



The greening potential of content delivery in residential community networks



Rosario G. Garroppo^a, Gianfranco Nencioni^a, Luca Tavanti^{a,*}, Bernard Gendron^b

^a *Dipartimento di Ingegneria dell'Informazione, Università di Pisa, Italy*

^b *CIRRELT and DIRO, Université de Montréal, Canada*

ARTICLE INFO

Article history:

Received 25 November 2013

Received in revised form 3 July 2014

Accepted 25 August 2014

Available online 4 September 2014

Keywords:

Energy efficiency

Residential community networks

Wireless mesh networks

Resource allocation

Routing

Optimisation

ABSTRACT

In this paper we focus on a Residential Community Network (RCN) in which the various home gateways have (limited) caching capabilities, can be shared among the users, and can exploit multi-hop wireless paths to communicate with each other. We analyse such a RCN from an energy-aware perspective, investigating its energy-saving potential in delivering the contents to the end users. To this aim, we define a resource allocation and routing scheme and, by means of an integer linear programming model, we solve to optimality the problem of associating the user terminals to the home gateways and computing the data flow paths. Accordingly, the resulting allocation pattern can serve as an upper bound for the global “greening” capability of the RCN. A comparison between our method and other techniques for resource sharing in RCNs sheds light on the pros and cons of both our scheme and the other approaches.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

The steep rise in the fruition and sharing of multimedia contents over the Internet has pushed for the definition and deployment of novel network architectures in order to efficiently deliver such contents to the end users. Nowadays, Content Distribution Networks (CDNs) are the most popular way for achieving this goal. By caching and replicating the contents on servers placed near to the network boundaries, CDNs allow to reduce the congestion of the Internet and improve the quality of the delivery service [1,2].

The evolution produced by CDNs, however, should not limit its scope solely to the core and edge parts of the network. In fact, novel access infrastructures, such as Residential Community Networks (RCNs), are emerging as a sort of physical peer-to-peer paradigm that complements the

content sharing one. In RCNs, the users build a (wireless) network among their home gateways in order to share the capabilities of the gateways and/or the contents hosted at each user's premises [3,4]. In this scenario, the problem of delivering multimedia content should embrace not only the Internet, as for current CDN approaches, but also the opportunities offered by the spontaneous RCN infrastructure [5].

To date, most of the research on CDNs has focused on the techniques and algorithms to place the contents on the edge network devices. Only occasionally has this work involved the access network or directly the customer premise equipment. Specifically, some authors have included in the CDN the residential gateway (GW), i.e. the device that connects the home network to one or more access networks to deliver services to the user. The idea is that the GW can be exploited, either as a replacement or as an addition to the CDN, to store the contents destined to the users themselves [6,7]. With reference to the RCN, some authors have explored the possible advantages of employing the

* Corresponding author.

E-mail address: luca.tavanti@iet.unipi.it (L. Tavanti).

GWs as the caching points for various types of contents [8,9].

We build our work on the system architecture that arises from joining the CDN paradigm with the RCN model. Specifically, we address a “Content-Delivery Wireless Community Network” (CDWCN), i.e. a wireless RCN whose main goal is the delivery of the contents to the RCN users. In this CDWCN, the GWs are able to perform limited caching, either spontaneously, as in file swarming [10], or on behalf of the service provider, as in [7]. Therefore the contents can be retrieved from the various GWs (if cached on them) or directly from the Internet. Central to the CDWCN is also the willingness of the users to share their GWs, both in terms of connectivity and of cached contents. Finally, the GWs are capable of building and managing wireless multi-hop paths to connect with each other – in other words, the RCN backbone is implemented as a wireless mesh network (WMN) [11].

Given these premises, the main contribution of the paper is an investigation on the potential of the CDWCN from an energy-aware perspective. In detail, we devise a resource allocation and routing scheme (named GRARC, i.e. Green Resource Allocation and Routing for CDWCN) that jointly considers all the distinctive CDWCN features mentioned above in order to minimise the overall energy consumption for delivering the contents to the end users. Accordingly, we formulate a mathematical programming model to analyse the possible energy savings of this scheme, also in comparison to other state-of-the-art approaches. The final goal and result of our work is to quantify the maximum overall power saving that can be achieved by the GRARC strategy. We can thus provide an indication on how much a CDWCN can be “green” if the optimal resource allocation and content routing are enforced. To the best of our knowledge, no previous work has performed a similar study.

Our analysis starts from a given content placement pattern (e.g. laid out by some CDN approach) and then aims at defining what GWs shall be powered on, what user terminals shall be associated to them, and what delivery strategy (i.e. routing the contents) shall be employed in order to minimise the overall energy consumption. In particular, we do not follow the greedy approach of minimising the energy consumed by the sole CDWCN, but we also account for the energy required by the Internet for delivering the data to the CDWCN users. Indeed, the sole minimisation of the CDWCN power consumption usually results in a very simple problem, whose solution, however, might even lead to an increase in the global energy usage (as it will be shown in Section 5.4). Also, since many works on content caching already exist, we do not deal with content placement on the GWs, but take it for granted that a suitable CDN strategy has been enforced (see e.g. [7,10,12,13]). Therefore our work does not integrate the definition of a CDN approach, but is complementary to it.

The paper is structured as follows. The next section presents the reference CDWCN architecture in more detail, together with an overview of the most closely related works in the field. Our GRARC scheme is described in Section 3, whereas Section 4 illustrates some alternative optimisation approaches to CDWCN. The performance




Fig. 1. Architecture of the CDWCN.

and impact of GRARC is studied and compared to these approaches in Section 5. The conclusions are drawn in Section 6.

2. Content-delivery wireless community networks

2.1. Architecture

The reference CDWCN architecture is shown in Fig. 1. The users' home gateways (GWs) are connected to the Internet by means of, for example, a DSL link. The GWs are also part of the CDN service, which decides the contents that are cached on each GW.

Then, each GW is equipped with two different wireless interfaces. One interface is used to serve the assigned user terminals¹ (UTs), while the other is used to connect with the neighbouring GWs. The two interfaces work on non-overlapping channels in order to avoid mutual interference. For example, it is common practice in WMNs to operate the inter-GW connections over Wi-Fi links in the 5 GHz band (i.e. IEEE 802.11a/n), whereas the GW-UT connections use the 2.4 GHz band (IEEE 802.11g/n).

Finally, we assume that a UT can be assigned (associated) to one GW only. Typically, the UTs of a given user will be associated to the GW of the same user. However, we do not prevent UTs from associating with neighbouring GWs if a radio link is available and the optimum allocation requires so.

Therefore, according to the illustrated CDWCN architecture, there are three possibilities for a GW to deliver a content to an assigned UT: (i) the content is cached on the GW itself; (ii) the content is cached on another GW that can be reached either directly or by means of a multi-hop connection; and (iii) the content must be downloaded from the Internet.

2.2. Related work

The capability and potential of the CDWCN or similar architectures, especially with regard to the energy issue, has recently received some attention from the research community.

For example, Rossi et al. [14] devised some empirical procedures and a distributed protocol for relocating the

¹ Note that a single human user may be in possession of many UTs.

user terminals of a RCN to the various GWs in order to switch off some GWs and thus save energy. Similarly, Goma et al. [15] formulated an optimisation model that takes advantage of the overlapping of home wireless LANs to aggregate the user traffic in as few GWs as possible. In both works, however, the possibility of transferring data among the GWs is not accounted for, nor is the caching capability of the GWs. Hence, the resulting CDWCN model would be very simplified.

Han et al. [4] assumed a neighbourhood model which is quite close to the CDWCN one, and showed how such local connectivity and storage can be exploited to reduce the traffic on the access network. Yet, the energy issue is out of the scope of their work.

