

The TREND Experimental Activities on “green” Communication Networks

Michela Meo, Yi Zhang, Yige Hu
Dipartimento di Elettronica e Telecomunicazioni
Politecnico di Torino, Italy
{firstname.lastname}@polito.it, S197483@studenti.polito.it

Ivaylo Haratcherev, Alberto Conte
Networking Technology Domain
Alcatel-Lucent Bell Labs, Nozay, France
{firstname.lastname}@alcatel-lucent.com

Antonio Cianfrani, Luca Chiaraviglio, Angelo Coiro
DIET Department
University of Rome “Sapienza”, Italy
{firstname.lastname}@diet.uniroma1.it

Edion Tego, Francesco Matera
Fondazione Ugo Bordoni
Rome, Italy
mat@fub.it

Filip Idzikowski, Łukasz Budzisz, Fatemeh Ganji
Telecommunication Networks Group
Technische Universität Berlin, Germany
{firstname.lastname}@tu-berlin.de

Raffaele Bolla, Olga Jaramillo Ortiz, Roberto Bruschi
Telecommunication Networks and Telematics Lab
RU of CNIT c/o University of Genoa, Italy
{firstname.lastname}@unige.it

Roberto Gonzalez, Carmen Guerrero
Departamento de Ingeniería Telemática
Universidad Carlos III de Madrid, Spain
rgonzal@it.uc3m.es, carmen.guerrero@uc3m.es

Stratos Keranidis, Giannis Kazdaridis, Thanasis Korakis
Department of Computer and Communication Engineering
University of Thessaly, Greece
{efkerani, iokazdarid, korakis}@uth.gr

Abstract— Aimed at answering important questions about the energy demand of current telecom infrastructure and the design of sustainable and energy-efficient future networks, the research of a number of European partners is brought together in the TREND project. In this paper we present the achievements of the Work Package coordinating the experimental activities of the project - WP4. Although not presenting a completely finished portrait yet, the results shown help building a better global view on the “big picture” in the field of energy-efficient networking.

Keywords—TREND; Experimentation; Energy; Power; Wireless Access networks; Backbone networks; Monitoring; Evaluation; 802.11a/g/n; Testbed; Measurement;

I. INTRODUCTION

Energy efficiency of communication networks is becoming an increasingly important topic, especially with the expected explosion of traffic to be carried by those networks. TREND is a Network of Excellence project focused at integrating the activities of various European partners – universities, manufacturers, operators, and research centers – that are doing research in the field of “greening” the telecom networks. The project is coordinated by Politecnico di Torino (PoliTo), Italy.

The paper presents the activities and the accomplishments done in the project’s 4th Work Package (WP4). The main goal of the WP4 is to implement and validate the approaches developed in the theoretical Work Packages (WP2 and WP3). The importance of the work done within WP4 is significant, since its outcome proves the effectiveness of those approaches.

Each section of the rest of the paper presents a separate activity within WP4. Despite that not all work has been finalized yet, thus showing somewhat incomplete view, the set of results presented helps unveiling the “big picture” of future “green” networks.

II. IMPLEMENTATION OF APS ON/OFF ALGORITHMS IN DENSE WIRELESS LAN

Dense Wireless LAN (WLAN) scenarios are typical for corporate or campus facilities [1]. While the high density of APs is needed to have enough capacity to guarantee good service to a potentially large population of users, it also introduces some redundancy and overlapping cells. A high number of devices fits the peak of user demand, but is superfluous during periods of low activity, e.g., at night, leading to considerable power wastage. In order to reduce the number of operating APs, some studies [2-4] propose strategies that lead to switch-offs of the inactive APs, but provide sufficient coverage to estimate the presence of potential users.

A. Cisco Router testbed

The first phase of the experiment is devoted to testbed experiments, aiming at identification of possible problems occurring with the technology. The laboratory testbed setup is represented in the Fig. 1.

