU consortium have at least one AP. We geographically limit our study by placing a circle of 750 m in radius at the city centre, taking into consideration all 409 APs residing within the circle. As a single-provider reference region we use the university campus, placing a circle of 450 m in radius at the university main hall, and taking into account the 278 APs within the circle.

Table 1 lists the key statistics for the two regions in January 2008. In addition to mobility we also include other statistics allowing other comparisons between the two different usage environments. A provider crossing refers to the event where a device roams from an AP of a particular network provider to an AP of another network provider during a session.

Table 1. Comparison of network usage at the city center and the university campus

Statistic	City center	University campus
APs	409	278
Devices	6031	2646
• Nokia multi-mode	929 (15.4%)	386 (17.5%)
• "Home AP"	5653 (93.7%)	2314 (87.5%)
• "Heavy user"	579 (9.6%)	159 (6.0%)
• "One time user"	4516 (74.9%)	1829 (69.1%)
• "Casual user"	936 (15.5%)	658 (24.9%)
Sessions	128480	44396
• Mobile	10376 (8.1%)	2205 (5.0%)
• Provider crossing	4465 (3.5%; 43.0%)	-

Table 2. Comparison of network usage between multi-mode mobile devices and other devices

Statistic	Nokia multi-mode devices	Other devices
Devices	929	5102
• "Home AP"	864 (93.0%)	4789 (93.9%)
• "Heavy user"	8 (0.9%)	571 (11.2%)
• "One time user"	862 (92.8%)	3654 (71.6%)
• "Casual user"	59 (6.4%)	877 (17.2%)
Sessions	4990	123940
• Mobile	518 (10.4%)	9858 (8.0%)
• Provider crossing	236 (4.7%; 45.6%)	4229 (3.4%; 42.9%)
• Per device	5.4	24.3
• Average duration (mins)	64	1001
• Median duration (mins)	6	37

We see that in the city center 4465 sessions, 3.5% of all 128480 sessions and 43.0% of the 10376 mobile sessions, involve a provider crossing. It is difficult to judge if these 4465 sessions justify the current layer 2 design, partially motivated by our desire to provide easy support for session mobility.

Although the amount of session mobility appears small at this point, we expect it to increase as multi-mode mobile devices become more popular. A PC is not really intended to be used in a mobile manner. Table 2 provides a compact comparison of the usage statistics of Nokia made multi-mode devices and other devices (largely PCs) within the same 750 m circle at the city center in January 2008.

We see that the multi-mode devices have a clearly smaller number of sessions per device (5.4 vs 24.3), and their median duration is much shorter (6 vs 37 minutes) as expected. The fact that many residents at the city centre have their PCs constantly connected to the network skews the averages. Multi-mode devices have a slightly larger proportion of mobile sessions (10.4% vs 8.0%), but the difference is not significant. The usage of multi-mode devices is much more active on the university

campus than in the city center with 9.7 sessions per device of much longer duration (average 154 min, median 28 min).

4.3 Roaming Performance

So, our network contains a simple built-in mobility support for WiFi devices. We conclude the experimental part by characterizing the roaming performance in our network with two simple experiments. With roaming we refer to the WiFi device roaming from a particular AP to another AP. We are particularly interested in the performance at the application layer, how long the application processes are "offline" and what is the expected packet loss.

The client device used in both experiments was a Lenovo X60S laptop with Ubuntu OS and an integrated Intel PRO/Wireless 3945ABG WiFi card. During the experiments the laptop was moved spatially at walking space to force roaming. The experiments were executed in different locations to study roaming in varying radio conditions and between different pairs of APs, from a Cisco 1200 series AP to another Cisco 1200 series AP, from a Strix OWS-2400 series mesh AP to a Cisco 1200 series AP and vice versa. The beacon rate of each AP was set to 100 ms. During the experiments the network was being normally used by arbitrary other clients, which may have had an impact on our results.

In the first experiment we measured the throughput of a TCP connection between the WiFi laptop and a server using the ***iperf*** TCP/UDP bandwidth measurement tool [13]. This experiment serves two purposes, firstly verifying session mobility, i.e. the TCP connection survives the roaming, and secondly quantifying the "offline" time of the TCP connection during roaming. The ***iperf*** tool was configured to use 128 KB TCP windows and to measure TCP throughput over 0.5 s intervals. While the TCP throughput was typically in the 1-5 Mbps range during successful data transfer, immediately after handover between two APs the connection was "offline" (throughput 0 bps) for a lengthy period ranging from 8 s to sometimes even over 20 s. This long latency is due to problems in packet routing. When the laptop moves between APs in two different BSSs, the layer 2 switch connecting the two BSSs is not informed. Eventually, it will learn of the new location of the laptop, when it receives a frame from the laptop. However, this self-learning property of the switches is not designed to support highly mobile users who want to maintain TCP connections while moving between APs in different BSSs.

