

Energy-efficient Wi-Fi Gateways for Federated Residential Networks

Original

Energy-efficient Wi-Fi Gateways for Federated Residential Networks / Rossi, Claudio; Borgiattino, Carlo; Casetti, CLAUDIO ETTORE; Chiasserini, Carla Fabiana. - STAMPA. - (2013), pp. 1-3. (IEEE WoWMoM 2013 Demo Session Madrid (Spain) June 2013) [10.1109/WoWMoM.2013.6583364].

Availability:

This version is available at: 11583/2507385 since:

Publisher:

IEEE / Institute of Electrical and Electronics Engineers Incorporated:445 Hoes Lane:Piscataway, NJ 08854:

Published

DOI:10.1109/WoWMoM.2013.6583364

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Energy-efficient Wi-Fi Gateways for Federated Residential Networks

Claudio Rossi, Carlo Borgiattino, Claudio Casetti, Carla-Fabiana Chiasserini
Department of Electronics and Telecommunications, Politecnico di Torino
Email: {carlo.borgiattino,claudio.rossi,claudio.casetti,chiasserini}@polito.it

Abstract—Cooperation among federated APs in dense urban areas can yield energy saving by allowing under-used devices to hand over their wireless stations (WS) to nearby APs and temporarily switch off while meeting user expectations in terms of throughput. We demonstrate the effectiveness and the benefits of our energy-efficient cooperative protocol through a real deployment emulating a residential scenario. The demo we propose is highly interactive, as users can generate traffic within a BSS through a wireless station, like a smartphone or a notebook, and observe, through a web interface, the protocol behavior and the network topology changes caused by the new traffic scenario.

I. INTRODUCTION AND NETWORK ARCHITECTURE

Today ICT is accountable for 2-4% of the worldwide carbon emissions and this number is projected to double by 2020 [1]. Telecoms infrastructure and devices account for 34% of the total ICT consumption, and the 95% of this share is due to home and access networks [2]. Within the EU FP7 FIGARO project [3], we address the emerging need for green technologies by proposing a cooperation scheme among “federated” residential Gateways (GW). A Federation is a logical overlay relationship among trusted home GWs with the purpose of content exchange and resource sharing. Federated GWs can communicate and coordinate with each other using an out-of-band channel, which runs through their backhaul Internet connection. Each GW offers local wireless access acting as Access Point (AP) and creating a 802.11 a/b/g/n BSS (Basic Service Set) over independently-managed (but possibly coordinated) frequency channels.

The cooperation scheme enables GWs within radio range of each other to manage their associated devices, or Wireless Stations (WS): (i) under low traffic conditions, a GW can hand over the associated WSs to nearby GWs and temporarily switch off; (ii) under high traffic conditions, a GW can selectively hand over one WS at a time in order to lower the congestion. When GWs are “off”, they no longer have wired/wireless connectivity and only run a low-cost, low-power radio interface, e.g., a IEEE 802.15.4 card, that can be used as wake-on WLAN interface [4].

II. LOAD ESTIMATION AND THE ASSOCIATION MANAGEMENT PROTOCOL

The implementation of the energy-efficient cooperation scheme for federated GWs consists of three main steps.

The first one is the estimation of the current load of the wireless channel and of its saturation throughput, through

passive traffic measurements as detailed in [5]. By comparing the current load to the saturation throughput S , a GW classifies its status as either *Light*, *Regular*, or *Heavy*. In the *Light* status, traffic likely comes from background communications to/from the WSs, prompting the GW to try and relocate them, switch itself off and save energy. The *Heavy* status, instead, characterizes an overloaded BSS, where some WSs should associate to other BSSs to benefit from load balancing. In the *Heavy* status a GW tries to relocate one WS at a time, starting from the one having the lowest bit rate, until the *Regular* status is reached. A GW in *Regular* status is too busy to switch itself off while it does not need to be relieved of some of its WSs. It might, however, accommodate relocated WSs within its BSS.

Passive traffic measurements account for the second step too, in which a GW trying to relocate one or more of its WS compiles a traffic profile of each WS, detailing the throughput of its active (downlink/uplink) traffic flows.

The third and final step amounts to an inter-GW communication in which a *requesting* GW, aiming at relocating one or more of its WSs, sends a handover request, along with the traffic profiles of relocatable WSs, to *candidate* GWs. A computation of the projected load (current estimated load and expected load from the incoming WS profile), and its comparison against the estimated saturation throughput S , allow a candidate GW to assess its suitability to provide help to others (i.e., if the additional WS does not plunge the GW in *Heavy* status). Responses returned by candidate GWs let the requesting GW identify a feasible relocation strategy. Among the feasible solutions that allow a GW to relocate its WSs, the allocation maximizing the average data rate of the WSs is selected. For a *Heavy* initiated request, if no viable relocation is found the requesting GW needs to wake up a neighboring “off” GW and repeat the requesting procedure. Further details about the protocol can be found in [5].