The testbed setup comprises several WLAN APs providing overlapping coverage areas. WLAN APs are connected to a Cisco switch equipped with EnergyWise Technology [5]. The

APs are powered through the IEEE 802.3af/at Power-over-Ethernet (PoE) standard interfaces of the router. A management PC connected to the router gathers measurements of a number of parameters for each AP: the amount of carried traffic, the number of associated users, and the power consumption. The management PC executes an energy saving algorithm using that data as input. The output of the algorithm determines the PoE ports of the router to be switched on/off, which turns on/off the corresponding APs. We investigate the APs switching transient phases, including measuring the duration of transients, and observing QoS degradation.

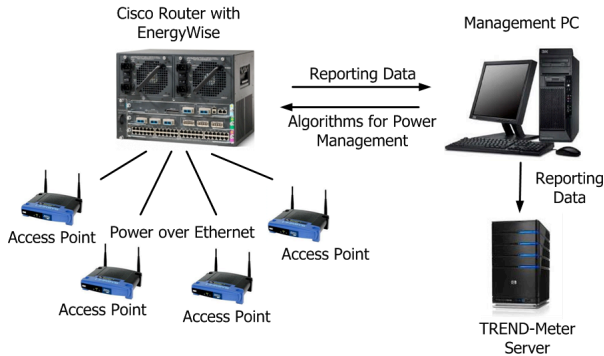


Fig. 1. Testbed setup

B. Energy saving strategy and results

Two energy saving strategies are implemented in the testbed:

1) Switching-on/off strategy based on number of associated users.

In this experiment, two APs are deployed in the same area, and our algorithm is based on the number of associated users, as in [3]. We set the maximum number of users (T_h) of an AP to be 5, and the hysteresis (T_1) in the algorithm to be 1 user. The hysteresis is used for preventing APs from turning on and off too frequently. It is defined that the $(k+1)$ -st AP is switched on when the number of users in the network reaches kT_h , but the number of active APs decreases from $k+1$ to k only when the number of users in the network becomes less than or equal to $kT_h - T_1$. The threshold T_h and T_1 can be set to any value depending on different environment and QoS requirements. Fig. 2 shows the total power consumption of the two APs when the number of associated user varies. The second AP is turned on when the number of associated users increases to 5, but turned off when the number of associated users decreases to 4.

We found out that around 60 s boot time is needed for an AP after it gets power from PoE. Therefore, in the power consumption curve of Fig. 2 there is a gradual increase before the power consumption of the second AP stays constant. It is also shown that there is a short latency (around 10-15 s) between the user behavior variation and the response of the AP. The latency comes from the control process, because during this period the management PC is doing the procedure of: 1) detecting the users' association/de-association by SNMP protocol; 2) running the control algorithm; 3) sending signals to turn on/off the second AP; 4) getting the new power consumption value. Also, the power consumption of an AP is almost constant no matter how many users are associated to it.

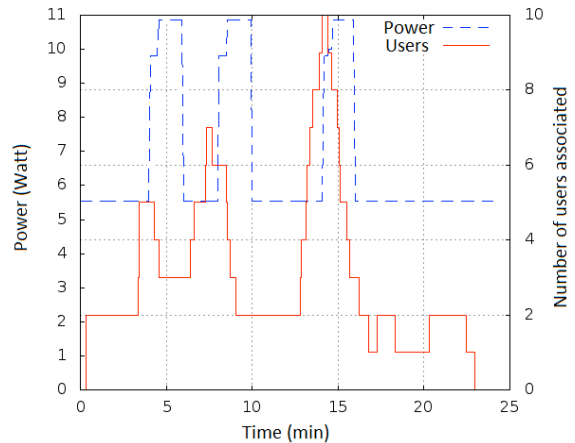


Fig. 2. Power consumption of two APs according to the on/off algorithm based on number of associated users

2) Switching-on/off strategy based on aggregate traffic.

