

Exploiting Regular Hot-Spots for Drive-thru Internet

Jörg Ott and Dirk Kutscher

Technologiezentrum Informatik (TZI), Universität Bremen,
Postfach 330440, 28334 Bremen, Germany,
{jo|dku}@tzi.uni-bremen.de

Abstract. IEEE 802.11 WLAN technology has become an inexpensive, yet powerful access technology that is targeted at mobile users that remain within reach of the hot-spot. Such hot-spots are usually provided by a wireless Internet service provider (WISP) in locations often frequented by travelers. Past measurements have shown that WLAN is even able to support mobile users passing by without stopping and “hopping” from one hot-spot to the next. The Drive-thru Internet project develops a disconnection-tolerant architecture that enables such unconventional usage of WLAN technology. In this paper, we focus on two prime aspects relevant for interworking with existing hot-spot installations: we investigate the impact of auto-configuration and authentication and present performance results for a driving user accessing the Internet via a hot-spot using different access link technologies. We finally suggest enhancements to hot-spot architectures to facilitate Drive-thru Internet access.

1 Introduction

Mobile users and nomadic computing are today supported by two classes of networks: On one hand, cellular networks (GSM, GPRS, UMTS) aim at providing ubiquitous connectivity, even across different service providers. However, their price-performance ratio is rather poor and temporary disconnections may occur despite the wide coverage for a variety of reasons. On the other hand, IEEE 802.11 WLAN hot-spots do not aim at seamless connectivity; their limited reach implies disconnection periods while the user is moving between locations. Usually, manual user interaction is required, e.g., to suspend and resume communication applications but also for reconfiguration [OK04c]. Nevertheless, the availability of high data rates at acceptable cost render WLAN technology an attractive alternative for mobile usage scenarios. Because of unlicensed operation and low investment and operational cost, WLAN has become an inexpensive commodity and the number of public WLAN hot-spot installations is ever-increasing: besides hotels, cafés and the like particularly airports, train stations, gas stations, and service areas are covered, i.e., places serving commuters and travellers on the road.¹

Numerous approaches are pursued that combine access to different service providers or integrate WLANs and cellular networks to enhance connectivity (particularly for WLANs), improve the achievable data rate, and minimize cost (for cellular networks) to keep users *always best connected* [RCC⁺04] [ZWS⁺03] [Lei01]. While

¹ Examples include Agip gas stations and MAXI service areas in Germany, Neste A24 gas stations in Estonia, and Texaco service stations in the UK as well as truck stops in the US.

such approaches render existing wireless technologies more attractive to users, the potential of temporary disconnection (from an affordable high performance link) remains.

In the Drive-thru Internet project [OK04a], we leverage conveniently located WLAN hot-spots to provide Internet services to mobile users, focusing on users moving at high speeds (in vehicles). The objective of Drive-thru Internet is to enable access to Internet services by using *intermittent connectivity*, i.e., connectivity that is only temporarily established while a user traverses the coverage area of a WLAN hot-spot. We have developed an architecture that allows existing and future applications to take advantage of such potentially short and unpredictable periods [OK04b].

A key requirement for Drive-thru Internet is the ability to operate in today's *existing* WLAN infrastructure, which mostly consists of public hot-spots. We allow for incremental deployment by avoiding dependencies on specific service providers and (modifications to) hot-spot architectures. To achieve this independence, we need to take the characteristics of commercial hot-spot installations into account and develop support functions that allow us to use an existing hot-spot "as is" for Drive-thru Internet.

This paper explores the issues related to obtaining Internet access from existing WLAN hot-spots in a way suitable for Drive-thru Internet: we investigate implications of access link characteristics and analyze steps towards efficient user authentication and access authorization. In section 2, we review representative architectures for commercial hot spots and outline the Drive-thru Internet concept in section 3. Related work concerning these two fields is discussed in the respective sections. We then report on our findings from real and experimental hot-spot settings and introduce our approach towards enabling the use of commercial hot-spots in section 4 where we also derive desirable hot-spot properties. Section 5 concludes this paper and suggests future work.

