Computer Networks 57 (2013) 1597–1610

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Dynamic green self-configuration of 3G base stations using fuzzy cognitive maps

Christian Facchini [a,b,*], Oliver Holland [b], Fabrizio Granelli [a], Nelson L.S. da Fonseca [c], Hamid Aghvami [b]

[a] Department of Information Engineering and Computer Science, University of Trento, Italy
[b] Centre for Telecommunications Research, King's College London, United Kingdom
[c] Institute of Computing, State University of Campinas, Brazil

ARTICLE INFO

Article history:
Received 25 August 2012
Received in revised form 29 December 2012
Accepted 9 February 2013
Available online 16 February 2013

Keywords:
Cognitive network
Fuzzy cognitive map
Green communications

ABSTRACT

Cellular networks are rapidly evolving towards the fourth generation, thus providing a global infrastructure for wideband mobile network access. Currently, most of the energy consumption of such technology is by cellular base stations, which are not energy efficient—at least in terms of the transmission energy to "from-the-socket" energy consumption ratio. This paper addresses the problem of energy efficiency in cellular networks by taking advantage of the principles of cognitive networking, which promotes the creation of intelligent networks capable of self-configuration with minimal human intervention. In particular, this paper uses the concept of fuzzy cognitive maps to decide upon opportunistic traffic and user reallocations between radio network equipment operating in different spectrum bands to enable power saving modes by some subsets of the radio network equipment, and to utilize spectrum of more appropriate propagation characteristics to save transmission energy. The feasibility and performance of the proposed approach is investigated through simulations. Significant energy savings of some 25–30% are shown over a 72-h period, and blocking rate under the concept is shown to remain reasonable albeit exhibiting a high variance.

1. Introduction

Global emissions of greenhouse gases (GHGs) represent a footprint on the world of the development of humanity, especially after industrialization. The ICT sector itself roughly accounts for 2% of today's global carbon footprint [1], but figures are expected to significantly increase in forthcoming years, with forecasts predicting levels in 2020 around three times what they were in 2002. However, ICT is also forecast to contribute, both directly and indirectly, to reducing global emissions of about five times its own footprint, potentially leading to approximately €600 billion savings [2] in energy costs.

The most significant direct effect is that the telecom infrastructure is expected to grow significantly, ultimately being responsible for 13% of the total sector footprint. Considering also that power generation in ICT is acknowledged to be one of the main causes behind the increase of man-made greenhouse gases, it is evident the importance of energy optimization in the telecom infrastructure.

In current cellular networks, base stations are usually kept powered on and operating all day long, pursuing the vision of an "always-on" network. As power consumption in such networks is mainly due to base stations, which account for almost 80% of the total [3], it is no wonder that

* Corresponding author. Address: Department of Information Engineering and Computer Science, University of Trento, via Sommarive 14, I-38123 Trento, Italy. Tel.: +39 0461 282062; fax: +39 0461 2093.

E-mail addresses: facchini@disi.unitn.it (C. Facchini), oliver.holland@kcl.ac.uk (O. Holland), granelli@disi.unitn.it (F. Granelli), nfonseca@ic.unicamp.br (N.L.S. da Fonseca), hamid.aghvami@kcl.ac.uk (H. Aghvami).

URL: http://disi.unitn.it/facchini (C. Facchini).

1389-1286/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2013.02.011

several research efforts have tried to intervene directly on the functioning of transceivers, at different levels of details.

Most of the works in the literature address the problem of energy consumption from a static point of view, by applying static optimization algorithms[3–5]. While this approach works fine in quasi-static or highly predictable scenarios, the same may not hold in cases where the environment is subject to dynamic changes or in cases of emerging and differentiated usage of the communication infrastructure. For example, even if it is known that the average number of users in a cell follows a predictable pattern, the actual number may be markedly different from the mean, thereby reducing the effectiveness of energy saving mechanisms implementing a static approach. In line with such considerations, the paper aims to explore the possibility of developing dynamic energy-reduction schemes.

In this context, the cognitive networking paradigm represents a potential approach to pursuing the objective of energy efficiency. Cognitive networking is a relatively recent research field, stemming from cognitive radio technology. It aims to extend the principles underlying cognitive radios and apply them to the whole communication protocol stack according to a network-wide perspective [6].