In this experiment, we have implemented an algorithm to turn on/off APs based on the aggregate traffic transmitted by the cluster of APs covering a common area. In the algorithm we turn on/off a specific number of APs depending on the total traffic of all APs and the threshold of throughput we set to the each AP. In the experiment, we generate traffic according to a trace based on real history data in PoliTo WiFi networks, in which the peak value (8 Mbps) is reached around 12 o'clock (as shown in Fig. 3). We have 10 APs covering this testing area. The threshold of the throughput of each AP is set to 800 kbps. The results in Fig. 3 show that the number of working APs is adaptive to the aggregated traffic. Note that the threshold of the throughput can be set to any value depending on different environment and QoS requirement. We use UDP traffic with fixed packet size and data rate in this experiment. We plan future tests with TCP traffic as well.

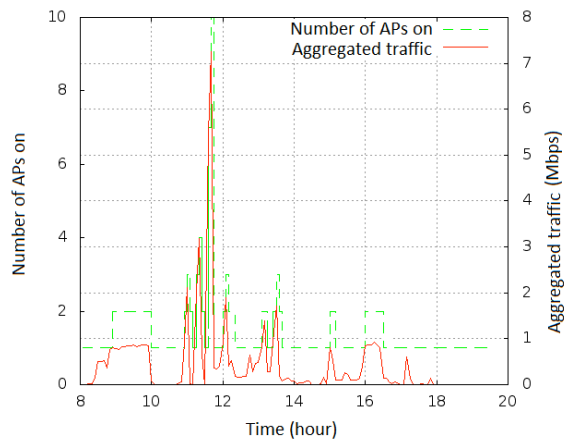


Fig. 3. Number of APs powered on based on the aggregate throughput.

III. ENERGY-AWARE ALGORITHMS OVER A TESTBED

During this activity we have configured a testbed [6] to prove the effectiveness of the algorithms FUFL and DUFL [7]. Our testbed considers both core and access parts of a network. The core part is a ring composed of three Juniper M10(i) routers interconnected by 50 km of fiber cable. The

logical topology is also a ring, where each logical link is composed of 2 parallel GbE optical links. The logical topology is centrally managed. Traffic generators form the access part and are placed at the border of the ring. They are composed of dedicated machines (Ethernet Testing Platform Spirent SPT- 3U and Anritsu MD1230B) and Linux PCs. Two types of traffic are used: i) random traffic with specified minimum and maximum values; ii) sine-like traffic with specified minimum and maximum values as well as the period length in seconds. The maximum value of the sine-like traffic that our generators can produce is 100 Mbps. Corresponding value for the random traffic is 1000 Mbps. The IP routing has been chosen so that all the logical links carry traffic, and that the load exceeds the capacity of a single GbE optical link. It gives us the opportunity to see what happens when the whole logical link or just one out of two parallel GbE links is switched off in the low demand hour. We focus on two classes of energy-savings approaches described in [7], namely FUFL (Fixed Upper Fixed Lower) and DUFL (Dynamic Upper Fixed Lower).

The first class is very simple and therefore attractive for network operators [8]. It involves neither changing of IP routing nor changing of the connectivity of the logical topology. The load on each GbE link constituting the logical link is monitored. A GbE link is switched off when load on the previous parallel GbE link goes down below W_D . It is switched on again when the load on the previous parallel GbE link goes above W_A . W_D and W_A are defined as utilization of a GbE link. We assume bin-packing of traffic in parallel links.

DUFL is more complex, as it allows changing of IP routing, which in turn may increase the number of idle interfaces in the network, and lead even to switching off the whole logical link. A logical link non-existing in the base network cannot be established though. For the sake of this work, we assume the following implementation of DUFL. The decision about an attempt to reroute traffic with the aim of deactivation/activation of a logical link is triggered by violation of the thresholds W_L and W_H , respectively. Both W_L and W_H are defined as utilization of a logical link. In particular, if the aggregated demand on a link goes below W_L , traffic is attempted to be rerouted on an alternative path. Idle logical links with optical interfaces are switched off. The original logical topology and routing is restored when W_H is violated on any logical link. DUFL and FUFL take as input monitored average load values over a period T_M . Validation of the violation of the thresholds is performed every T_L seconds. We set $T_M = 10$ s, and $T_L = 10$ s. The thresholds are set to the following values for demonstration of the power saving approaches: for FUFL $W_D = 0.977$ and $W_A = 0.985$, for DUFL $W_L = 0.4885$ and $W_H = 0.9925$. Power consumption of interfaces has been measured using a commercial power meter.