2 Wireless LAN Hot-Spots

Early WLAN hot-spot installations concentrated on connectivity and did generally not provide sophisticated user authentication mechanisms. However, emerging regulatory requirements, identified WLAN security issues and, in particular, commercial interests have led to the development of *WLAN hot-spot architectures* that provide a whole set of functions beyond the basic provisioning of network access. [HKR⁺03] provides an overview of typical hot-spot features including: *enabling WLAN access* (WLAN association etc.), *provisioning the hot-spot* (device and user authentication), *IP layer management* (auto-configuration, DNS, NAT, etc.), *providing access to a hot-spot LAN for local information services*, *providing WAN access*, and *providing accounting information*.

How these functions are distributed across different components in a hot-spot depends on its specific architecture. An example is shown in figure 1: The wireless access link can be composed of a *service set* of WLAN access points providing coverage for the hot-spot area. In addition, there is an access controller that controls client access and a local access router (e.g., a DSL router) among other network elements. The *user authentication* function consists of a local front end in the hot-spot and, for larger hot-spot operators, usually an external AAA server. So far, no common authentication method has been established. Instead, different incompatible methods are in use, such as web-based login, SIM-cards [ACDS03], and IEEE 802.1X. For commercial hot-spots, the




Fig. 1. Sample hot-spot architecture

most common method is *web-based login* which is recommended by the Wi-Fi Alliance [ABS03] under the name *Universal Access Methods* (UAM) and works as follows:

The user's mobile device connects to the hot-spot WLAN that has SSID broadcast enabled and does not use WEP. A DHCP server supplies the necessary IP and DNS parameters but access to the Internet is still disabled at this point. The hot-spot access control function intercepts the first HTTP request from the user's device and redirects the user's web browser to the operator's login page. The login page prompts the user to enter her credentials. If the user is authorized network access is granted, typically by means of MAC or IP address filtering. Terminating the use of a hot-spot (and thus stopping accounting) may also be done via the web browser and is usually complemented by inactivity detection. The advantage of UAM is that it imposes very little requirements on user equipment. The user can use commodity WLAN NICs and is not required to install any special-purpose software, e.g., VPN clients.²

WLAN roaming is still not generally available. However, there is an increasing number of hot-spots where the local operator is not identical to the actual WISP. E.g., at some airports the user is presented a selection of WISPs in the initial UAM web page and can select the preferred WISP before logging in. So called hot-spot aggregators [Wir03] such as *Boingo Wireless* and *iPass* have roaming agreements with selected hot-spot operators and offer some form of international roaming – at selected hot-spot sites.

For commercial hot-spots, one of the more important characteristics with respect to mobile usage scenarios is the *tariff model* employed by the WISP, where essentially three different models can be distinguished: volume-based, connection-time-based and flat-rate tariffs. In most countries, true flat-rate access is currently rather uncommon.³ With respect to billing, pre-paid and post-paid models can be differentiated: Mobile phone operators that run a larger hot-spot network often apply accounting models for GSM-based Internet access to WLAN: network usage is accounted for on a time basis,

² It should be noted that different strategies for providing users with credentials for the UAM login exist: For example, users may purchase vouchers that carry an authentication code that has to be entered into the HTML form of the login page. For operators with existing business relationships with their WLAN users, e.g., for mobile phone operators, it is considered beneficial to bill the user on her regular mobile phone bill; therefore the credentials are sent in text messages (SMS) to the user's GSM phone.

³ In some countries, e.g., the US, flat-rate WLAN access is already available. In addition hot-spot aggregators often offer flat-rate tariffs, however with some restrictions with respect to specific hot-spots. E.g., Boingo Wireless offer flat-rates in general, but many international, i.e., outside the US, hot-spots are considered *premium sites* where additional charges apply.

and the user is billed on the regular mobile phone bill, which means the time budget is not limited in advance. For the pre-paid model, there are significant differences with respect to the accounting granularity: For voucher-based authentication, the user is typically authorized to use the WLAN at a location for a certain duration, e.g., one hour, without being able to suspend the session. In contrast to this one-time access model, some WISPs allow for using the time budget more flexibly, in multiple sessions. The granularity differs and ranges from one minute to one hour.⁴

Summarizing, we can state that while more and more WLAN hot-spots are established, there are still some open technological and deployment challenges that have to be overcome in order to use WLAN hot-spots as a ubiquitous infrastructure. The existing infrastructure does, in general, not provide global roaming and relies on a set of different authentication, accounting and billing strategies.⁵

3 Drive-thru Internet

In the Drive-thru Internet project, we exploit WLAN connectivity from hot-spots along the road to provide Internet access to mobile users passing by who will experience intermittent connectivity. Our past work concentrated on the WLAN link: we carried out extensive laboratory and field measurements investigating the communication characteristics between a mobile node in a car and a fixed one co-located with the access point(s) under a variety of conditions. The results were convincing and clearly proved the feasibility of WLAN as last hop access even at higher speeds: in our current IEEE 802.11g setup we are able to obtain some 1800 m of connectivity [OK04b].