III. TESTBED ARCHITECTURE

As shown in Figure 1, the testbed is a federated network composed by three GWs, six WSs, a web server and a router. Each GW acts as AP of a 802.11g network operating on a different channel in the 2.4 GHz band, and secured with WPA-PSK, which is part of the 802.11i standard and it is widely adopted in residential networks. Every WS knows all access keys in the federated network, hence it can associate



Figure 1. Testbed architecture.

to any GW. We included four laptops and two smart-phones as WSs, so as to create the heterogeneity of a real-life residential environment.

When “on”, GWs notify the web server about the WSs associated to them, along with the WSs traffic profile. The web server has two functions: (i) it provides contents to WSs and (ii) it graphically shows the testbed status over time in a web page. Specifically, such testbed monitoring interface shows the actual association of the WSs, the GWs current load (L), the aggregated throughput, the saturation throughput (S), and L/S , which is used to determine the GWs status. On top it displays the total energy consumption within the federation and, when the protocol is started, the total energy saving with respect to the case where all GWs are always “on”. The layout of the monitoring interface is given in Figure 2.

A. Hardware and software description

GWs feature an Alix PC Engines motherboard, equipped with an AMD Geode 500 MHz processor, one IEEE 802.11 b/g compliant Wistron DCMA-82 Atheros wireless card and one omnidirectional antenna with a gain of 5 dBi. Each Alix runs OpenWrt Backfire, a Linux distribution for embedded devices, while the WSs can run any operating system.




Figure 2. Testbed monitoring interface.

The passive traffic measurements needed by the protocol are implemented *on the GW* within the mac80211 module of the Linux wireless driver *compact-wireless 2011-21-01* [6]. Note that, since we modified only the mac80211 module and not the GW hardware, such measurements work on any device. Specifically, at the BSS level we keep track of: average size of the packet payload, average SNR, average data rate, average packet error rate (PER), number of associated WSs. For each WS, at the MAC layer we measure: all throughput contributions (elastic/inelastic, downlink/uplink traffic) with the corresponding number of handled packets, the average packet payload, the average data rate, the average SNR. The PER is computed as the estimated fraction of the erroneously received/transmitted packets. For received packets we count the CRC errors (at the PHY and MAC layer), while for transmitted packets we count all unsuccessful transmission attempts at the physical layer. This results in a worst case PER estimation, as collisions are also included in the count ¹. All measurements are made available to the application by the mac80211 module every 2 s.

B. Testbed scenario

Among all realistic traffic patterns, we consider three kinds of traffic flows: *mice TCP*, *elephant TCP* and *UDP*. A *mice TCP* flow represents a whispering WS, e.g., a user browsing the web without any kind of background traffic. We emulate whispering stations by implementing the traffic model for mobile web browsing proposed in [7], where the mean web page dimension is 4 KB and the mean reading time (interval time between two consecutive web page requests) is 15 s. An *elephant TCP* flow consists of a bulk HTTP download of a content locally stored at the web server. Finally, we introduce *UDP* flows with a bit rate similar to widely used audio/video peer-to-peer real-time applications, like Skype. All flows are established between WSs and the web server.

Initially we consider two WSs associated to each GW. Then, we show the behavior and performance of the offload procedure in the following cases: (1) one or more GWs are in Light state and try to get rid of their WSs in order to switch off; (2) one or more GWs are in Heavy state and try to relocate their WSs without waking up additional GWs, (3) one or more GWs are in Heavy state and a GW in “off” state has to be woken up to accommodate for the relocation. During the demo users will be allowed to interact with one client, either a mobile phone or a notebook, adding traffic to the BSS to test the reaction of our scheme, e.g., a user could surf the web or start a bulk TCP-based transfer. The system behavior will be shown by the testbed monitoring interface (see Figure 2).

¹Collisions cannot be discriminated from errors due to harsh channel conditions without changing the WS software or the 802.11 protocol.

C. Energy consumption model

We consider that the power consumption of an Alix equipped with an 802.11g radio interface is equal to $P_i = 3.68$ W in idle mode [8]. Depending on the data rate, the consumption in transmission mode is $P_t \in [0.24, 0.44]$ W and in receive mode is $P_r \in [0.27, 0.38]$ W [8]. The consumption of the low-cost, low-power interface (assumed to be an 802.15.4 radio) is $p_s = 186$ μ W in sleep mode and $p_a = 165$ mW in receive/transmit mode [9]. As proposed in [8], we take into account the fraction of energy consumed by each packet crossing the protocol stack, distinguishing between transmission, $P_{xt} = \gamma_{xt}R_t$, and reception, $P_{xr} = \gamma_{xr}R_r$. The parameter R_t (R_r) is the packet transmission (reception) rate, while γ_{xt} and γ_{xr} are hardware-dependent factors. Thus, we compute the energy consumed by a GW in an observation period T as follows:

$$T \cdot [t_{on}(P_i t_{on,i} + (P_r + P_{xr})t_{on,r} + (P_t + P_{xt})t_{on,t} + p_s) + t_{off}p_a]$$

where t_{on} (t_{off}) is the time fraction during which the GW is “on” (“off”), and $t_{on,i}$, $t_{on,r}$ and $t_{on,t}$ are, respectively, the idle, receive and transmit time fraction during the “on” period.

