Performance-Oriented Handover Incentives: A Proof-of-Concept Demonstration*

Xenofon Fafoutis
fontas@ics.forth.gr

Kostas Mathioudakis
kostasm@ics.forth.gr

Vasilios A. Siris†
vsiris@ics.forth.gr

Institute of Computer Science
Foundation for Research and Technology - Hellas
P.O. Box 1385, GR 711 10, Heraklion, Crete, Greece

ABSTRACT
In competitive environments, where parties act in their own self-interest, there are cases where access points can improve their own performance by serving clients of neighboring access points that belong to a different self-interested party. As a result, handover incentives arise due solely to the performance improvements that each party can achieve. We have implemented a procedure that allows an access point to identify when such incentives arise and if so, perform the handover. The demo illustrates the operation of this procedure, and the corresponding performance gains achieved with the handover.

1. BACKGROUND
In IEEE 802.11 networks, the concurrent operation of low and high transmission rate users at the same channel significantly degrades the performance of the high-rate users. In our previous work [1][2], we have identified and investigated cases where self-interested IEEE 802.11 WLANs have incentives to serve neighboring clients, since this handover mitigates the aforementioned degradation, yielding performance improvements to both self-interested parties. Moreover, we proposed an analytical model that can estimate the throughput of each client in different contexts. This model can be used to predict when the performance-oriented handover incentives arise, i.e., when handovers are expected to improve the throughput of *both parties*; this is a win-win situation for both involved parties. The prerequisite for the handover incentives to arise is that two or more access points operate in the same channel and in the same contention area. Indeed, it is common that the available channels are not sufficient to assign orthogonal channels to different access points, especially in dense residential areas. Moreover, non-overlapping channels will be further reduced as more wireless networks operating in unlicensed bands are deployed over time.

We stress that the key difference to prior work is the focus on competitive parties that act in their own self-interest and there is no other cooperation between them, such as monetary exchange or other forms of enforcement. The incentives for cooperation arise due solely to the performance improvements each self-interested parties obtains from cooperating to perform handovers. In contrast, prior work - on one hand - aims to improve the aggregate performance when all network devices work towards a common goal, thus assuming and requiring altruistic behavior of all parties, and - on the other hand - focuses on ways to enforce cooperation.

2. HANDOVER DECISION PROCEDURE
The handover decision procedure for each access point is depicted in Figure 1. The decision procedure includes three phases: the monitoring phase, the handover decision phase and the handover execution phase. In the first phase, access points perform passive measurements, to identify low transmission rate stations, and to obtain the measurements required by the handover decision module; this information includes the number and transmission rate of the clients served by the neighboring access points. Assuming that there are no hidden nodes, this information can be obtained by counting unique MAC addresses and extract their rates from the PLCP header. Additionally, the access point needs to estimate the rate that the low rate clients would operate at, if handovers are performed. Measuring the received signal strength of the frames the low rate nodes transmit can be used for this estimation. If the handover decision module indicates that a handover is expected to be beneficial, then the access point continues with the execution of the handover.

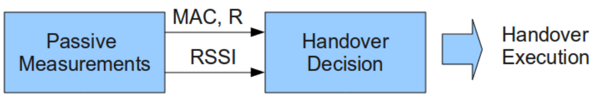

Figure 1: Handover Decision Procedure.

3. DEMONSTRATION SCENARIO
Next we describe the demonstration scenario. The scenario is able to demonstrate the performance benefits of exploiting the handover incentives, and at the same time has the smallest infrastructure requirements.

3.1 Topology and Traffic
Consider the case of two access points, APL and APH, which belong to different self-interested owners, Figure 2. APL sends UDP traffic under saturation conditions to node CLX at a low transmission rate. Node CLX is a client of APL; hence the actions of APL target to improve its throughput. APH sends UDP traffic under saturation conditions to node CL at a high rate. However, CLX is closer

*This work was supported in part by the EC in the 7th Framework Program through project EU-MESH, ICT-215320, `http://www.eu-mesh.eu`

†Vasilios A. Siris is also with Athens University of Economics and Business

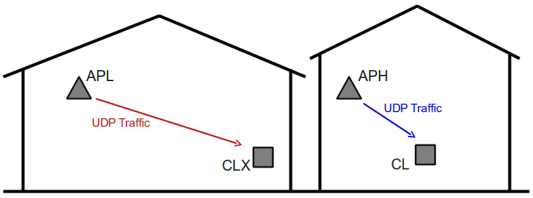

Figure 2: The initial case.

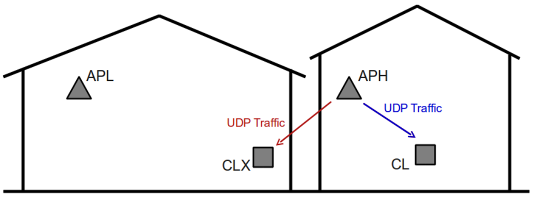

Figure 3: The handover case.

to APH and would be able to operate at a higher rate, if it was associated to it. Additionally, we assume that both access points operate at the same channel and all access points and nodes are in the same contention area. Because of the low-rate transmissions to CLX, the performance of CL also degrades. The handover decision procedure can identify when both parties have performance-oriented incentives for allowing handovers of the low-rate client (CLX) of APL, to the neighboring access point APH. In this case, Figure 3, APH sends traffic at a high rate to both its own client (CL) and the ex-low-rate client of APL (CLX).