Fig. 4 shows the total power consumption and the power saving for the testbed running both FUFL and DUFL on the logical topology over 1000 seconds. The total power consumption values correspond to power consumed by all active GbE interfaces together with the routers according to the measurements performed in the testbed. Power consumption varies more frequently with DUFL than with

FUFL, since our implementation of DUFL is more aggressive in turning off the network interfaces – it attempts to switch off the whole logical links. This in turn produces in general higher power saving compared to FUFL.

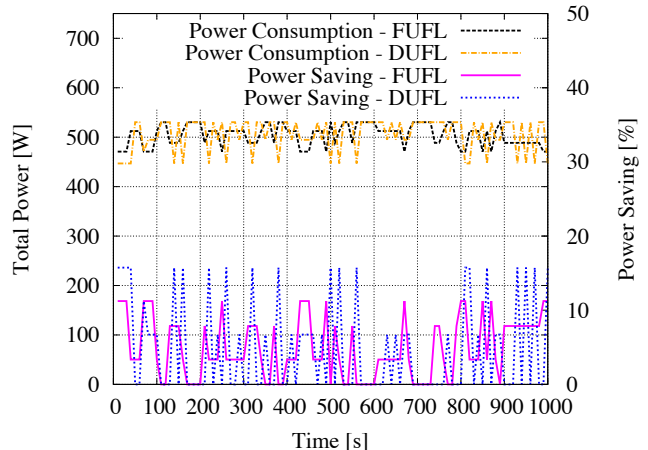


Fig. 4. Power consumption and power saving for FUFL and DUFL.

IV. THE TREND-METER

The main goal of the Trend-Meter [9] is to collect data of power consumption and utilization from a variety of devices located in the Internet: a centralized server gets measurements from the devices hosted in different Trend partners' locations. As a second goal, Trend-Meter aims at consolidating these measurements to study whether there are similarities in the patterns of power and utilization of the devices. Additionally, collected data from Trend-Meter is made publicly available, with friendly graphical representation of the information. The architecture is composed of three main units: a device back-end for collecting the measurements, a server back-end collecting and storing the information from all the devices, and a server front-end to display the information on a web site. The set of monitored devices includes one server machine, a set of personal computers, a university datacenter, a router deployed in a campus, a software router, and a femto cell.

An example of produced output is reported in Fig. 5. The figure reports the histogram for the utilization of the campus router. In this case, the power consumption (not reported here) is almost constant, meaning that this device does not exploit power saving functionalities. At the same time, however, the utilization of the device is instead quite variable, meaning that power saving approaches should be deployed for such devices.

A. Integration with monitoring systems

Network monitoring represents one of the main components for a network management tool. Existing network monitoring software tools control many features, such as traffic levels, hardware and applications states. However, little attention is paid to monitoring the power consumption of network devices. During this activity, we have integrated the measurements of the Trend-Meter inside software for infrastructure monitoring. Among the different solutions currently available in the market, we have selected Nagios,

since this software is an open-source platform for monitoring IT devices.