Throughput measurements showed a bell-shaped curve that led us to identify a *three-phase model*: during the *entry* and *exit phases* communication is possible at limited performance due to low link layer bit rate, link layer retransmissions, and packet losses. A *production phase* with stable connectivity and constantly high throughput approaching stationary conditions is available for up to 1000 m [OK04a] [OK04b]. With link data rates of 54 Mbit/s for large parts of the production phase and net data rates on top of TCP of up to 16 Mbit/s, we achieved transfer volumes of more than 70 MB in a single pass at 120 km/h. While traffic (and weather) conditions influenced the performance, we obtained a minimum of 20–30 MB across all measurement settings. Further experiments with UDP and TCP background traffic and with two mobile nodes in different cars showed that the TCP flows to the mobile nodes obtained a reasonable share of the available link capacity when competing with fixed nodes or each other. Congestion control adapted quickly to the changing link and traffic conditions. Altogether, our past measurements showed that IEEE 802.11 WLAN is a suitable communication substrate for providing high-performance network access to fast moving mobile nodes.

Nevertheless, the short and unpredictable connectivity periods pose various challenges to applications usually built under the assumption of rather stable network access conditions [OK04c]. We have devised the Drive-thru architecture [OK04b] to deal with

⁴ For example, Deutsche Bahn offers a pre-paid tariff, where users purchase a budget of 8 hours that can be used at different Deutsche Bahn hot-spots with a granularity of one minute. However, most WISPs account for network usage at a coarser granularity.

⁵ A detailed discussion of challenges and future directions for hot-spots is provided in [BVB03].

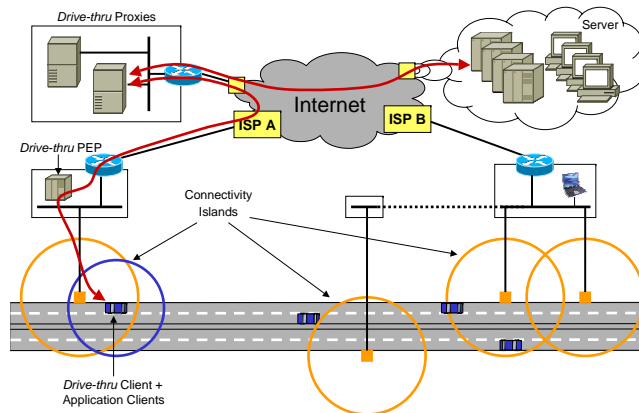


Fig. 2. Overview of the Drive-thru Internet Architecture

unstable connectivity as well as transport and application layer timeouts. As shown in figure 2, we apply an enhanced variant of *connection splitting* [BKG⁺01] and introduce two intermediaries: The *Drive-thru client* is co-located with the mobile node and the *Drive-thru proxy* is placed somewhere in the fixed network.⁶ These two entities terminate transport (and application) connections to the application client (e.g., an e-mail client or a web browser) and the corresponding server, respectively, and protect the application entities from the intermittent nature of connectivity. Drive-thru client and proxy use a TCP-based “session” protocol—the *Persistent Connection Management Protocol* (PCMP) [OK05]—to provide transport connections that persist across connectivity islands and allow for continuous exchange of larger data volumes.