In the second experiment a client process in the WiFi laptop generated every 20 ms a 160-byte payload containing a timestamp and a running sequence number and the payload was encapsulated into an UDP packet. This type of traffic would be very typical for a VoIP call, for example. The client process sent the packet into an UDP socket for immediate transmission to a server process, which stored all arriving packets. By analyzing the gaps in the arriving packet sequence we could quantify the "offline" time the client process experienced due to roaming. Typical minimum "offline" time of the client process during a roaming was 2.6 seconds corresponding to about 130 lost packets. The maximum "offline" time measured several times was staggering 5.6 seconds with 280 lost packets! In this case the latency is due to the upper layers of the protocol stack, as layer 2 handover takes clearly less than 100 ms.

The experiments show what kind of roaming delays a client process running on a standard off-the-shelf WiFi laptop can expect to experience in our network, when no special measures are taken to speed up the handover.

5 Discussion

We presented a simple layer 2 network design for merging the WiFi zones of multiple organizations into a large multi-provider municipal WiFi network that in terms of IP addressing appears as a single IP subnet. We also provided a simple design for mobile APs so that they are part of the same IP subnet. The design provides built-in session mobility support for the WiFi clients using the network, without any specific support either in the network or in the client. The latter requirement is partially motivated by the fact that a large proportion of the users of our network are one time visitors. We cannot expect them to install any particular client software just for the sake of having session mobility in our network.

We analyzed the session mobility of the over 12000 devices using our network in January 2008. The proportion of mobile sessions is below 10%. About 40% of the mobile sessions at the city center involved a "provider crossing", where our layer 2 design comes handy in terms of keeping the connections alive. We also provided a brief comparison of the usage statistics of multi-mode mobile devices and WiFi laptops. Finally, we characterized the roaming performance of a standard WiFi laptop in our network with two simple experiments.

We believe that our "KISS" layer 2 network design can serve as a good model for similar municipal multi-provider WiFi networks. The design effectively reduces a multi-provider network into a single-provider case, which requires the providers to agree on common IP addressing. This could impose severe limitations in terms of AAA (Authentication, Authorization and Accounting). However, since we currently do not authenticate users nor charge them for the network usage, AAA is not really an issue for us.

The most intense debate we have had has focused on the geographically large panOULU VLAN of the City of Oulu, which spans pretty much every public office and service point. The large VLAN is suspect to large amounts of broadcast traffic and broadcast storms. We have measured the broadcast traffic at the central switch to average about 37 kbps over 24-hour periods. We have had two major broadcast storms in the city's panOULU VLAN due to a faulty OS in newly acquired AP's. The software bug in the OS resulted in an incoming IGMP packet being reflected into all outgoing ports. Consequently, a single IGMP packet was replicated into millions in a blink of an eye, forcing a temporary shutdown of the VLAN.

In the near future we will expand the WiFi coverage to key locations in nearby ten townships. We are also currently building a Mobile WiMAX network around the City of Oulu, to be part of the open and free wireless access provided by the panOULU network.

Acknowledgments. The authors wish to thank the members of the panOULU consortium for their support.

References

1. @450 Wireless Broadband network, http://www.450laajakaista.fi/
2. Balazinska, M., Castro, P.: Characterizing mobility and network usage in a corporate wireless local-area network. In: The First International Conference on Mobile Systems, Applications, and Services, pp. 303–316 (2003)
3. Henderson, T., Kotz, B., Abyzov, I.: The changing usage of a mature campus-wide wireless network. In: Tenth Annual International Conference on Mobile Computing and Networking, pp. 187–201 (2004)
4. Ojala, T., Hakanen, T., Mäkinen, T., Rivinoja, V.: Usage analysis of a large public wireless LAN. In: 2005 International Conference on Wireless Networks, Communications and Mobile Computing, pp. 66–667 (2005)
5. panOULU network, http://www.panoulu.net/
6. Pelletta, E., Lilieblad, F., Hedenfalk, M., Pehrson, B.: The design and implementation of an operator neutral open wireless access network at the Kista IT-University. In: 12th IEEE Workshop on Local and Metropolitan Area Networks, pp. 149–150 (2002)
7. SeattleWireless, http://seattlewireless.net/
8. FreiFunk, http://freifunk.net
9. Perkins, C.: Mobile IP: Design Principles and Practices. Addison-Wesley, USA (1998)
10. Snoeren, A.C., Balakrishnan, H.: An end-to-end approach to host mobility. In: 6th ACM/IEEE International Conference on Mobile Computing and Networking, pp. 155–166 (2000)
11. Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., Schooler, E.: SIP: Session Initiation Protocol, RFC 3261 (2002)
12. panOULU Luotsi, http://luotsi.panoulu.net
13. Iperf, http://sourceforge.net/projects/iperf