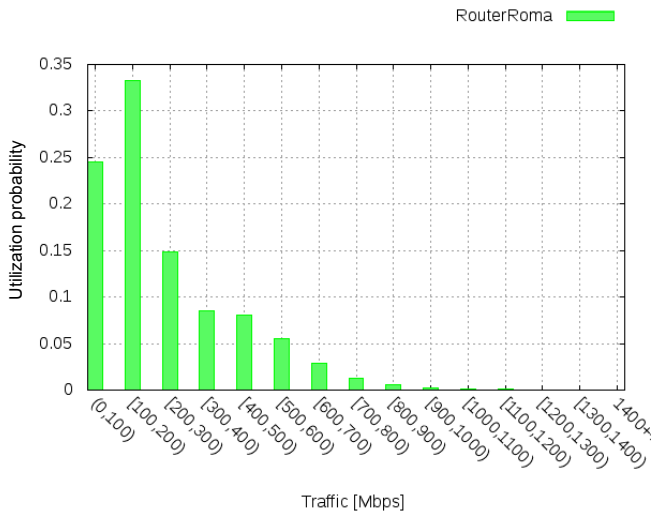


Fig. 5. Histogram of the traffic distribution for the Campus router

We have divided the work into two main tasks: a) implementation of the TREND-Meter plugin to export the data into the Nagios format; b) integration of the exported data inside Nagios.

Fig. 6 reports a schematic description of the architecture. Nagios runs the TREND-Meter plugin whenever there is a need to check the status of one of the devices monitored by the TREND-Meter. The TREND-Meter plugin then interacts with the TREND-Meter database to retrieve the information on the device status.




Fig. 6. Integration of TREND-Meter in Nagios.

Host	Service	Status	Last Check	Status Information
DataCenter1HU	Power Consumption [W]	OK	08-09-2013 16:33:45	11271.6620521173
	Utilization [Mbps]	OK	08-09-2013 16:34:34	0.10742468485342
Femtocell	Power Consumption [W]	WARNING	08-09-2013 16:32:23	0
	Utilization [Mbps]	WARNING	08-09-2013 16:33:58	0
PCPoli	Power Consumption [W]	OK	08-09-2013 16:32:47	26.8484848484848
	Utilization [Mbps]	OK	08-09-2013 16:34:36	0.0023706666666667
PolitoSubnet	Power Consumption [W]	OK	08-09-2013 16:34:12	3075
	Utilization	OK	08-09-2013 16:36:01	1.4734725
SWRouter	Power Consumption [W]	OK	08-09-2013 16:34:50	169.534125576765
	Utilization [Mbps]	OK	08-09-2013 16:36:26	0.563535235294118
routerRoma	Power Consumption [W]	WARNING	08-09-2013 16:32:45	0
	Utilization [Mbps]	WARNING	08-09-2013 16:34:20	0

Fig. 7. The Trend-meter data inside Nagios.

The data retrieved by the TREND-Meter has been integrated inside Nagios. In particular, the information about power consumption and load is stored, together with the status of each device (on/off). This information is then processed inside Nagios in order to obtain consolidated results, like charts and alerts. An example of output is reported in Fig. 7. In

this case, Nagios displays warnings for the measurements that are not available. Additionally, the value of power consumption and utilization is stored in the status information.

V. MEASURING THE ENERGY COST OF NETWORKING INTENSIVE APPLICATIONS

The network-intensive nature of some networking applications affects greatly the energy consumption of the host. In order to analyze this energy behavior we performed experiments running two different protocols used for downloading content from the Internet: Hypertext Transport Protocol (HTTP) and BitTorrent (BT). HTTP is a client-server protocol in which the download is carried out over a single TCP connection, while BitTorrent is a p2p protocol in which the download is simultaneously done from different sources.

We focus our analysis on the CPU, a computer component with higher energy requirements. Current CPUs already implement advanced power management features, which can be set through the Advanced Configuration and Power Interface (ACPI) standard [10]. In this section we analyze its energy consumption in presence of HTTP and BT

A. Evaluation Methodology

The testbed is composed of: The System Under Test (SUT), a Linux workstation acting as Gateway; and a multi-channel Data Acquisition (DAQ) device for collecting data from DC power consumption probes. For measuring the energy consumption of the CPU, we used a riser board for ATX power connectors [11], allowing putting some current and voltage probes on the CPU supply rails. In order to evaluate the CPU utilization, we used Oprofile [12], an open source tool that realizes a continuous monitoring of system dynamics with a frequent and quite regular sampling of CPU HW registers. To evaluate the impact of downloading different contents we decided to consider the following files: two Debian distributions – Debian-6.0.7-amd64-CD-1.iso (676 MB) and Debian-6.0.5-amd64-DVD-1.iso (4.64 GB), and the most popular film in the principal BitTorrent tracker during the test: Parker.avi (2.5 GB). The access link was emulated by setting the data rate between the “gateway” and the SUT to 10, 100, and 1000 Mbps. Due to the time demands for performing each experiment, it has not been possible, in our tests, to calculate confidence intervals.