Valancius et al. [7] proposed NaDa, a distributed platform that uses ISP-controlled home gateways (GWs) to provide computing and storage services to the end-users. NaDa adopts a peer-to-peer philosophy, but the GWs are not shared among the users and the connection among them is made available and managed by the ISP. Therefore, despite addressing the energy-saving topic, Valancius' work can hardly fit the CDWCN context. The same can be said for the work of den Hartog et al. [6], who assessed the pros and cons of caching various types of content in the GWs. The evaluation is in terms of bandwidth usage and blocking probability, without caring neither for the energy consumption nor for the availability of a RCN.

Manetti et al. [8] designed a centralised and a peer-to-peer architecture to deliver contents in community networks. However, the focus of the work is mostly on the protocol design and content indexing aspects, with no mention to the power consumption issues. Also, neither architectures are specific for wireless RCNs (WCNs).

The wireless RCN paradigm is, instead, central to Alasaad et al. [9], who studied the problem of energy consumption in sharing a given content over the WCN. In Alasaad's view, the WCN is at the same time origin and destination of a single "viral" content. Also, only some nodes can perform caching, and the content is retrieved by means of peer-to-peer procedures that are agnostic on the physical distance between the peers. Consequently, the achieved results cannot be applied to the CDWCN system under study in this paper.

Taghizadeh et al. [16] studied cooperative caching policies for reducing the provisioning cost in spontaneous networks formed by mobile devices. Although the optimisation model accounts for multi-hop paths, the energy consumption is not considered at all. Shevade et al. [17] devised an algorithm for placing the contents on the Access Points serving a vehicular network. The main focus of the authors is on content dissemination, based on the prediction (and exploitation) of the vehicle trajectories. No consideration is given to the energy topic.

Finally, some research on home router sharing has been performed within the ongoing SmartenIT project [18]. In particular, Seufert et al. [19] proposed a framework that targets traffic offloading from mobile networks to WiFi, content prefetching and caching on the home routers, and content delivery. Router sharing and a trust mechanism are also part of the framework. The authors' main focus is on how information from online social networks

and from user mobility can be exploited to prefetch and store popular contents on the home routers in order to reduce the load on the mobile networks. Energy reduction, however, is seen solely as a positive, but not quantified side effect. Also, the existence of direct wireless connections among the routers is not considered.

In summary, there is no previous study that faces the same problem as the one we present in the paper. We recall that our goal is to minimise the energy consumption by jointly considering and acting on the three aspects that are distinctive of the CDWCN: sharing the access to the GWs, sharing the cached contents, and transporting the contents among the GWs. Therefore, the cited works either address single aspects of the problem, or target substantially different systems.

3. Green Resource Allocation and Routing for CDWCN

In this section we describe GRARC, an approach that exploits the distinctive features of the CDWCN to make it "greener". Internet connectivity sharing, i.e. the possibility to associate any UT to any reachable GW, allows to power off some GWs. Caching on the GWs can be exploited to save Internet bandwidth and energy. The mesh service enables the dissemination of the content cached in any GW to any UT in the CDWCN.

More precisely, this study is aimed at assessing the potential of the CDWCN in terms of energy savings, i.e. what is the maximum degree of energy that the optimal resource allocation permits to save. For this purpose we have built a static system model so that, for any given "snapshot" of the system, we can derive the most energy-efficient allocation. Accordingly, we refer solely to the instantaneous characteristics of the components of the system. For example, we characterise the contents by means of the bit rate, rather than their size. Some considerations on how the system dynamicity may affect the optimal allocation are given in the concluding section.

To reach our goal, we build a mathematical program that takes as input the content placement (on the GWs), the device connectivity, and the content demands, and decides (outputs) the UT–GW associations and the routing of the contents. Fig. 2 gives a graphical view of the abstraction of the CDWCN used to formulate the




Fig. 2. Sketch of the CDWCN abstraction. Solid lines are physical links, dashed lines are logical links.

mathematical model. There are three kinds of nodes (UTs, GWs, and contents) and three kinds of links (GW–UT, GW–GW, and content–GW). The GW–UT and GW–GW are physical (wireless) links, and therefore are characterised by a data rate dependent on the modulation and radio propagation rules. Conversely, the content–GW links are logical, because we assume that every content on the Internet can be accessed by every GW with no data rate restrictions. Note that if a GW is turned off, this reduces both the total caching capability and the bandwidth and connectiveness of the RCN.

3.1. Notation

In the addressed problem, we consider six sets:

- \mathcal{G} , the set of deployed gateways (GWs);
- \mathcal{U} , the set of user terminals (UTs);
- \mathcal{C} , the set of available contents;
- \mathcal{D} , the set of content demands (c, u) , $u \in \mathcal{U}$, $c \in \mathcal{C}$;
- \mathcal{E}' , the set of GW–UT edges $\{(g, u) : r_{gu} > 0\}$, $u \in \mathcal{U}$, $g \in \mathcal{G}$;
- \mathcal{E}'' , the set of GW–GW edges $\{(g, h) : r_{gh} > 0\}$, $g, h \in \mathcal{G}$,

where r_{gu} and r_{gh} are the data rates between a GW g and a UT u and between a GW g and a GW h , respectively. Note that: (i) the edges are directed, i.e. $(g, h) \neq (h, g)$, (ii) \mathcal{E}' and \mathcal{E}'' are two disjoint sets, i.e. $\mathcal{E}' \cap \mathcal{E}'' = \emptyset$, and (iii) we have not made the set of content–GW edges explicit because it is not relevant to the model formulation.

The known parameters of the problem are the following:

- P^{GW} [W], the power consumption of a GW;
- E^I [J/b], the average energy for retrieving a bit from the Internet;
- E^W [J/b], the average energy for transmitting a bit through a wireless link (either GW–GW or GW–UT);
- r_{ij} [b/s], the average data rate between the vertexes $i \in \mathcal{G}$ and $j \in \{\mathcal{U} \cup \mathcal{G}\}$;
- b_c [b/s], the rate needed for retrieving the content c ;
- t_{cg} , a binary flag that is set to 0 if the content c is cached in GW g , to 1 otherwise.

We then define the following sets of binary variables:

- q_g , which is set to 1 if GW g is powered on, to 0 otherwise;
- x_{gu} , which is set to 1 if UT u is assigned to GW g , to 0 otherwise;
- y_{hg}^{cu} , which is set to 1 if content c is delivered to UT u through edge (h, g) , to 0 otherwise;
- z_g^{cu} , which is set to 1 if UT u retrieves the content c from GW g (i.e. either c is cached in g or c is downloaded from the Internet by g), to 0 otherwise; note that u needs not to be assigned to g .