Drive-thru Internet differs fundamentally from other approaches using wireless (LAN) technologies for communication to and between mobile users in vehicles: Often users in airplanes, buses, or trains [TMC04] are served as a group via classic WLAN with WWAN connectivity to the outside. MAR [RCC⁺04] and IPonAir [ZWS⁺03] implement the seamless integration of different networks for individual users, as mentioned in the introduction. Fleetnet [BFW03] requires a dedicated wireless infrastructure with specific lower layer protocols, primarily targets new applications and does not address short-lived connectivity whereas Drive-thru Internet leverages existing WLAN hot-spots for existing applications. Other approaches such as Hocman [EJÖ02] focus on occasional very short inter-vehicle communications without fixed network access. Finally, in its support for disconnected operation, the Drive-thru Internet approach bears similarities to *Disruption/Delay-tolerant Networking (DTN)* [Fal03] and Drive-thru clients and proxies conceptually resemble DTN routers. The major difference is that DTN fundamentally assumes an asynchronous communications model (for newly developed applications) while we focus on *existing* applications and particularly need to embrace synchronous and interactive communications as much as possible.

⁶ In-between Drive-thru client and proxy, a *Performance Enhancing Proxy*—the Drive-thru PEP—in a hot-spot may be used to decouple link layer characteristics of the wireless network from those of the access link and backbone [BKG⁺01].

4 Real-World Hot-Spots in Drive-thru Environments

Our past measurements assumed an ideal hot-spot environment to validate basic operation: only a single, predefined SSID and static IP addresses were used, we did not perform any authentication with the access point, and we assumed that the Drive-thru proxy was located in the hot-spot so that the data did not need to pass through an access link that could become a bottleneck in communications. This section presents our investigations and findings when moving towards real-world hot-spot architectures: We have carried out experiments with different hot-spots in airports, train stations, and other locations in different cities. We have traced packet exchanges from entering a hot-spot to completion of the authentication process to measure contents and timing of packet exchanges, and we have collected login pages from different WISPs to analyze authentication forms. We have finally replicated parts of a hot-spot infrastructure (using the web-based authentication engine *NoCat*) in the lab and on the road.

4.1 Hot-spot Access

Accessing a network through a hot-spot incurs several steps: First of all, the mobile node needs to detect a radio carrier and determine the SSID(s) of the available WLANs (*association*). Next, DHCP is used for *IP auto-configuration*. Packet traces obtained in different commercial hot-spot environments have shown that DHCP completes in 0.1–2s if no retransmissions are needed (those may add up to 5s each), with subsequent ARP requests to ensure uniqueness of the assigned IP address taking another 0.5–2s. Thus, the total delay incurred by DHCP in static scenarios is usually less than five seconds. We have validated the applicability of DHCP in our Drive-thru environment by measurements on the road at 120km/h and have observed some 2–8s for DHCP.

Finally, the mobile node needs to authenticate with a hot-spot service provider (*authentication*) as described in section 2. For basic hot-spot architectures (a single WISP, UAM-based authentication mechanisms [ABS03]), we are able to perform automated authentication with rather simple means (see section 4.3 below). Packet traces of web-based authentication processes in commercial hot-spots have shown that DNS lookup, HTTP redirection, TLS setup, and login page retrieval completed fairly quickly as does the authorization once the user credentials were entered into the web form (less than 6s in total). Our experiments with an automated tool have confirmed these observations: access to a hot-spot is usually granted or denied in less than five seconds including DNS requests and redirection of the initial HTTP request (together usually less than 0.5s) as well as retrieval of the login page, form submission, and retrieval of the confirmation page. Altogether, automatic configuration and authentication using standard procedures may easily complete within 5–10s (i.e. during the entry phase) and hence leave all of the production phase to exchanging user data.

4.2 Hot-spot Communications

After completing authentication, the Drive-thru client on a mobile node starts exchanging data with its peer, i.e., the Drive-thru proxy in the fixed network. Without any access link constraints, the throughput characteristics follow the three phase model outlined in

section 3. To determine the impact of different access links on the communication characteristics in a Drive-thru environment, we have created the setup depicted in figure 3a): A mobile node (a laptop with two Ethernet interfaces) with a fixed node (another laptop with an Ethernet interface). The link simulator runs a software link-layer bridge (*satnat*) for which delay, data rate, and queue size can be configured independently in each transmission direction. For reference measurements without simulator the latter one is bypassed (dashed line in the figure). Mobile and fixed node both run the *tcpx* tool (short for *TCP eXchange* [OK04a]) developed by us that carries out a configurable data exchange pattern with its peer (here: sending or receiving only at the maximum achievable rate using 1460 byte segments).