B. Experimental Results

Fig. 8 shows the average energy consumption per file of BT and HTTP, for all test cases. From this figure we see that the energy consumption depends on the amount of data downloaded, i.e., in the case of the IsoDVD it consumes up to 120 kJ. Fig. 9 reports the average energy consumption to download a MB of data. From this figure we see that the energy consumption is also dependant on the link speed. This figure also shows that consumption decreases with higher speeds. This effect is more evident when the link changes from 10 to 100 Mbps, where we can obtain energy savings up to 60%, and 35% when the link changes from 100 to 1 Gbps. The previous results are directly related to the time needed to download a file, as shown in Fig. 10. Note that in the case of the ISOs, BT needs less time than HTTP to download the data

allowing lower energy consumptions, i.e. 40% less in the case of the DVD downloaded over the 100 Mbps link. Even if BT protocol is more complex than HTTP due to the additional processing we can still obtain energy gains. In the case of the movie the behavior is different. BT consumes more energy than HTTP. This is because in this case the download is more competitive with a higher number of peers querying for the same resources, forcing to start more connections to obtain the same amount of resources.




Fig. 8. Average Energy Consumption of BT and HTTP with different downloading contents.




Fig. 9. Average Energy Consumption per MB of the different downloading contents with BT and HTTP

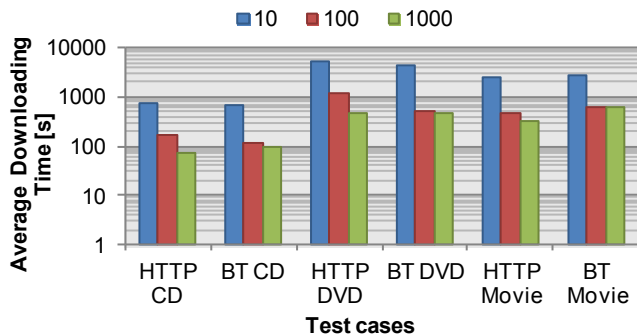


Fig. 10. Average Downloading time of the different downloading contents with BT and HTTP.

VI. ONLINE ENERGY CONSUMPTION MONITORING OF WIRELESS TESTBED INFRASTRUCTURE THROUGH THE NITOS EMF FRAMEWORK

Development of energy-efficient protocols and algorithms requires in-depth understanding of the power consumption characteristics of real world devices. To this aim, energy efficiency analysis is performed by the research community, mainly focusing on the development of power consumption

models. However, recent studies [13] have highlighted the inability of existing models to accurately estimate energy consumption even in simple scenarios, where the operation of a single device is analyzed. In this section, we present the innovative NITOS Energy consumption Monitoring Framework (EMF) able to support online monitoring of energy expenditure, along with the experiment execution.

A. NITOS EMF Framework

The NITOS EMF framework is built on a distributed network of low-cost, but highly accurate energy monitoring devices, named NITOS Advanced Chassis Manager (ACM) card, and is fully integrated with the large-scale wireless NITOS testbed [14]. NITOS ACM acts as a voltage sampling device able to monitor voltage drop on a current shunt resistor that is placed in series with the wireless Network Interface Card (NIC) to be measured. The ACM is composed of both open-source and custom-made hardware components.

The Arduino Mega board is the core module of the developed hardware. To provide network connectivity, we equipped the core module with the Arduino Ethernet shield. The shield also supports external storage capabilities, through the available micro SD slot, providing for online measurement logging and posterior network transfer for further processing. Towards improving measurement accuracy, we also use the Texas Instruments INA139 [15] current shunt monitor, which provides for amplification of low differential voltage signals.