3.2. Mathematical programming model

The objective function of our problem consists in minimising the total consumed power of the CDWCN:

$$P^{GW} \sum_{g \in \mathcal{G}} q_g + \sum_{(c, u) \in \mathcal{D}} b_c \left\{ E^I \sum_{g \in \mathcal{G}} (t_{cg} z_g^{cu}) + E^W \sum_{(h, g) \in \mathcal{E}''} y_{hg}^{cu} \right\}, \quad (1)$$

subject to:

$$q_g \in \{0, 1\}, \quad \forall g \in \mathcal{G}, \quad (2)$$

$$x_{gu} \in \{0, 1\}, \quad \forall (g, u) \in \mathcal{E}', \quad (3)$$

$$y_{hg}^{cu} \in \{0, 1\}, \quad \forall (h, g) \in \mathcal{E}'', \quad \forall (c, u) \in \mathcal{D}, \quad (4)$$

$$z_g^{cu} \in \{0, 1\}, \quad \forall g \in \mathcal{G}, \quad \forall (c, u) \in \mathcal{D}, \quad (5)$$

$$x_{gu} + \sum_{h: (g, h) \in \mathcal{E}''} y_{hg}^{cu} - z_g^{cu} - \sum_{h: (h, g) \in \mathcal{E}''} y_{hg}^{cu} = 0, \quad \forall (c, u) \in \mathcal{D}, \quad \forall g: (g, u) \in \mathcal{E}', \quad (6)$$

$$\sum_{h: (g, h) \in \mathcal{E}''} y_{hg}^{cu} - \sum_{h: (h, g) \in \mathcal{E}''} y_{hg}^{cu} - z_g^{cu} = 0, \quad \forall (c, u) \in \mathcal{D}, \\ \forall g: (g, u) \notin \mathcal{E}', \quad (7)$$

$$\sum_{g \in \mathcal{G}} z_g^{cu} = 1, \quad \forall (c, u) \in \mathcal{D}, \quad (8)$$

$$\sum_{g: (g, u) \in \mathcal{E}'} x_{gu} = 1, \quad \forall u \in \mathcal{U}, \quad (9)$$

$$x_{gu} \leq q_g, \quad \forall (g, u) \in \mathcal{E}', \quad (10)$$

$$y_{hg}^{cu} \leq q_g, \quad \forall (h, g) \in \mathcal{E}'', \quad \forall (c, u) \in \mathcal{D}, \quad (11)$$

$$y_{gh}^{cu} \leq q_g, \quad \forall (g, h) \in \mathcal{E}'', \quad \forall (c, u) \in \mathcal{D}, \quad (12)$$

$$z_g^{cu} \leq q_g, \quad \forall g \in \mathcal{G}, \quad \forall (c, u) \in \mathcal{D}, \quad (13)$$

$$\sum_{u: (g, u) \in \mathcal{E}'} \left\{ \frac{\sum_{(c, u) \in \mathcal{D}} b_c}{r_{gu}} \cdot x_{gu} \right\} \leq \rho q_g, \quad \forall g \in \mathcal{G}, \quad (14)$$

$$\sum_{h: (h, g) \in \mathcal{E}''} \sum_{(c, u) \in \mathcal{D}} \left\{ \frac{b_c}{r_{hg}} \cdot (y_{hg}^{cu} + y_{gh}^{cu}) \right\} \leq q_g, \quad \forall g \in \mathcal{G}. \quad (15)$$

The objective function (1) is composed by three terms: the power consumption of powered-on GWs, the power consumption for retrieving contents from the Internet, and the power consumption for transferring contents among the GWs. As for the constraints, Eqs. (2)–(5) define the binary variables, Eqs. (6) and (7) are the flow conservation constraints (see below for a brief explanation), Eq. (8) impose that each UT u retrieves the content c from exactly one GW (single source of demand (c, u)), Eq. (9) impose that each UT must be assigned to exactly one GW, Eqs. (10)–(12) impose that, if a GW is powered off, no UT or GW can be connected to it, Eq. (13) impose that no content can be retrieved from a powered-off GW, and finally Eqs. (14) and (15) are the classical capacity constraints (one for each wireless interface of the GW – we assume that the wired link to the Internet has unlimited bandwidth). In particular, in (14) the constant parameter $\rho \in (0, 1]$




Fig. 3. Depiction of the flow conservation constraints (6). The same figure applies to (7) by removing user u , the (g, u) link, and the x_{gu} variable.

may be used to set a limit on the GW usage, thus explicitly leaving some spare bandwidth for data traffic that do not rely on the CDN (such as e-mail or voice over IP).

With regard to the flow conservation constraints (6) and (7), a pictorial explanation is given in Fig. 3. Flow conservation is imposed in both cases at GW g . Note how we have employed the x_{gu} variable to indicate (in addition to the $g - u$ assignment) that demand (c, u) is carried over the (g, u) link. In a more complete and formal version, we should have used some y_{gu}^{cu} variables, but it is also true that $x_{gu} = y_{gu}^{cu} \forall (g, u) \in \mathcal{E}' \forall (c, u) \in \mathcal{D}$, and thus we can simplify the model by using solely the x_{gu} variables. Also note that, since we do not perform demand splitting (i.e. no multi-path routing), each link (g, h) shall either carry the whole demand or nothing. Thus, we can correctly formulate the flow conservation constraints by means of binary variables.

The above formulated model is an Integer Linear Programming (ILP) problem. Thus it can be solved by means of a general-purpose Mixed-Integer Linear Programming solver, such as the IBM ILOG CPLEX Optimizer.²

4. Other CDWCN optimisation approaches

In order to provide a measure for the performance of our approach, we consider other techniques that can be applied to a CDWCN system to optimise its operations under different point of views. These are described in the following. Note, however, that the first one (named SHAN) does not perform any optimisation, but is just the currently most diffused approach to content delivery. Therefore it is used as the reference benchmark for the others.

4.1. Single Home Area Networks (SHAN)

This approach represents the current situation of most Internet users, who have their own private (home) network and do not take part in the CDWCN. Therefore, the GWs are not shared and not connected with each other, and every UT downloads the content from the Internet by means of its own GW only. Also, the GWs do not perform content caching.

In this case, the power consumption of the network can be readily computed as:

$$P_{\text{SHAN}} = |\mathcal{G}| \cdot P^{\text{GW}} + E^I \sum_{(c,u) \in \mathcal{D}} b_c. \quad (16)$$

4.2. Internet Sharing (IS)

In this case the users do share their GWs (as in [14,15]), so that each UTs can download the content from the Internet by means of the most convenient GW. Yet, we assume that no caching occurs at the GWs and no mesh network is set up. We thus have an objective function that is aimed at minimising the number of powered-on GWs:

$$\min \sum_{g \in \mathcal{G}} q_g, \quad (17)$$

subject to constraints (2), (3), (9), (10), and (14).

The power consumption of IS is:

$$P_{\text{IS}} = P^{\text{GW}} \sum_{g \in \mathcal{G}} \bar{q}_g + E^I \sum_{(c,u) \in \mathcal{D}} b_c, \quad (18)$$

where \bar{q}_g is the value that q_g takes once the ILP has been solved. Note that, since we assume that all contents are downloaded from the Internet (no caching), there is no point in using the GW–GW links.

4.3. Internet and Cache Sharing (ICS)

This technique enhances the IS approach with caching on the GWs. Still, it does not take advantage of the inter-GW connections. As a result, the objective of ICS is to minimise the following power consumption:

$$P_{\text{ICS}} = P^{\text{GW}} \sum_{g \in \mathcal{G}} q_g + E^I \sum_{(c,u) \in \mathcal{D}} b_c \sum_{g \in \mathcal{G}} (t_{cg} \cdot z_g^{cu}), \quad (19)$$

subject to constraints (2), (3), (5)–(10), (13), and (14).