Fig. 3. Measurement setup for investigating the impact of access links

We have chosen seven different settings and measured each with a fixed and a mobile sender: a reference setting with no access link constraints (*LAN*), dial-up access (*ISDN BRI*), a leased line equivalent (*ISDN PRI*), *DSL* access at two rates, and bi-directional satellite links without rate limits (*satellite 1*, to investigate the influence of pure delay) and at a low *DSL* rate (*satellite 2*, as available from DVB-RCS satellite service providers). The data rate limits (*uplink* and *downlink*) are defined from the hot-spot's viewpoint as indicated in figure 3. We have carried out two measurements for each direction on the road resulting in 28 measurements and one measurement each in the lab. These settings are listed on the left hand side of table 1: downlink and uplink limit refer to the data rate limit on the access link from and to the Internet, respectively; and RTT was measured in the lab on an uncongested simulated link using *ping*. The measurements on the road were taken at a speed of 120 km/h. The access point was connected to an external antenna mounted on a fence pole at about 2 m height, the WLAN card of the mobile node was connected to an omni-directional antenna on the car's rooftop. These tests were carried out between 10:00 and 14:00 on a weekday. Varying traffic conditions have influenced individual measurements (e.g., led to lower goodput in one transmission direction) but have not affected the general observations.

On the right hand side, the table shows a summary of the measurement results: *duration* indicates the average length of the connectivity period at the *tcpx* level, i.e., from connection setup to the last received segment; *volume* shows the average net data volume transferred in a single pass and *rate* the resulting effective net data rate for a single hot-spot. Finally, we define a degree of effectiveness (*effect.*) to denote the ratio of the result on the road compared to an ideal static lab setting in order to quantify the performance loss due to vehicle mobility.

<i>Parameters</i>					<i>Results</i>			
Link type	Sender	Downlink limit	Uplink limit	RTT	Duration	Volume	Avg. Rate	Effect.
LAN	fixed	—	—	1 ms	60s	59.1 MB	7.9 Mbit/s	54%
LAN	mobile	—	—	1 ms	64s	51.4 MB	6.4 Mbit/s	69%
ISDN BRI	fixed	64 kbit/s	64 kbit/s	30 ms	69s	498 KB	58 kbit/s	97%
ISDN BRI	mobile	64 kbit/s	64 kbit/s	30 ms	63s	447 KB	56 kbit/s	92%
ISDN PRI	fixed	1920 kbit/s	1920 kbit/s	20 ms	60s	11.1 MB	1.5 Mbit/s	80%
ISDN PRI	mobile	1920 kbit/s	1920 kbit/s	20 ms	62s	12.6 MB	1.6 Mbit/s	88%
DSL 1	fixed	768 kbit/s	128 kbit/s	50 ms	61s	5.3 MB	700 kbit/s	95%
DSL 1	mobile	768 kbit/s	128 kbit/s	50 ms	64s	821 KB	102 kbit/s	83%
DSL 2	fixed	4096 kbit/s	384 kbit/s	37 ms	68s	16.5 MB	1.95 Mbit/s	63%
DSL 2	mobile	4096 kbit/s	384 kbit/s	37 ms	59s	2.45 MB	335 kbit/s	91%
Satellite 1	fixed	—	—	515 ms	71s	1.66 MB	188 kbit/s	72%
Satellite 1	mobile	—	—	515 ms	63s	929 KB	118 kbit/s	89%
Satellite 2	fixed	768 kbit/s	128 kbit/s	515 ms	66s	1.68 MB	205 kbit/s	83%
Satellite 2	mobile	768 kbit/s	128 kbit/s	515 ms	65s	611 KB	76 kbit/s	67%

Table 1. Overview of measurement settings and results

As can be seen from the table, the duration of the TCP connection remains roughly constant at 60–70 s per pass (the connectivity island extends to some 2 km, see also figure 4). From the degree of effectiveness, we observe that mobility does not appear to have a negative impact on the overall performance for a given access link: despite the highly variable connectivity, the access link is kept filled most of the time and a moving vehicle will exploit on average 80% of its capacity. Only when the available bandwidth grows beyond several megabits per second (as in settings LAN and DSL 2)⁷ and makes up a significant fraction of the available WLAN capacity, the poor performance of the entry and exit phases gain weight. The reason can be seen from figure 4 showing a few representative plots: if the access link’s capacity amounts only to a small fraction of the WLAN’s itself, even the low performance in the entry and exit phases already suffice to fill the access pipe—which is the case almost immediately after the TCP connection is established. The temporary throughput drops also seen in the figure contribute further to the lower effectiveness; they are due to link layer retransmissions and packet losses caused by other vehicles blocking signals or causing interference. Nevertheless, such losses are recovered quickly as long as the RTT is sufficiently low—which holds for all scenarios except for satellites.