To turn the hardware platform into a functional framework that can be fully integrated with the NITOS testbed, we developed some software as well. Through modifications of the Arduino software, we increased the sampling rate from 10 to 63 KHz with 10-bit resolution, while only reducing the accuracy by approximately 11% [16]. Moreover, we developed a tiny Web Server that operates on each individual ACM card, in order to provide for remote control and transfer of captured data.

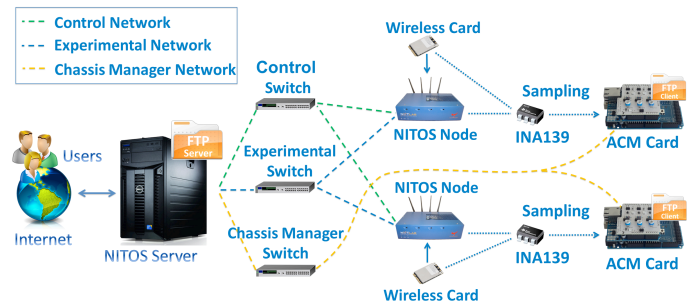


Fig. 11. Integration of the NITOS EMF framework with the overall testbed architecture.

The integration of the NITOS EMF framework with the overall testbed architecture is illustrated in Fig. 11. NITOS currently consists of 50 wireless nodes, deployed outdoors at the exterior of the University of Thessaly campus building. Two Gigabit Ethernet switches interconnect the nodes with NITOS server, namely the *Control* switch that provide for control of experiment execution and measurement collection and the *Experimental* switch, which can be used for conducting wired experiments. A third Gigabit Ethernet

switch, namely the *Chassis Manager*, is dedicated to controlling the operational status of the nodes.

B. Energy monitoring of realistic wireless experiments

In this section we report a representative experiment demonstrating the capabilities of the NITOS EMF framework.

The experimental scenario includes 3 stations (STAs) that are associated with the same AP, which is operating on Ch.1 of the 2.4 GHz band and are simultaneously uploading a file of 25 MBs. All the nodes feature IEEE802.11a/g compatible Atheros AR5424 chipset. The three associated STAs are configured at different PHY-layer Transmission Rates (TX_R) – the first at 54 Mbps, the second at 18 Mbps, and the last at 24 Mbps. Another pair of collocated nodes is generating interference, with the station node (STA_{INT}) transmitting on uplink at the application layer traffic load of 15 Mbps. Fig. 12(a) illustrates the topology along with the PHY-layer rate settings, configured for each specific NIC. The experiment is repeated 6 times, where in each different run we configure the AP_{INT} to operate on a different channel between Ch.1 and Ch.6 of the 2.4 GHz band. During each different run, we monitor the energy consumed by each NIC and plot the collected results in Fig. 12(b).

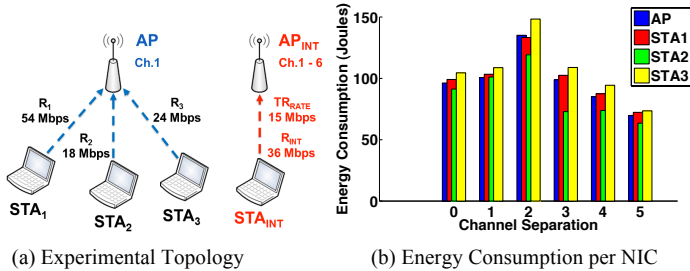


Fig. 12. Energy monitoring of realistic experiment using 802.11a/g NICs.

As long as STA_{INT} operates on the same channel with the 3 STAs, 802.11 efficiently controls channel contention, while on the other hand as channel separation increases, transmissions of STA_{INT} are not always detected, resulting in frame collisions and subsequent frame retransmissions. The overall effect is that the file transmission duration increases for each individual node and thus the resulting energy expenditure is impacted. As AP_{INT} moves from Ch.3 to Ch.5, we notice that the energy consumption tends to decrease for all NICs, resulting in the lowest monitored values in the case that the interfering link operates on Ch.6, as it no more interferes with the 3 STAs.

