A Pegasus over WiFi and WiMax Demo: Connectivity at high speeds

Nikolaos Frangiadakis (designated corresponding author), Nick Roussopoulos
{ntg,nick}@cs.umd.edu
Department of Computer Science,
University of Maryland,
College Park, Maryland, USA

***Abstract*—We present Pegasus, a system providing wireless connection roaming at high velocities over multiple interfaces. Our system operates over WiMax, as well as over "in situ" WiFi networks, while at the same time offering transparency to user level applications by allowing a single IP address per user, operating in a decentralized mode. One of our most important findings is that in a system where handoffs delay minimization is crucial, network information for user locations and used paths can be used for effective and balanced utilization of the available bandwidth. We exploit this by caching DHCP connections, which we store on the server, as well as selecting superior Access Points (AP) to use when clients handover. We use the our working testbed implementation on a simulation of mobility by cycling the client's among APs, and derive the resulting bandwidth for large file downloading and HTTP connections.**

I. Introduction

We will soon be surrounded by ubiquitous wireless networks and equipped with devices able to use possibly more than one of them. Protocols and systems should be generic and efficient enough to support a vast array of applications remaining perpetually connected as users move. A major challenge is the proper usage of all available mediums to enhance the provided services in an effective, seamless and non-disruptive way.

Measurements and ongoing research have shown that WLAN connection for moving vehicles is feasible. However none of the previous works suggests a solution addressing a complete array of the challenges in vehicular WLAN communications. In the same time, technologies like WiMax will further enhance future communications between mobile users. Our system, Pegasus amends this by providing wireless connection roaming at high velocities. To the best of our knowledge, it is the first system that operates on WiMax and over in situ WiFi networks, while at the same time offers transparency to user level applications by allowing a single IP address per user, and does not impose additional requirements to existing infrastructures. Pegasus offers simple deployment, improved scalability, and is the first able to operate over secure in situ networks. Furthermore, it remains efficient under intermittent connectivity conditions and supports heterogeneous network mediums for increased robustness.

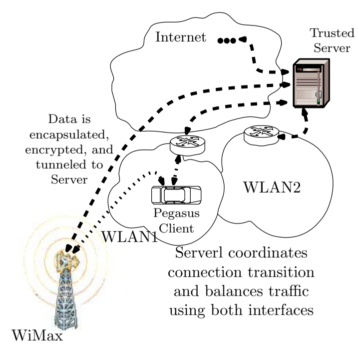

Fig. 1. Pegasus system architecture

II. Pegasus

A. System Architecture

Figure 1 shows an overview of Pegasus. It is a system built to support multiple interfaces for rapidly moving devices. It can sustain fast 802.11 access point connection switches. It transparently switches between 802.11 access points, and connections utilizing multiple mediums, such as WiFi and WiMax. It presents a constant IP and persistent connectivity appearance to users. Clients are supported by Pegasus Servers (PegSvc), which manage all connections. Each server may host some or all of the modules required for connection managing, location management, as well as authentication, authorization and auditing functions. User traffic is tunneled through PegSvc. There is no additional requirement, or modification for any client applications, "in situ" APs, and we do not modify any of the Protocols involved. The client just installs a driver that creates a new interface.

Therefore, in Pegasus, every WLAN is independently managed, so we deal with different ISPs, private address spaces and NATs. To handle such heterogeneity, client's traffic is tunneled to the server, which can be operated by a third party and acts as a multiplex point for all client Internet communications. PegSvc attempts to predict the client movement through deployed WLANs and offers choices for the next access point connection. The switch from one AP to another will not sever the ongoing client application sessions. Moreover, since the server acts as fixed peer to the non-mobile connection endpoints, it buffers network packets to smooth possible connectivity dead spots. All of the tunneled traffic is encrypted to offer extra security for the client data.

B. Implementation

Our implementation makes use of the Linux standard driver for wireless devices, thus it is compatible with any hardware supported by Wireless Tools by Tourrilhes. One of the new ideas rests with caching DHCP connections on the server side and reusing them to allow for fast handovers between APs. Our interface can split traffic and distribute them in parallel to multiple interfaces thus combining the bandwidth and profiting from the existence of multiple wireless networks at minimal cost.

We have simulated a mobile environment in our testbed. Pegasus server is a Pentium III with Ubuntu Linux deployed in public domain. For "in situ" WLANs, we installed Linksys 54g access points that are available in any store configured with default factory settings. Every access point runs a firewall, NAT, and DHCP for its private network. For the open access points the DHCP cache on the server is populated dynamically by clients that connect to every WLAN and send the DHCP connections to server. The client used for measurements is a regular laptop running Linux with ipw3945 Intel wireless card which is a standard for Dell laptops. To simulate movement, the client switches its wireless connection from one access point to the next and the connectivity period to each WLAN depends on the simulated driving velocity.