Several conclusions can be drawn from these observations: Today’s available terrestrial access technologies are suitable to provide Drive-thru Internet services in regular WLAN hot-spots. As long as queues in the access routers (both at ISPs and in the hot-spot) are kept short so that the RTT remains low, there is no need for performance enhancing proxies (PEPs) in the Drive-thru connectivity island: with terrestrial access networks, the variations of the wireless link’s communication characteristics can be handled “end-to-end”, i.e., between Drive-thru client and proxy. This is of particular importance because it implies that a regular hot-spot may be used “as is” and depen-

⁷ For the setting DSL 2 with a fixed sender, the maximum transmission rate was limited by the RTT and the small TCP default window size of 16 KB on the receiving machine so that even under ideal conditions only a maximum transmission rate of some 3 Mbit/s was possible.




Fig. 4. Measured data rate (fixed sender, 120 km/h)

dependencies on WISPs are avoided. Only for satellite links, it is advisable to use PEPs to decouple the WLAN from the satellite link characteristics and provide optimized repair strategies for the long latency link, particularly because of non-congestion induced packet losses on the access link (as further experiments have confirmed).

Finally, we have validated a complete hot-spot scenario using a setup as depicted in figure 3b)—which includes another Laptop with two Ethernet interfaces to act as an access router and an authentication server. Our measurements for the DSL 1 setting showed that DHCP and automated authentication always completed in 5s and that the subsequent data exchange yielded results similar to the above (duration 63s, volume 5.2 MB, avg. rate 658 kbit/s): the incurred goodput penalty appears smaller than the usual performance variation (of up to 1 MB). In summary, automatically accessing a UAM-compliant hot-spot is feasible and takes less than ten seconds and thus completes in the entry phase leaving the production phase available for actual data exchange.

4.3 Implementation

Based upon our findings, we have developed a number of software components for the mobile node. They support the Drive-thru client in its task to manage connectivity with the Drive-thru proxy in existing hot-spots. The *ConnectivityDetector* monitors the network interfaces in the mobile node to determine when link layer connectivity becomes available and when it disappears. It communicates information about available access points, their SSIDs, and the signal strength to all interested parties. An *enhanced DHCP client* is triggered by the *ConnectivityDetector*, performs auto-configuration, and signals the completion of the IP stack configuration to other components. The *AutoAuthenticator* is triggered by the *ConnectivityDetector* and the DHCP client and is responsible for obtaining network access. It attempts to create an HTTP connection to the Drive-thru Proxy in order to check for Internet connectivity. As soon as the retrieval succeeds, authentication is considered complete. Usually, the initial request is redirected to a local web page via HTTP/TLS, which is retrieved and analyzed to find the login submission

form following UAM conventions. In addition to the SSID and DHCP-assigned domain name, the authentication server certificate and the login page are used as hints to identify the WISP which is then matched against a local database. If found, the *AutoAuthenticator* submits the form with the corresponding credentials filled in. It checks the response and performs further actions (such as maintaining a connection for a logout page) to keep the access enabled before retrieving the initially requested resource. The *AutoAuthenticator* continuously sends local notifications about (changes to) the authentication status of a hot-spot.

The *Drive-thru client* is independent of the underlying connectivity establishment process. But it takes the above triggers about connectivity as hints when to attempt to set up a new connection to its Drive-thru Proxy and when to wrap up communications for a connectivity island. If a WLAN link becomes available it initiates/resumes communications right away (rather than waiting for a particular SNR threshold to be reached) and continues until connectivity is lost again.

Our implementation is currently tailored to work with the most common hot-spot scenarios. Laboratory tests of the individual software components, first non-mobile tests of the *AutoAuthenticator* with real-world hot-spots, and field tests using our experimental hot-spot setup have been successful. Nevertheless, various issues remain to be addressed: enhancing the *AutoAuthenticator* to deal with more complex login pages, supporting selected WEP-protected networks and integrating 802.1x-based authentication. Two further open issues are user policies defining when to use which hot-spots and which WISP and, as a prerequisite, to unambiguously (and quickly) identify which WISPs are available in a given hot-spot—these are subject to our current research.