VII. CONCLUSIONS

Forecasted expansion of traffic to be carried by the communication networks will result in tremendous increase of their electrical power consumption, if no energy-efficiency improvements are made in the meantime. This is unacceptable, both from operators' point of view (surge of OPEX), and from the "greening" infrastructure perspective.

In this paper we show experimental work about making the communication networks more energy-efficient – an effort by

multiple collaborating partners. Those experiments and testbeds are aimed at validating theoretical concepts, assessing algorithms' effectiveness, and providing valuable data for further developments. Instead of focusing on a single aspect or a network element or technology, we present "greening" research covering various points of the communication chain, together with the work allowing to obtain the global energy "pulse" of the communication network (Trend-Meter). Thus, a global picture can be constructed, which is a key in designing the communication systems of the future.

ACKNOWLEDGMENT

The work presented here is within the framework of the Network of Excellence TREND project, which has received funding from the European Union Seventh Framework Programme, under grant agreement n. 257740.

REFERENCES

- [1] W. Lemstra and V. Hayes, J. Groenewegen, "The Innovation Journey of Wi-Fi: The Road To Global Success," Cambridge Univ. press, 2011.
- [2] F. Ganji, L. Budzisz, and A. Wolisz, "Assessment of the power saving potential in dense enterprise WLANs" TKN Group, TU Berlin, Tech. Rep. TKN-13-003, Apr. 2013.
- [3] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, and M. Meo, "A Simple Analytical Model for the Energy-Efficient Activation of Access Points in Dense WLANs," Proc. e-Energy, Passau, Germany, Apr. 2010.
- [4] A. P. Couto da Silva, M. Meo, and M. Ajmone Marsan, "Energy-Performance Trade-off in Dense WLANs: a Queuing Study," Computer Networks, vol. 56, no. 10, pp. 2522–2537, July 2012.
- [5] Cisco EnergyWise Technology, <http://www.cisco.com/en/US/products/ps10195/index.html>, Last access: June 2013
- [6] A. Valenti, A. Rufini, S. Pompei, F. Matera, S. Di Bartolo, C. Da Ponte, D. Del Buono, and G. T. Belefli, "QoE and QoS comparison in an anycast digital television platform operating on passive optical network," Proc. Networks, Rome, Italy, Oct. 2012.
- [7] F. Idzikowski, S. Orłowski, C. Raack, H. Woesner, and A. Wolisz, "Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings," Optical Switching and Networking, Special Issue on Green Communications, vol. 8, no. 3, pp. 181-200, July 2011.
- [8] F. Idzikowski, L. Chiaraviglio, R. Duque, F. Jimenez, and E. Le Rouzic, "Green Horizon: Looking at Backbone Networks in 2020 from the Perspective of Network Operators," Proc. ICC, Budapest, Hungary, June 2013.
- [9] L. Chiaraviglio, R. Bruschi, A. Cianfrani, O. Jaramillo Ortiz, and G. Koutitas, "The TREND Meter: Monitoring the Energy Consumption of Networked Devices," International Journal of Business Data Communications and Networking (IJBDNC), IGI Global, Special Issue on Green Networking and Computing, on press, 2013.
- [10] "The Advanced Configuration and Power Interface (ACPI)", <http://www.acpi.info/>, Last access: June 2013
- [11] "The ATX specification 2.1", http://www.formfactors.org/developer/specs/atx2_1.pdf, Last access: June 2013
- [12] "Oprofile", <http://oprofile.sourceforge.net/>, Last access: June 2013
- [13] A. Garcia-Saavedra, P. Serrano, A. Banchs, and G. Bianchi, "Energy consumption anatomy of 802.11 devices and its implication on modeling and design," Proc. CoNEXT, Nice, France, Dec. 2012.
- [14] "NITOS Wireless Testbed", <http://nitlab.inf.uth.gr>, Last access: June 2013
- [15] "Texas Instruments INA139", <http://goo.gl/rPQLB>, Last access: June 2013
- [16] "Atmega ADC accuracy vs clock speed", <http://goo.gl/qTlhx>, Last access: June 2013