4.4. Traffic oriented CDWCN (TO)

Under this approach the inter-GW connections can be profitably employed to transfer the cached contents within the CDWCN instead of retrieving them from the Internet. In other terms, TO implements a cooperative caching mechanism [20]. The objective of TO is to minimise the traffic downloaded from the Internet:

$$\min \sum_{g \in \mathcal{G}} \sum_{(c,u) \in \mathcal{D}} t_{cg} \cdot b_c \cdot z_g^{cu}, \quad (20)$$

subject to constraints (2)–(15). Note, however, that such problem formulation would not prevent the formation of loops while taking the routing decisions (for GRARC, it is the objective function (1) that prevents them by minimising the energy consumption). Therefore, a loop-free TO formulation is based on the following objective function:

$$\min \sum_{g \in \mathcal{G}} \sum_{(c,u) \in \mathcal{D}} t_{cg} \cdot b_c \cdot z_g^{cu} + \gamma \sum_{(g,h) \in \mathcal{E}''} \sum_{(c,u) \in \mathcal{D}} y_{gh}^{cu}, \quad (21)$$

where the goal of the second term is to minimise the number of hops for all demand routes (which clearly leads to

² <http://www.ibm.com/software/commerce/optimization/cplex-optimizer>.

loop avoidance), and γ is used to weigh such term so that the primary minimisation objective is not altered. In our case, γ is computed as $\frac{b_c^{min}}{2(|\mathcal{C}|-1)}$, with b_c^{min} being the smallest content rate. Such γ guarantees that the second term is always smaller than the minimum variation of the first one.

The power consumption of the TO technique is:

$$P_{TO} = P^{GW} \sum_{g \in \mathcal{G}} \bar{q}_g + \sum_{(c,u) \in \mathcal{D}} b_c \left\{ E^I \sum_{g \in \mathcal{G}} \bar{z}_g^{cu} t_{cg} + E^W \sum_{(h,g) \in \mathcal{E}''} \bar{y}_{hg}^{cu} \right\}, \quad (22)$$

where, as usual, \bar{z}_g^{cu} and \bar{y}_{hg}^{cu} are the values assigned to the variables by solving the ILP problem. Note that TO does not account for the energy in its objective function, thus it does not care about turning off the GWs.

Also note that, since minimising the traffic downloaded from the Internet might be related with the cost of the Internet connection, TO might also be seen as a cost-minimisation strategy [16].

5. Performance evaluation

The performance of GRARC has been tested over a series of 11 network scenarios and compared to the performance of the techniques presented in Section 4. Each scenario is characterised by different values for the input parameters. The set of parameters and their values are reported in Table 1, with S being the GW cache size and α measuring the content popularity. More information and discussion on the input parameters is given in Section 5.1. Note that the first scenario is used as the reference one: starting from it, we have changed one parameter value per scenario (in bold in the table) in order to estimate its impact on the resource allocation performed by the various techniques. For each scenario, we generated and solved twenty instances, averaging the results.

The factor ρ in (14) is used to account for non-cacheable traffic, such as email or VoIP. By means of ρ we can limit the bandwidth of the GW–UT links used to deliver cacheable contents, thus leaving some spare bandwidth for data traffic that does not rely on the CDN. A first set of results is obtained for the “ideal” case $\rho = 1$ (i.e. no bandwidth limit), whereas a second set is presented for $\rho = 0.65$, a value that is in line with the current and future network scenarios [21]. By comparing the two sets we can deduce the impact of such a “reserved” bandwidth on the performance of the CDWCN.

5.1. Scenario and instance generation method

GWs and UTs have been placed in a fictitious test area of varying size (in order to keep the GW density constant). In each instance, the positions of the GWs and UTs have been randomly determined. However, to guarantee a minimum degree of rationality, the test field has been divided into a regular grid and the GWs and UTs have been evenly distributed across the grid squares. This strategy ensures enough uniformity to mimic a residential neighbourhood

and to avoid heavily unbalanced instances. A sample GW and UT placement in the test area is shown in Fig. 4.

The energy consumed by every GW has been set to $P^{GW} = 15$ W, which is a typical value for residential wireless routers. The energy consumption for retrieving a bit of content via Internet has been set to $E^I = 39$ μ J/b [22], and the energy consumption for transferring a bit of content over a wireless link to either $E^W = 26$ μ J/b [23] or $E^W = 0.02$ μ J/b [24].

It is worth citing the fact that the energy consumption figures available in the literature differ greatly. We have therefore opted for an average value for E^I and P^{GW} , whereas for the wireless connections, which are a peculiar characteristic of the system, we used two extreme values. The higher E^W can be associated to old or energy inefficient devices, whereas the lower E^W is representative of modern and energy efficient equipment.

To determine the GW–GW and GW–UT data rates (i.e. the various r_{gu} and r_{gh}), we employed a simplified version the COST-231 path loss model [25], which allows to account for various propagation aspects, such as the presence of walls and other obstacles, and the use of a realistic path loss exponent. The values of all parameters have been extracted from real measurements [26,27] and data sheets (such as [28]). Then, the data rates have been extracted from the computed signal-to-noise ratio (σ_{ij} , with $i \in \mathcal{G}$ and $j \in \{\mathcal{U} \cup \mathcal{G}\}$) according to:

$$r_{ij} = \min\{\psi_{ij}(\sigma_{ij}), r_{max}\}, \quad (23)$$

where $\psi_{ij}(\sigma_{ij})$ is a proper function that might depend on several parameters, such as the specific modulation and coding schemes and the medium access overhead – in our case we referred to the work of Zhang et al. [29], and r_{max} is used to cap r_{ij} to the maximum rate achievable by the physical link, say $r_{max} = 54$ Mbps (as per the IEEE 802.11g).

5.2. Content model

The content demand has been modelled according to a content popularity that follows the widely adopted Zipf distribution [30]. The probability of demanding the i -th most popular content in \mathcal{C} is computed as:

$$P(X = i) = \frac{1/i^\alpha}{\sum_{j=1}^{|\mathcal{C}|} 1/j^\alpha}, \quad (24)$$

Table 1
Parameter values for the tested scenarios.

Scenario	$ \mathcal{G} $	$ \mathcal{U} / \mathcal{G} $	$ \mathcal{C} / \mathcal{U} $	$S/ \mathcal{U} $	$S \mathcal{G} / \mathcal{U} $	α
1	30	4	1000	2.5	0.075	0.65
2	10	4	1000	2.5	0.025	0.65
3	50	4	1000	2.5	0.125	0.65
4	30	2	1000	2.5	0.075	0.65
5	30	6	1000	2.5	0.075	0.65
6	30	4	100	2.5	0.75	0.65
7	30	4	10,000	2.5	0.0075	0.65
8	30	4	1000	0.25	0.0075	0.65
9	30	4	1000	25	0.75	0.65
10	30	4	1000	2.5	0.075	0.5
11	30	4	1000	2.5	0.075	0.8




Fig. 4. Placement of GWs (red triangles) and UTs (blue discs) for the first instance of the first scenario. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

where the popularity exponent α varies from 0.5 to 0.8 (as shown in Table 1) based on the empirical observation reported in [31,32].

The cached contents in the GWs have been modelled by means of the same Zipf distribution. We reckon that this is a sensible behaviour for a “smart” caching strategy, which would presumably store with the highest probability the most popular contents [30]. For each GW we set a cache of size S and we assumed a number of cached contents equal to the cache size (the system is in the steady state). From the numbers in Table 1, it emerges that S varies from a minimum of 30 (scenario 8) to a maximum of 3000 contents (scenario 9), being the average 300. These values are indeed compatible with home gateways including or attached to a storage unit (e.g. internal hard disc or NAS). We have considered a somewhat cooperative caching strategy, in which the contents to be cached are selected so as the same contents are not cached in adjacent GWs. To have an idea of the probability of finding a content cached somewhere in the CDWCN, the reader can refer to column $S|G|/|U|$ in Table 1, which gives an upper approximation of the ratio of the total cached contents (because some content replication may occur) with respect to the total content population.

Finally, based on [33], we have considered the discrete distribution of the data rate of the contents (b_c) shown in Fig. 5.