3.2 Implementation

The demonstration is implemented on two multi-radio nodes operating as access points. Each node contains two mini-PCI wireless interfaces (WLM54AG, 802.11a/b/g). Each wireless interface is connected to a Triband APXtender 5 GHz, 2.2 dBi indoor antenna. Finally, each mesh node contains a 10/100 Ethernet interface. The nodes run Ubuntu OS, with Linux kernel 2.6.24-19. The wireless device driver is MadWiFi (version r4100), and the driver configuration is performed using the Wireless Extensions API.

All nodes are connected to a wired network, over which we trace the progress of the demonstration and update accordingly a graphical web interface. In our current implementation, both APs operate on a single ESSID, although this is not restrictive. We use the driver's access lists to ensure that each client is assigned to the access point determined by the handover decision procedure. The scenario begins by starting two UDP flows using the *iperf* traffic generator. Every few seconds, the clients report their received throughput to the web interface via the Ethernet interface. For the monitoring phase, we use the secondary wireless adapter of our access points, which is set to monitoring mode, and operates on the same channel as the other interface. Alternatively, a virtual interface can be used, thus avoiding the need for two wireless interfaces. The traffic analysis tool *tshark* is used to extract the information required for the handover decision, which includes the MAC address and rate of the neighboring AP's client, and the received signal strength of its ACK frames; the received signal strength is used to estimate the transmission rate if the client (CLX) associates with the specific access point (APH). The above information is passed to the handover decision procedure, which estimates whether the handover is expected to be beneficial. If the answer is positive for both self-interested access points, then the handover is executed. Execution of the handover requires direct communication of the access points. For the handover execution phase, we use the Ethernet channel, through which APH sends a message to APL requesting the handover of CLX. Then, both access points update their access lists accordingly, so that when CLX gets disassociated from APL, it can associate to APH. The throughput improvements are shown in the web interface. The steps described above are shown in Figure 4.

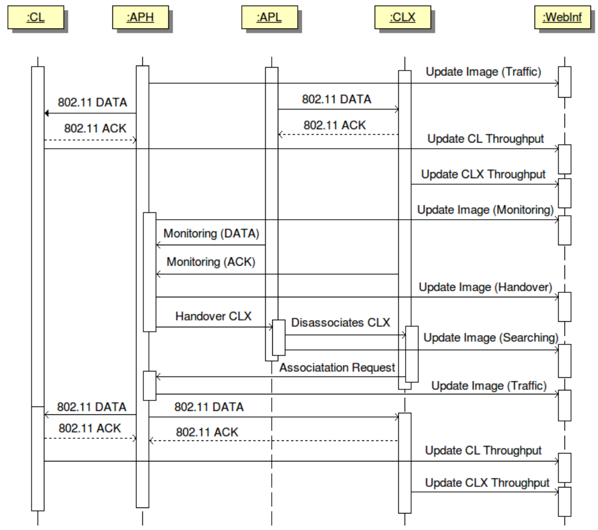

Figure 4: Sequence diagram of the demonstration.

3.3 Results of a demonstration run

As an example, we show the throughput measurements of each self-interested access point operating at 802.11b, for the case where CLX transmits at 1 Mbps before the handover, and after the handover is served by APH at 11 Mbps. Figure 5 shows the throughput before and after the handover, as it appears in the demonstrations's web interface. Observe that that APH improves the performance of its own client CL, by serving the neighboring client CLX, which belongs to a different self-interested party (APL), while at the same time the performance of CLX also improves.

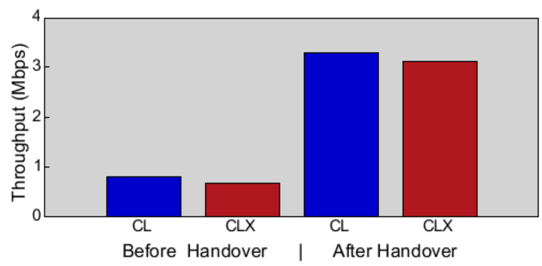

Figure 5: Demonstration results: throughput before and after the handover.

4. REFERENCES

[1] X. Fafoutis and V. A. Siris. Handover Incentives for WLANs with Overlapping Coverage. In *Proc. 7th Int'l Conf. Wired / Wireless Internet Communications (WWIC)*, Enschede, Netherlands, May 2009

[2] X. Fafoutis and V. A. Siris. Handover Incentives: Revised Model with extensions for Uplink Traffic. In *1st Int'l Workshop Mobility in Future Internet (MiFI)*, Chania, Greece, June 2010

APPENDIX

A. DEMO REQUIREMENTS

The demonstration infrastructure is comprised of two multi-radio nodes operating as access points and two laptops operating as clients. Additionally, the background tracing traffic requires a switch for wired connectivity. These are the minimum infrastructure requirements. However, the demonstration should be ideally shown on a large monitor. As a result, the space requirements are minimal. Lastly, the demonstration requires the exclusive use of a wireless channel, preferably at the band of 2.4 GHz, without any adjacent channel interference.