To benchmark performance we use TTCP tests and a web browsing session. The TTCP application runs unaware of the ongoing physical connection transitions and measures end-to-end TCP bandwidth. In order to emulate different application behaviors we test with several TTCP configurations; we simulate large continuous data transfers, and multiple smaller data transfers. For the web browsing sessions we record response times and the number of times a web page comes back with status "404 Page Not Found".

III. Related work

[1] studies the performance of TCP and UDP in wireless network scenarios from immobile clients. The Drive-thru Internet project by Ott and Kutscher [2] studied the behavior of network connections over 802.11b and 802.11g from a moving car and classified WLAN connection period to "entry", "production", and "exit" stages. In their more recent work [3], they show that they can avoid TCP start up overheads by using proxies, and hiding short period of disconnection from the transport layer. In Pegasus, we do not modify TCP, but provide constant connectivity appearance to the client. Gass et al [4] used off-the-shelf 802.11b wireless equipment between the client and the AP and concluded that packet losses are low within 150 meters of the access point for a wide speed range (5-75 mph).

I-TCP [5] is a split connection approach that introduces a transport layer intermediary for splitting a TCP connection between a fixed and a mobile host into two connections to mitigate disruptive effects of handovers. The Snoop protocol [6] provides a more transparent support layer, and relies on a dedicated agent that "snoops" on the TCP communication on the path between the mobile and fixed station, buffers and retransmits TCP segments. CAMA [7] and Mobile Router [8] explored using multiple wireless mediums. The first utilized cellular communications for control messaging purposes, while Mobile Router concentrated on allowing different client types to connect to a common router on a commuter bus.

Already a number of cities have plans to cover most of their area or large parts with hot-spots. This would create islands of connectivity Pegasus could take advantage of. It also is possible that in the near future users will decide to join private hot-spots enhancing each other communication capabilities [9].

As IEEE Wi-Fi evolve, handoff cost is taken seriously into account and recent standards [10] aim for Fast Roaming / Fast BSS Transition. So far most of this work is related to the physical and MAC layers. However, when we consider the added cost of higher layers and security, we understand that for short range handoffs additional efficiency is vital. Thus, the problem of optimal Access Point selection becomes crucial.

IV. Operation

A. Fast WiFi Handovers

We achieve efficiency by reusing a DHCP cache globally. Pegasus caches all of the DCHP connections from all of the clients in a global cache, and continuously reuses them. Since DHCP is bound to the client MAC address, users change their MAC address to a value handed-in by the PegSvc in order to reuse DHCP connections. Once a user moves on to the next connection, he changes the MAC address again. This concept of recycling acquired DHCP connections and using a different DHCP identity at each independent access point island is the core concept that allows Pegasus to achieve its efficiency and scalability.

In fact, Figure 2 displays the client TCP performance for continuous transfers in the absence of WiMax. We see that Pegasus can allow us to maintain connections if a WiMax connection is disrupted for some period of time in a city. We also see that Pegasus overhead is insignificant. Note that we used an old Pentium III with Ubuntu Linux deployed in public domain. Our implementation is based on the Click Modular Router.

Figure 3 demonstrates the performance of a Pegasus for short http transfers. Transfer rates for the shorter segments vary because the transfers are more susceptible to connection transitions. Some of the segments do not experience transitions at all. Nevertheless, PegSvc with no DHCP clearly illustrates the most stable behavior where all of the data is eventually delivered.

B. WiMax Interface

While WiMAX should deliver 70 Mbit/s and operate in ranges of 31 miles/50 kilometers, it can only do one or the

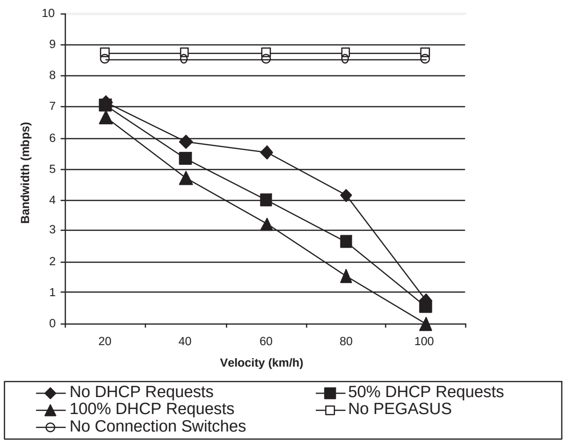

Fig. 2. Client TCP performance for continuous transfers

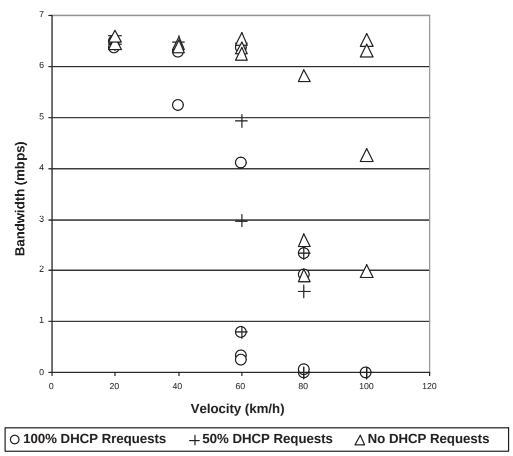

Fig. 3. Client TCP performance for short transfers

other. In urban environments users may not have line-of-sight and therefore receive 10 Mbit/s over 2 km [11]. This bandwidth would have to be shared between the users and allow for varying upload and downloading rates.