D. The handover problem

Despite the broad spectrum of schemes and protocols aimed at providing seamless or faster handover in 802.11 networks, none of them has reached wide adoption. Solutions based on Mobile IP are complex because they rely on many dependencies, requiring both new hardware and new software to be deployed. This would imply additional costs for both hardware vendors and network operators, with no practical business models that justify these additional expenses [10]. The Inter Access Point Protocol (IAPP), standardized as 802.11f [11], cannot be used to provide faster re-authentication with 802.11i-based security standards, as the latter does not allow security context transfer between GWs. 802.11i provides some alternatives to reduce handover delays, namely Pairwise Master Key Caching (PMK Caching) and pre-authentication over the distribution system. The main problem of these schemes is that they are WS initiated, which makes them unsuitable for a gateway-centric solution like the one we propose. The recent 802.11v [12] IEEE standard includes a BSS transition management that enables a GW to request WSs to handover to another GW, or to indicate a set of preferred GWs to a WS, based on network load balancing needs. However, the GWs must be part of the same Extended Service Set (ESS) and the WSs should support this scheme. In conclusion, nowadays there is no inter-domain network-managed seamless handover solution able to be deployed without changing both GWs and WSs, and without requiring new network element such as the home agent of Mobile IP. As a result, current GWs and

WSs come with no fast handover technology ready to use. Thus, since we want to propose an easily deployable and adoptable solution, we relax the need of a seamless handover and handle the handover as follows. For each WS to be relocated, the selected candidate GW inserts the WS’s MAC address in its white list, while all other GWs (including the offloading one) insert the WS’s address in their black list. This triggers the handover of the WS: the network manager of the WS will automatically scan for GWs and it will have no choice but to associate to the selected candidate GW. We recall that the WS network manager knows all network parameters necessary to associate to any GWs within the federated network. In the demo, we show that a handover takes a couple of seconds, thus causing a very short service disruption for the users; however, interrupted flows must be restarted at the application level. Future work will focus on the design and the implementation of a seamless handover technique and will specifically address this issue.

IV. CONCLUSION AND FUTURE WORK

Our testbed shows a practical approach for real networks aiming at providing considerable energy savings without (i) significantly affecting user experience nor (ii) requiring changes in the WS hardware or software. Savings can be extended to ISPs by adopting DSLAM line aggregation as shown in [13]. Furthermore, our framework provides guidelines for the design of the GW hardware so as to benefit from potential energy savings. Our future work will focus on how to achieve a seamless handover in heterogeneous networks and on how to account for the presence of interference in the capacity estimation.

REFERENCES

- [1] Global e-Sustainability Initiative (GeSI), <http://www.theclimategroup.org/assets/resources/publications/Smart2020Report.pdf>.
- [2] R. Bolla, R. Bruschi, K. Christensen, F. Cucchietti, F. Davoli, S. Singh, “The potential impact of green technologies in next generation wireline networks - is there room for energy savings optimization?,” *IEEE Communications*, vol.49, no.8, 2011.
- [3] EU FP7 FIGARO project, <http://www.ict-figaro.eu/>.
- [4] N. Mishra, K. Chebrolu, B. Rama, A. Patha, “Wake-on-WLAN,” *International Conference on World Wide Web (WWW)*, Edinburgh, Scotland, May 2006.
- [5] C. Rossi, C. Casetti, C.-F. Chiasserini, “Energy-efficient wireless resource sharing for federated residential networks,” *IEEE WoWMoM*, San Francisco, CA, June 2012.
- [6] Wireless driver, <http://linuxwireless.org/>
- [7] G.-F. Zhao, Q. Shan, S. Xiao, C. Xu, “Modelling web browsing on mobile Internet,” *IEEE Communications Letters*, vol.15, no.10, 2011.
- [8] A. Garcia-Saavedra, P. Serrano, A. Banchs, G. Bianchi, “Energy Consumption Anatomy of 802.11 Devices and its Implication on Modelling and Design,” *Co-NEXT’12*, Nice, France, Dec. 2012.
- [9] <http://www.libelium.com/support/waspmote>
- [10] Mobile IP Deployment Difficulties <http://www.ietf.org/proceedings/83/slides/slides-83-mobopts-0>
- [11] IAPP <http://gedes.ifo.edu.br/wp-content/uploads/802.11F-20031.pdf>
- [12] 802.11v <http://standards.ieee.org/findstds/standard/802.11v-2011.html>
- [13] E. Goma, M. Canini, A. Lopez Toledo, N. Laoutaris, D. Kostic, P. Rodriguez, R. Stanojevic, Pablo Yague Valentin, “Insomnia in the access or how to curb access network related energy consumption,” *ACM SIGCOMM*, Toronto, Canada, Aug. 2011.