4.4 Desirable hot-spot properties

From the above discussions, we can devise a set of desirable properties to be considered for future hot-spot installations. At the WLAN link layer, dedicated access points with antennae mounted outdoor, e.g. on the roof of a building, are needed to offer sufficient connectivity. Directional antennae may be used pointed towards the road, and separate access points with non-overlapping channels are recommended to serve stationary users without interference. At the IP layer, traffic of stationary and Drive-thru users may be separated on the access link, e.g., by means of diffserv classes (or even different access links) to ensure a guaranteed bandwidth share for passing users. In most cases, it should be easy to satisfy these requirements and “upgrade” an existing hot-spot by simply adding dedicated access points and traffic management functions.

Furthermore, hot-spot service identification and user authentication aspects are crucial elements for successful deployment. A Drive-thru node must be able to automatically determine 1) whether or not a hot-spot is meant for public access; 2) which services are offered⁸; and 3) from which WISPs Internet access is available at what cost. While our *AutoAuthenticator* shows that parsing HTML pages is workable for selected hot-spot operators, this short-term solution may fail if web pages do not follow the UAM conventions or contain JavaScript code, if multiple WISPs are offered via

⁸ Such services might include *Internet access* but also local services, e.g., availability of a Drive-thru PEP, download of local resources, advertisements, pre-selected content, etc.

a single web page, or if the initial redirect does not lead to the login page. However, other authentication mechanisms may need to be supported as well which ultimately requires implementing a trial-and-error decision chain to determine the access method and WISP—which takes precious time to complete and is error-prone. It is preferable to have the hot-spot operator disseminate information about available services, service providers, tariff models, and access methods using a standardized protocol and service description language. The Wi-Fi Alliance suggests a *smart client* authentication protocol that uses service announcements [ABS03]; other approaches to optimized service announcements are in discussed, e.g., in [BWSF03] and [KO03].

Finally, tariff models for WLAN access and accounting practices need to match the short hot-spot access periods, even special terms for short access durations (e.g., less than 120s) could be offered, and, of course, flat-rate service plans. Roaming agreements between WISPs would improve accessibility of hot-spots for individual users.

5 Conclusions

Drive-thru Internet introduces a new paradigm for mobility support and enhances the use of WLAN infrastructure previously restricted to stationary users. While initial measurements of WLAN performance in mobile scenarios have already shown the potential of the Drive-thru approach, in this paper, we have validated that our approach is workable with real-world hot-spot installations. We have proven that autoconfiguration and even “manual” authentication procedures can be completed in an automated fashion within a few seconds. Our measurements have shown that data rate adaptation to the changing WLAN characteristics is achievable end-to-end across common terrestrial access links. In summary, today’s existing hot-spots are, in principle, usable without fundamental modification. But we have also identified areas for improvements: of particular interest are service announcements that allow a mobile node to unambiguously determine the available WISPs, authentication methods, and tariff options—and to allow a hot-spot operator to provide this very information to attract customers. Furthermore, most current (time-budget-based) tariff models are suboptimal for the usually short Drive-thru access periods, and WLAN roaming is still not generally available.

On the hot-spot side, our future work includes providing service announcements, considerations for Drive-thru PEPs beyond plain connection splitting, and proper traffic differentiation. For the mobile node, we are advancing our *AutoAuthenticator* to understand service announcements, allow for policy-based WISP selection, and support additional authentication methods to achieve widespread applicability. Future measurements will include mobile user and background traffic via access links modeled following web and emails access traffic patterns. Ultimately, simulations are needed to assess the overall system behavior with many hot-spots and mobile users.

As the pervasiveness of 802.11 is still increasing—meanwhile the technology is considered as a low-cost area-wide network infrastructure for cities [Mooc04]—it is likely that the mobile usage scenario we have developed in the Drive-thru Internet project will gain popularity. With support for automated hot-spot association and disconnection-tolerant networking at hand, WLAN hot-spots could easily evolve into connectivity oases for stationary *and* mobile users.