C. WiFi AP selection

We have very strong indications that utilizing both interfaces can allow our system to greatly enhance the performance its WiFi part. In fact, we should expect that not all APs available to a device at a time will be equally effective. In fact, even if the direction in which a user moves is not considered, 10% of the APs will be responsible for 90% of the provided bandwidth. By suggesting APs via WiMax we can increase the effectiveness of WiFi connections and balance the load to APs.

In our current Implementation, Pegasus makes use of WiMax to tunnel traffic through it. However, Pegasus could become more integrated to WiMax by running as a part of it. In this case, connections would be handled by a different radio interface depending on location and capacity. Such architecture should be expected to be common in future networks[12]. In our case it would signify using IPv6 ontop of our system to handle handovers [13].

V. Conclusion

Pegasus is a system that provides wireless connection roaming at high velocities over multiple interfaces. Our implementation offers connections over “in situ” WiFi networks, as well as WiMax. It does not require any modification of AP settings or the TCP protocol. We implemented Pegasus over the standard linux WiFi library so that it can be used with most available WiFi devices. By making use of the techniques described in this paper, we have significantly reduced times required for handover in the WiFi domain and integrated the traffic from WiMax, making possible to implement and demonstrate an efficient system that seamlessly connects over the two media.

References

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, “Link-level measurements from an 802.11b mesh network,” in SIGCOMM ’04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications. New York, NY, USA: ACM Press, 2004, pp. 121–132.

[2] J. Ott and D. Kutscher, “Drive-thru internet: Ieee 802.11b for ”automobile” users,” in INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies, vol. 1, 2004.

[3] J. ”Ott and D. . Kutscher, “A disconnection-tolerant transport for drive-thru internet environments.” in INFOCOM. IEEE, 2005, pp. 1849–1862.

[4] R. Gass, J. Scott, and C. Diot, “Measurements of in-motion 802.11 networking,” in WMCSA ’06: Proceedings of the Seventh IEEE Workshop on Mobile Computing Systems & Applications. Washington, DC, USA: IEEE Computer Society, 2006, pp. 69–74.

[5] A. Bakre and B. R. Badrinath, “I-tcp: indirect tcp for mobile hosts,” in Proceedings - International Conference on Distributed Computing Systems. Vancouver, Can: IEEE, Piscataway, NJ, USA, 1995, p. 136. [Online]. Available: citeseer.ist.psu.edu/bakre95itcp.html

[6] H. Balakrishnan, S. Seshan, E. Amir, and R. H. Katz, “Improving tcp/ip performance over wireless networks,” Proceedings of the Annual International Conference on Mobile Computing and Networking, MOBICOM, pp. 2–11, 1995.

[7] B. Bhargava, X. Wu, Y. Lu, and W. Wang, “Integrating heterogeneous wireless technologies: a cellular aided mobile ad hoc network (cama),” Mob. Netw. Appl., vol. 9, no. 4, pp. 393–408, 2004.

[8] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee, “Mar: a commuter router infrastructure for the mobile internet,” in MobiSys ’04: Proceedings of the 2nd international conference on Mobile systems, applications, and services. New York, NY, USA: ACM Press, 2004, pp. 217–230.

[9] E. Efstathiou and G. Polyzos, “Self-organized peering of wireless lan hotspots,” in European Transactions on Telecomunications, vol. 16, no. 5 (Special Issue on Self-Organization in Mobile Networking), Sept/Oct 2005. [Online]. Available: http://mm.aueb.gr/research/P2PWNC/publications/papers/ett2005.pdf

[10] “IEEE std 802.11r-2008 (amendment to IEEE std 802.11-2007 as amended by ieee std 802.11k-2008),” pp. c1–108, 15 2008.

[11] A. Yarali, S. Rahman, and B. Mbula, “Wimax: The innovative broadband wireless access technology,” JCM, vol. 3, no. 2, pp. 53–63, 2008.

[12] G. Lampropoulos, A. Salkintzis, and N. Passas, “Media-independent handover for seamless service provision in heterogeneous networks,” Communications Magazine, IEEE, vol. 46, no. 1, pp. 64–71, January 2008.

[13] J. Yuan, C. Bo, and Z. Chen, “An improved fhmipv6 handover scheme for mobile wimax,” Wireless Communications, Networking and Mobile Computing, 2008. WiCOM ’08. 4th International Conference on, pp. 1–4, Oct. 2008.