5.3. Results with $\rho = 1$

In this section we evaluate the performance figures of GRARC in comparison to the alternative CDWCN solutions presented in Section 4. This first set of results refers to $E^W = 0.02 \mu\text{J}/\text{b}$; a second (smaller) set is presented later on for $E^W = 26 \mu\text{J}/\text{b}$. Both sets have obtained with $\rho = 1$. Further experiments with $\rho = 0.65$ are described in the next Section. The performance of the approaches is assessed in terms of several parameters, which are described in the following.




Fig. 5. YouTube video encoding rate distribution.

The most prominent measure of energy efficiency is the amount of power consumed by the various approaches. Fig. 6 depicts the power consumption of the various techniques normalised to the power consumption of SHAN.

In all scenarios, GRARC achieves the best results. Specifically, GRARC requires on average 21.5% less power than SHAN, with efficiency peaks for scenarios 6 and 9. These are the cases where either the number of contents is low or the cache size is the largest. This means that GRARC is more efficient when more contents (in relationship with the total number of available contents) are cached on the CDWCN, because it can either retrieve them from the local cache (i.e. the cache of the GW the UT is associated with), or fetch them from other GWs by means of the energetically cheap wireless connections.

The net effect of local caching can be seen in the difference among IS and ICS. In the majority of scenarios, ICS performs just slightly better than IS, as it gains an extra 5.2% of power saving. However, when the ratio $|C|/|U|$ is low, the cache size is large, or the popularity exponent is high, caching is more profitable, because it can improve the performance of IS by up to 13.9%. Clearly, this derives from a higher probability of finding a content cached on the local GW.

The energy saving of the IS strategy is roughly constant across all scenarios, with an average value of 10.4%. Of course, IS is insensitive to many of the problem parameters, such as the cache size or the content popularity. The only scenario in which IS gains something more is the number 4, where it can exploit the low number of UTs to switch off more GWs.

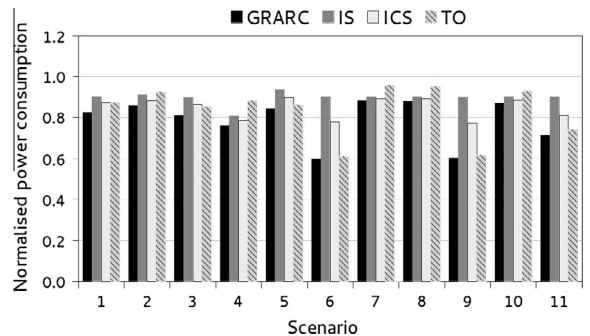


Fig. 6. Total normalised power consumption as a function of the test scenario – $E^W = 0.02 \mu\text{J}/\text{b}$.

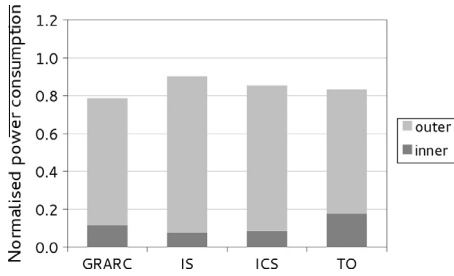


Fig. 7. Power consumption anatomy – $E^W = 0.02 \mu\text{J}/\text{b}$.

The TO approach presents a very similar trend with GRARC, but with definitely worst performance. In many scenarios it can save very little energy with respect to SHAN. As it will be clearer later, the excessive use of the mesh connections to fetch cached contents from distant GWs strongly reduces the gains achieved by the non-use of the Internet links. The sole exceptions are scenarios 6, 9, and (partly) 11, in which TO can exploit the higher probability of finding a content in a nearby GW.

The next figure (Fig. 7) shows from which part of the system the various approaches drain the energy. In detail, with the term “inner” we mean the energy used to power the GWs and to transfer contents over the wireless mesh, whereas the term “outer” indicates the energy necessary for retrieving a content through the Internet connection.

The bars reveal that for all approaches most of the energy is consumed by the Internet. However, we can distinguish two kinds of behaviour. TO puts a considerable share of consumption on the inner part (21.2%), thus being a sort of “selfless” approach, whereas the others (IS, ICS, and to a lesser extent GRARC) limit this quota to much lower values, thus being more “selfish”. More in detail, TO tends to maximise the use of the inter-GW connections, as it turns off no GW. Therefore, this selfless technique exploits the inner part of the CDWCN at the maximum, in order to save energy on the external network. Among the selfish techniques, which do turn off some GWs (look also at Fig. 8), IS is indeed the most selfish one, because it tries to power off as many GWs as possible, thus putting all the energy burden on the Internet. GRARC is somewhere in between these two extremes, as it moves a fair amount of energy to the inside of the network (14.7%, versus 8.4% of IS and 10% of ICS).

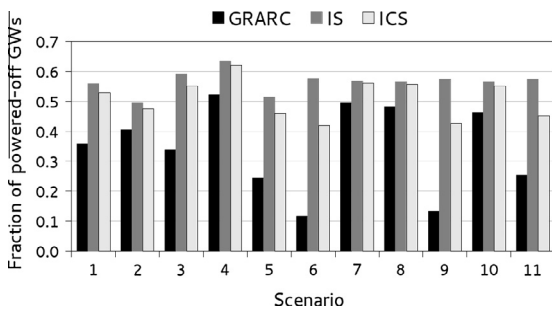


Fig. 8. Fraction of powered-off GWs – $E^W = 0.02 \mu\text{J}/\text{b}$.

Given that GRARC is the most efficient approach (as shown in Fig. 6), the conclusion is that, in order to achieve the globally most energy-efficient behaviour, it is necessary to strike a balance between a completely selfish approach and a more selfless one. The exact trade-off point depends mostly on the ratio between the energy parameters (i.e. E^I , E^W , and P^{GW}). Note, for example, that an upper bound on the maximum number of hops that are energetically convenient can be readily computed as $\lfloor E^I/E^W \rfloor$.

To have a clearer picture of this aspect, we can refer to Fig. 8, which reports the number of GWs switched off. Note that IS represents the maximum number of GWs that can be powered off. The bars for IS and ICS are in many cases quite close to each other, while GRARC have much lower values. However, the gap between GRARC and IS/ICS is largely variable. In some scenarios, it becomes greater, whereas in others the three approaches are relatively close to each other. This last case corresponds to scenarios in which the probability to have a demanded content cached is the smallest. Therefore, both GRARC and ICS have very little margin to improve the energy-efficiency other than turning off as many GWs as possible. Conversely, when that probability is higher (scenarios 6, 9, and 11), both GRARC and ICS can take advantage of the cache to reduce the Internet traffic (and thus the power consumption). In addition, GRARC can also exploit the mesh connections to retrieve the contents from other GWs. That is why GRARC turns off (much) less GWs than ICS: they are needed for both delivering the contents in their cache and to forward the traffic.

Also note that, though GRARC does not power off all the possible GWs (34.6% versus 56.5% of IS), it can nevertheless achieve an overall energy consumption that is lower than IS. This clearly indicates that the selfish approach of minimising the number of powered-on GWs is not the most energy-efficient solution.

Further insights on the features of the various techniques can be gained by means of the hit rate (η), which measures the capability to find and retrieve the demanded contents from within the CDWCN. The hit rate is defined as:

$$\eta = \frac{\sum_{(c,u) \in \mathcal{D}} \sum_{g \in \mathcal{G}} (1 - t_{cg}) \cdot \bar{z}_{cg}^{cu}}{|\mathcal{D}|} \quad (25)$$

Fig. 9 plots the hit rate achieved by all techniques except IS (whose hit rate is always zero, given that IS does not make use of caching). Obviously, TO obtains the highest η values, but GRARC also performs well. In detail, GRARC yields an average η of 18.6, being the one by TO 21.3. For the most favourable scenarios (i.e. 6 and 9) both TO and GRARC can fetch nearly half of the contents from the cache. ICS has a comparable performance only in a few scenarios, which clearly states that the lack of inter-GW connections heavily penalises ICS in those scenarios where the hit rate could be significant.