References

- [ABS03] B. Anton, B. Bullock, and J. Short. Best Current Practices for Wireless Internet Service Provider (WISP) Roaming, Version 1.0. Wi-Fi Alliance, February 2003.
- [ACDS03] Ahmad A, R. Chandler, A. A. Dharmadhikari, and U. Sengupta. SIM-Based WLAN Authentication for Open Platforms. *Technology@Intel Magazine*, 2003.
- [BFW03] Marc Bechler, Walter J. Franz, and Lars Wolf. Mobile Internet Access in FleetNet. In *13. Fachtagung Kommunikation in verteilten Systemen, Leipzig*, 2003.
- [BKG⁺01] J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby. Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations. RFC 3135, June 2001.
- [BVB03] A. Balachandran, G. M. Voelker, and P. Bahl. Wireless Hotspots: Current Challenges and Future Directions. In *Proceeding of WMASH'03*, 2003.
- [BWSF03] Marc Bechler, Lars Wolf, Oliver Storz, and Walter J. Franz. Efficient Discovery of Internet Gateways in Future Vehicular Communication Systems. In *Proceedings of the 57th IEEE VTC 2003 Conference*, Jeju, Korea, 2003.
- [EJÖ02] Mattias Esbjörnsson, Oskar Juhlin, and Mattias Östergren. The Hocman Prototype - Fast Motor Bikers and Ad-hoc Networking. Proceedings of MUM, 2002.
- [Fal03] Kevin Fall. A Delay-Tolerant Network Architecture for Challenged Internets. Proceedings of ACM SIGCOMM 2003, 2003.
- [HKR⁺03] John Hammond, Bart Kessler, Juan Rivero, Chad Skinner, and Tim Sweeney. Wireless Hotspot Deployment Guide. Technical report, Intel, December 2003.
- [KO03] Dirk Kutscher and Jörg Ott. Dynamic Device Access for Mobile Users. In *Proceedings of the 8th Conference on Personal Wireless Communications*, 2003.
- [Lei01] Gosta Leijonhufvud. Multi access networks and Always Best Connected, ABC. November 2001. MMC Workshop.
- [Mooc04] Philadelphia Mayor's office of communications. Mayor John F. Street Announces Appointment Of Wireless Philadelphia Executive Committee. Press release, 2004.
- [OK04a] Jörg Ott and Dirk Kutscher. Drive-thru Internet: IEEE 802.11b for „Automobile“ Users. In *Proceedings of the IEEE Infocom 2004 Conference, Hong Kong*, 2004.
- [OK04b] Jörg Ott and Dirk Kutscher. The “Drive-thru” Architecture: WLAN-based Internet Access on the Road. In *Proceedings of the IEEE Semiannual Vehicular Technology Conference May 2004, Milan*, May 2004.
- [OK04c] Jörg Ott and Dirk Kutscher. Why Seamless? Towards Exploiting WLAN-based Intermittent Connectivity on the Road. In *Proceedings of the TERENA Networking Conference, TNC 2004, Rhodes*, June 2004.
- [OK05] Jörg Ott and Dirk Kutscher. A Disconnection-Tolerant Transport for Drive-thru Internet Environments. In *Proceedings of the IEEE Infocom 2005 Conference*, 2005.
- [RCC⁺04] Pablo Rodriguez, Rajiv Chakravorty, Julian Chesterfield, Ian Pratty, and Suman Banerjee. MAR: A Commuter Router Infrastructure for the Mobile Internet. In *Proceedings of the ACM Mobile Systems, Applications and Services Conference (ACM Mobisys 2004)*, June 2004.
- [TMC04] TMCnet.com. Clic TGV Brings Wifi Onboard France's High Speed Trains. available online at <http://www.tmcnet.com/submit/2004/Jan/1022655.htm>, January 2004.
- [Wir03] Boingo Wireless. Toward Ubiquitous Wireless Broadband. White Paper, 2003.
- [ZWS⁺03] M. Zitterbart, K. Weniger, O. Stanze, S. Aust, M. Frank, M. Gerharz, R. Gloger, C. Görg, I. Gruber, S. Hischke, P. James, H. Li, C. Pampu, C. de Waal, W. Weiß, D. Westhoff, J. Wu, D. Yu, and X. Xu. IPonAir – Drahtloses Internet des nächsten Generation. PIK, Vol 26, No 4, October 2003.