To complete the hit rate picture, Fig. 10 shows the decomposition of the hit rate in its local and remote parts, where the term “local” refers to contents cached on the GW a UT is currently assigned to ($\bar{x}_{gu} = 1$), and “remote” to contents cached on other GWs in the CDWCN. As




Fig. 9. Hit rate of GRARC, TO, and ICS for the various test scenarios – $E^W = 0.02 \mu\text{J}/\text{b}$.

already anticipated, TO maximises the remote hit rate, closely followed by GRARC. The reason for which GRARC retrieves less contents from the other GWs is that retrieving contents from distant GWs might be more energy consuming than downloading them directly from the Internet (as proved by the fact that GRARC saves much more energy than TO). Obviously, the remote hit rate of ICS is zero, because it does not use the mesh connections.

The last performance metric we report is about the reduction of Internet traffic that each approach can achieve. We thus define the “traffic saving” parameter, θ , as:

$$\theta = \frac{\sum_{g \in \mathcal{G}} \sum_{(c,u) \in \mathcal{D}} (1 - t_{cg}) \cdot b_c \cdot z_g^{cu}}{\sum_{(c,u) \in \mathcal{D}} b_c} \quad (26)$$

Note that θ is normalised to the total traffic necessary to retrieve all demanded contents from the Internet.

The values of θ for the various approaches and scenarios are reported in Fig. 11 (apart from IS, whose θ is always zero). The general trend is quite similar to that of Fig. 9, but the bars of GRARC and TO are now slightly lower. This comes from these approaches fetching contents with a smaller than average b_c , especially TO. Indeed many “light” contents are easier and more efficient to route over the mesh than few “heavy” contents, and therefore, in order to exploit the mesh, GRARC and TO tends to favour lighter contents.

The impact of a higher E^W (26 $\mu\text{J}/\text{b}$) on the performance of the approaches can be analysed by looking at Figs. 12–15. Starting from Fig. 12, which reports the normalised power consumption, we can note that IS and ICS achieve very similar results to the previous case, whereas

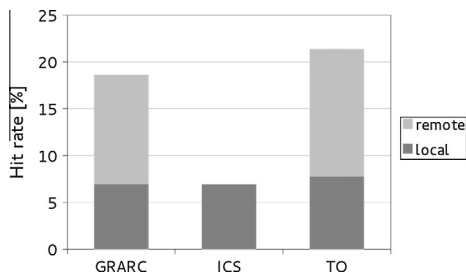


Fig. 10. Hit rate decomposition – $E^W = 0.02 \mu\text{J}/\text{b}$.

the performance of GRARC and TO is greatly reduced. In particular, GRARC behaves almost like ICS. Going to the sheer numbers, GRARC yields an average power consumption of 0.843 instead of 0.785. The performance of TO degrades even more, to the point that it often leads to consuming even more power than the reference SHAN case. Clearly, since TO’s strategy is not energy-driven, but only traffic-driven, this method exploits the mesh network at its maximum, retrieving cached contents even from the most distant GWs. As the energy consumption of the wireless connections is not so convenient with respect to downloading directly from the Internet, TO ends up with causing an energy wastage.

This thesis is confirmed also by Fig. 13, which shows the fractions of inner and outer power consumptions. TO is the sole method whose inner power consumption is considerably higher than the average, being the behaviour of GRARC more similar to the one of IS and, above all, ICS.

The impact of E^W on the number of active GWs is illustrated in Fig. 14. Again, the behaviour of GRARC is now very close to ICS. In fact, GRARC turns off more GWs (49.2% instead of 34.6%), and reaches almost the same level of ICS (51.3%). This means that GRARC heavily diminishes the amount of contents retrieved from other GWs. A proof of that can be found in Fig. 15, which shows how the hit rate of GRARC goes down towards the values of ICS.

In summary, given the high values of E^W , with respect to E^I , the use of the mesh network is now much less convenient, and therefore a simpler approach, such as ICS, might also be acceptable from an energy point of view.

The last performance figure we present is about the CPU times for solving GRARC, which are resumed in Table 2. The computation times differ greatly on the basis of E^W and of the problem “size”, i.e. the values of $|\mathcal{G}|$ and $|\mathcal{U}|/|\mathcal{G}|$ (which have also been reported in the table for convenience). For $E^W = 0.02 \mu\text{J}/\text{b}$ GRARC is much slower in finding a solution, given that the search space is much larger than for $E^W = 26 \mu\text{J}/\text{b}$. Indeed, in the latter case the use of the wireless connections is convenient up to $E^I/E^W = 1$ hop, so the search for convenient GWs is limited to direct neighbours only. Conversely, for $E^W = 0.02 \mu\text{J}/\text{b}$, all GWs are potential candidates for fetching the contents. Also, much less routing decisions must be taken for $E^W = 26 \mu\text{J}/\text{b}$. As for the problem size, it clearly appears that the computation times heavily depend on the number of GWs and UTs. The smallest scenarios are solved faster and the largest much more slowly, in an almost exponential dependence on $|\mathcal{G}|$ and $|\mathcal{U}|/|\mathcal{G}|$. The other parameters play a minor role in the computational complexity.

5.4. Results with $\rho = 0.65$

The effect of the parameter ρ on the various approaches has been analysed by performing a series of computational tests with $\rho = 0.65$. The difference with the case $\rho = 1$ can be seen in Table 3, which reports the power consumption, averaged over all scenarios, for both the examined E^W values. It can be seen that, for all methods, there is an almost negligible increment in the power consumption. All other




Fig. 11. Traffic saving of GRARC, TO, and ICS for the various test scenarios – $E^W = 0.02 \mu\text{J}/\text{b}$.

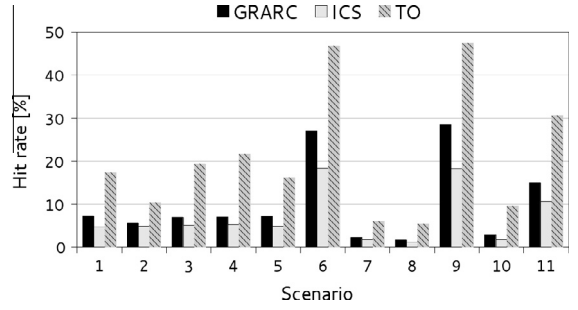


Fig. 15. Hit rate of GRARC, TO, and ICS for the various test scenarios – $E^W = 26 \mu\text{J}/\text{b}$.




Fig. 12. Total normalised power consumption as a function of the test scenario – $E^W = 26 \mu\text{J}/\text{b}$.




Fig. 13. Power consumption anatomy – $E^W = 26 \mu\text{J}/\text{b}$.




Fig. 14. Fraction of powered-off GWs – $E^W = 26 \mu\text{J}/\text{b}$.

performance metrics, such as fraction of powered-off GWs, hit rate, and traffic saving, have also a very similar trend to the power consumption. The numbers can be seen in Tables 4–6.

Table 2
CPU times (in seconds) for solving GRARC.

Scenario	$E^W = 0.02 \mu\text{J}/\text{b}$	$E^W = 26 \mu\text{J}/\text{b}$	$ \mathcal{G} $	$ \mathcal{U} / \mathcal{G} $
1	405	13.5	30	4
2	0.7	0.4	10	4
3	13,541	221	50	4
4	59	3.2	30	2
5	918	117	30	6
6	875	8	30	4
7	179	26.2	30	4
8	255	16.5	30	4
9	1160	7.5	30	4
10	210	36.2	30	4
11	355	11.5	30	4

Table 3

Total average normalised power consumption for the cases $\rho = 0.65$ and $\rho = 1$.

	$E^W = 0.02 \mu\text{J}/\text{b}$		$E^W = 26 \mu\text{J}/\text{b}$	
	$\rho = 0.65$	$\rho = 1$	$\rho = 0.65$	$\rho = 1$
GRARC	0.794	0.785	0.853	0.843
IS	0.905	0.896	0.905	0.896
ICS	0.860	0.849	0.860	0.849
TO	0.840	0.836	1.015	1.009

Table 4

Average fraction of powered-off GWs for the cases $\rho = 0.65$ and $\rho = 1$.

	$E^W = 0.02 \mu\text{J}/\text{b}$		$E^W = 26 \mu\text{J}/\text{b}$	
	$\rho = 0.65$	$\rho = 1$	$\rho = 0.65$	$\rho = 1$
GRARC	0.311	0.346	0.455	0.492
IS	0.515	0.565	0.513	0.563
ICS	0.474	0.509	0.473	0.513
TO	0	0	0	0

Table 5

Average hit rate [%] for the cases $\rho = 0.65$ and $\rho = 1$.

	$E^W = 0.02 \mu\text{J}/\text{b}$		$E^W = 26 \mu\text{J}/\text{b}$	
	$\rho = 0.65$	$\rho = 1$	$\rho = 0.65$	$\rho = 1$
GRARC	18.2	18.6	9.53	10.1
IS	0	0	0	0
ICS	6.06	6.92	6.18	6.94
TO	20.7	21.3	20.2	20.9

Table 6Average traffic saving for the cases $\rho = 0.65$ and $\rho = 1$.

	$E^W = 0.02 \mu\text{J}/\text{b}$		$E^W = 26 \mu\text{J}/\text{b}$	
	$\rho = 0.65$	$\rho = 1$	$\rho = 0.65$	$\rho = 1$
GRARC	0.179	0.182	0.097	0.100
IS	0	0	0	0
ICS	0.063	0.069	0.063	0.069
TO	0.195	0.201	0.191	0.195

6. Concluding remarks

From the performance analysis of our GRARC method and its comparison with other resource sharing approaches, several considerations can be drawn. At first, GRARC can yield a substantial energy saving for both the CDWCN users and the Internet. In our tests it achieved an average improvement of 21.5% (with peaks as high as 40%) with respect to the Single Home Area Network (SHAN) model. The combination of Internet and Cache Sharing (ICS) can also be profitable, as its energy savings are sometimes close to those of GRARC. However, ICS loses ground in those scenarios where the hit rate could be significant, i.e. exactly in those cases where a good caching technique could have been enforced. Also, ICS generates more Internet traffic than GRARC. This may make a difference in those cases when Internet bandwidth is a precious resource, e.g. in terms of money or due to a narrowband Internet connection. In these cases, TO offers the best solution, as its application ensures the maximum reduction of Internet traffic. However, since TO is not energy-aware, it might also cause higher power expenditure than SHAN. On the other hand we have proved how GRARC can constitute a viable trade-off approach, because it is more efficient than TO in choosing the contents to be retrieved from the other GWs. To complete the review, the selfish IS approach of maximising the number of switched-off GWs can offer some energy saving to the CDWCN users, but this is neither the most energy-efficient solution, nor a fair one, as it puts all the energy burden on the Internet. On the contrary, keeping more GWs active can yield a higher overall gain if smart resource allocation techniques such as GRARC are employed.

Clearly, the performance of GRARC depends on many parameters, such as the ratio of the cached contents (S) to the total number of available and requested contents, and the ratio between the energy consumption for retrieving a bit from the Internet (E^I) and the energy for transmitting a bit over a wireless link (E^W). Specifically, GRARC is more efficient when $S/|C|$ and/or E^I/E^W are high, because it can take full advantage of the wireless mesh network.

In perspective, if the current trend of improvement in the spectrum efficiency of the wireless connections keeps going at a faster pace than energy awareness in the wired Internet, it is easy to envision even higher energy gains from the adoption of GRARC or similar techniques. Also, the reduction in the power consumption of home gateways and the possibility to have large storage options at cheap prices, may give further strength to GRARC. In such a scenario, keeping the GWs active would be less penalising in

terms of energy, and allow for a higher amount of contents to be retrieved from within the CDWCN. Also note that reducing the Internet traffic is relevant both for network providers, as it translates in lower costs, and for the applications, since these can enjoy better quality of service in terms of throughput and delay.

As the next steps to fully exploit the benefits promised by GRARC, it will be necessary to adapt it to the dynamic context of actual deployments. Therefore, future works must be targeted at creating a version of GRARC that can be implemented in the central controller where the allocation takes place. Specifically, from the purely modelling perspective, more elements should be accounted for, such as the size of the contents, the arrival and termination of requests, the possibility of re-routing some requests. From a more practical point of view, future activities might include the definition of more efficient solving techniques, including heuristics, to enable very fast solving times, and the integration with content deployment strategies.

Acknowledgments

This work was supported by the GATECOM and Green-Net research projects, both funded by the Italian Ministry of Research and University (MIUR).

The authors would like to thank Dr. Loris Gazzarrini for his help in setting up a fast and efficient instance generator.

The authors would also like to thank the anonymous reviewers for their very helpful suggestions.

References

- [1] G. Pallis, A. Vakali, *Insight and perspectives for content delivery networks*, *Commun. ACM* 49 (2006) 101–106.
- [2] S. Spagna, M. Liebsch, R. Baldessari, S. Niccolini, S. Schmid, R. Garroppo, K. Ozawa, J. Awano, *Design principles of an operator-owned highly distributed content delivery network*, *IEEE Commun. Mag.* 51 (2013) 132–140.
- [3] C. Szabó, K. Farkas, Z. Horváth, *Motivations, design and business models of wireless community networks*, *Mobile Networks Appl.* 13 (2008) 147–159.
- [4] D. Han, D. Andersen, M. Kaminsky, D. Papagiannaki, S. Seshan, Hulu in the neighborhood, in: *Third International Conference on Communication Systems and Networks (COMSNETS)*, 2011, <http://dx.doi.org/10.1109/COMSNETS.2011.5716501>.
- [5] R. Canonico, C. Guerrero, A. Mauthe, *Content distribution infrastructures for community networks*, *Comput. Networks* 53 (2009) 431–433.
- [6] F. den Hartog, B. Bastiaans, M. Blom, M. Pluijmaekers, R. van der Mei, *The use of residential gateways in content delivery networking*, in: *Proc. Australian Telecommunication Networks and Applications Conference (ATNAC)*, Sydney, 2004.
- [7] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, P. Rodriguez, *Greening the internet with nano data centers*, in: *Proc. 5th International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2009, pp. 37–48, <http://dx.doi.org/10.1145/1658939.1658944>.
- [8] V. Manetti, R. Canonico, W. De Donato, G. Ventre, A. Mauthe, G. Tyson, *Next generation CDN services for community networks*, in: *Proc. Third International Conference on Next Generation Mobile Applications, Services and Technologies (NGMAST)*, 2009, pp. 89–94, <http://dx.doi.org/10.1109/NGMAST.2009.45>.
- [9] A. Alasaad, S. Gopalakrishnan, H. Nicanfar, V. Leung, *Green content distribution in wireless mesh networks with infrastructure support*, in: *IEEE International Conference on Communications (ICC)*, 2012, pp. 5896–5900, <http://dx.doi.org/10.1109/ICC.2012.6364736>.

- [10] S. Androutsellis-Theotokis, D. Spinellis, A survey of peer-to-peer content distribution technologies, *ACM Comput. Surv.* 36 (2004) 335–371.
- [11] G. Aggelou, *Wireless Mesh Networking*, McGraw-Hill Professional, 2008.
- [12] S. Borst, V. Gupta, A. Walid, Distributed caching algorithms for content distribution networks, in: Proc. IEEE INFOCOM, 2010, <http://dx.doi.org/10.1109/INFOCOM.2010.5461964>.
- [13] Z. Al-Arnaout, Q. Fu, M. Frea, A divide-and-conquer approach for content replication in WMNs, *Comput. Networks*, 57 (2013) 3914–3728.
- [14] C. Rossi, C. Casetti, C.F. Chiasserini, An Energy Efficient Protocol for Gateway-Centric Federated Residential Access Networks, *Computing Research Repository (CoRR)*, 2011. <http://arxiv.org/abs/1105.3023>.
- [15] E. Goma, M. Canini, A. Lopez Toledo, N. Laoutaris, D. Kostić, P. Rodriguez, R. Stanojević, P. Yagüe Valentin, *Insomnia in the access: or how to curb access network related energy consumption*, *Comput. Commun. Rev.* 41 (2011) 338–349.
- [16] M. Taghizadeh, K. Micinski, S. Biswas, C. Ofria, E. Torng, Distributed cooperative caching in social wireless networks, *IEEE Trans. Mobile Comput.* 12 (2013) 1037–1053.
- [17] U. Shevade, Y.-C. Chen, L. Qiu, Y. Zhang, V. Chandar, M.K. Han, H.H. Song, Y. Seung, Enabling high-bandwidth vehicular content distribution, in: Proc. 6th International Conference on emerging Networking EXperiments and Technologies (CoNEXT), ACM, New York, NY, USA, 2010, pp. 23:1–23:12, <http://dx.doi.org/10.1145/1921168.1921199>.
- [18] B. Stiller, D. Hausheer, T. Hofeld, Towards a socially-aware management of new overlay application traffic combined with energy efficiency in the internet (SmartenIT), in: A. Galis, A. Gavras (Eds.), *The Future Internet, Lect. Notes Comput. Sci.*, vol. 7858, Springer, Berlin Heidelberg, 2013, pp. 3–15, http://dx.doi.org/10.1007/978-3-642-38082-2_1.
- [19] M. Seufert, V. Burger, T. Hossfeld, HORST – home router sharing based on trust, in: Network and Service Management (CNSM), 2013 9th International Conference on, 2013, pp. 402–405, <http://dx.doi.org/10.1109/CNSM.2013.6727865>.
- [20] J. Zhao, P. Zhang, G. Cao, C. Das, Cooperative caching in wireless P2P networks: design, implementation, and evaluation, *IEEE Trans. Parallel Distrib. Syst.* 21 (2010) 229–241.
- [21] Cisco Systems, Inc., *The Zettabyte Era – Trends and Analysis*, White Paper, 2013.
- [22] J. Baliga, R. Ayre, K. Hinton, W.V. Sorin, R.S. Tucker, Energy consumption in optical IP networks, *J. Lightwave Technol.* 27 (2009) 2391–2403.
- [23] R. Balani, Energy Consumption Analysis for Bluetooth, Wi-Fi and Cellular Networks, NESL Technical Report TR-UCLA-NESL-200712-01, University of California, Los Angeles, 2007.
- [24] D. Halperin, B. Greenstein, A. Sheth, D. Wetherall, *Demystifying 802.11n power consumption*, in: Proc. International Conference on Power Aware Computing and Systems, USENIX Association, 2010.
- [25] European Commission, COST 231 – Digital Mobile Radio Towards Future Generations Systems, Final Report, 1999.
- [26] R. Wilson, Propagation Losses Through Common Building Materials: 2.4 GHz vs 5 GHz, Technical Report E10589, Magis Networks, Inc., 2002.
- [27] R. de Francisco, Indoor channel measurements and models at 2.4 GHz in a hospital, in: Proc. IEEE Global Telecommunications Conference (GLOBECOM), 2010, pp. 1–6, <http://dx.doi.org/10.1109/GLOCOM.2010.5683218>.
- [28] Cisco Systems, Inc., Aironet 3600 Series Access Point, 2013. <<http://www.cisco.com/en/US/products/ps11983/index.html>>.
- [29] J. Zhang, K. Tan, J. Zhao, H. Wu, Y. Zhang, A practical SNR-guided rate adaptation, in: Proc. IEEE Conference on Computer Communications (INFOCOM), 2008, pp. 2083–2091, <http://dx.doi.org/10.1109/INFOCOM.2008.274>.
- [30] G. Haßlinger, F. Hartleb, Content delivery and caching from a network provider's perspective, *Comput. Networks* 55 (2011) 3991–4006.
- [31] P. Gill, M. Arlitt, Z. Li, A. Mahanti, Youtube traffic characterization: a view from the edge, in: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, IMC '07, 2007.
- [32] X. Cheng, C. Dale, J. Liu, Statistics and social network of YouTube videos, in: Proc. 16th International Workshop on Quality of Service (IWQoS), 2008.

- [33] P. Ameigeiras, J.J. Ramos-Munoz, J. Navarro-Ortiz, J. Lopez-Soler, Analysis and modelling of YouTube traffic, *Trans. Emer. Telecommun. Technol.* 23 (2012) 360–377.



Rosario G. Garroppo received the Laurea (M.Sc.) degree “cum laude” in Telecommunication Engineering in 1995 and the Ph.D. degree in Information Engineering in 1999 from the University of Pisa. Since December 2001, he is Assistant Professor at the University of Pisa. His research interests are in traffic modeling, QoS guarantee in mobile networks, performance evaluation of telecommunication networks by means of simulation tools, traffic control in wireless networks, green networking.



Gianfranco Nencioni received the Master's degree in Telecommunication Engineering “cum laude” from the University of Pisa in 2008, and the Ph.D. in “Ingegneria dell'Informazione” from the University of Pisa in 2012. In the fall of 2011, he was a visiting Ph.D student at the Computer Laboratory of the University of Cambridge, supervised by Prof. Jon Crowcroft. He is currently a postdoctoral fellow at the Department of Ingegneria dell'Informazione of the University of Pisa. His research activity regards energy-aware routing and network design in both wired and wireless environments.



Luca Tavanti received the M.Sc. degree in Telecommunications Engineering in 2000 and the Doctorate degree in Information Engineering in 2007, both from the University of Pisa. From 1999 to 2003 he was first with BT Labs in Ipswich (UK), and then with Simtel spa in Florence (Italy). Since January 2007 he is a research fellow (non tenure track) at the Dept. of Ingegneria dell'Informazione of the University of Pisa. He has authored many journals and conference articles mostly in the area of wireless networks with main focus on the issues of routing, quality of service (QoS), and efficient resource allocation.



Bernard Gendron is a Professor at the Department of Computer Science and Operations Research at the University of Montréal, Canada. His research interests include integer and combinatorial optimization, large-scale optimization and parallel computing, and location and network design problems applied to transportation logistics and telecommunications. He is affiliated to the Centre for Research on Transportation at the University of Montréal and has held positions of Visiting Professor at MIT, EPFL, Pisa, Versailles and Valenciennes. During his recent sabbatical (2003), he worked at ILOG, Paris. He is Chair of the Montreal Chapter of the Canadian Operational Research Society (CORS), and has served as Chair of CORS (2004–2005) and Chair of the Section on Transportation Science & Logistics of INFORMS (2002–2003). He was awarded the CORS Practice Prize (2004) and the CORS Service Award (2006).