Works in this area are traditionally targeted at reducing management complexity or optimizing Quality-of-Service-related metrics. In this article, we propose to adapt cognitive networking principles to address the problem of energy saving in cellular networks. More precisely, we employ a mathematical tool called Fuzzy Cognitive Maps (FCMs) to analyze causal relationships between energy consumption in base stations and other variables characterizing the cell, in order to identify the most appropriate run-time decisions to reduce energy consumption while maintaining a suitable performance level. It should be noted that a dynamic approach such as the one proposed in this paper can, in principle, substitute "static" schemes, but can also be used in conjunction to them.

The rest of the paper is structured as follows. In Section 2 we review related works. Section 3 is devoted to the description of the proposed cognitive architecture, which is analyzed in detail in Section 4. We validate our scheme through simulations: the simulation scenario together with the details of simulations are described in Section 5, while results are presented and discussed in Section 6. A discussion of real-world implementation of the proposed system is presented in Section 7, and our concluding remarks are offered in Section 8.

2. Related works

As mentioned in the introductory section, several works try to cope with the problem of energy consumption by turning on and off network components—base stations, radio modules, parts of the network itself.

For instance, Tipper et al. [4] observe that powering down transceivers may not be convenient in certain scenarios, e.g. for security reasons or because it does not comply with regulatory constraints, and advance the idea to *dim* cellular networks. Specifically, they propose to lower base stations transmission power, diminish the number of frequency slots available, and reduce high data rate services so as to achieve, in order, coverage, frequency, and service dimming.

Other works aim to reduce energy consumption by turning on and switching off base-station based on the traffic load of a cell [5,7]. Such approach is based on the observation that the traffic load in real-world cellular networks alternates busy periods with quiet periods in a periodic fashion, to the point that it can be approximated with a sinusoid [8]: during low-peak traffic periods the system is underutilized and energy can be saved by switching off inactive base stations, provided that coverage is guaranteed by neighboring cells. This scheme is shown to potentially reduce energy consumption of about 25–30% [7].

Following an even finer degree of control, Saker et al. illustrate a scheme to save energy by reducing the number of active transceivers based on the current traffic load [3]. Results show that by implementing sleep modes in a mixed 2G/3G network, it is possible to save up to 66% of the power used in a traditional network, still being able to retain a blocking rate as low as 0.2%.

Our work differs from [5,7,9] in that it focuses on a single base station, rather than multiple base stations in a network. We also avoid putting to sleep a base station in its entirety. Instead, we allow a finer degree of control by activating or deactivating subsets of the transceiving modules that compose the base station. Similarly to [4], we consider powering off radio modules at the higher frequencies if the number of customers is small enough to exclusively fit the lower band. However, we assume that radio coverage remains unchanged as we do not allow transmission power to be dimmed. Instead we allow the base station to switch operational mode, from omni-directional to tri-sectorized and viceversa, depending on the context. In general, differently from all the mentioned approaches, our scheme aims to independently adapt to the context variations, using minimal a priori information, while also discovering cause-effect relationships among the variables constituting the problem.

Only a few works exist that merge together the cognitive networking paradigm and green communications. One such example is the architecture developed within the End-to-End Efficiency (E^3) European project [10]. Although the E^3 architecture aims primarily to maximize spectrum and radio resources utilization while reducing configuration complexity, it can in principle optimize the power consumption thanks to long- and short-term decisions taken by different modules (Dynamic Self-Organizing Network Planning and Management module and Self-x for Radio Access Networks module) to be installed in the network.

However, rather than with respected to the E^3 architecture, it is more appropriate to position our approach with respect to the reasoning techniques employed in the approaches advanced by the cognitive networking research community. Examples of reasoning formalisms proposed thus far in the literature include neural networks [11], Bayesian and Markov networks [12], and optimization algorithms in general [6].

The work carried out in this paper is also different from what presented in [13]. Simulations in the cited paper assume that operations related to energy saving automatically take place every time a user terminal joins or leaves a cell. In this paper we introduce augmented reasoning capabilities for the base station, so that more refined decisions can be taken, by considering not only terminals joining or leaving the cell, but also others, such as blocking rate and signal to noise interference.

We use Fuzzy Cognitive Maps (FCMs), a tool that makes reasoning to be based on cause-effect dependencies and that already yielded encouraging results [14]. FCMs are graphically represented through directed labeled graphs, in which edges are causal relationships that tie together two variables (technically referred to as "concepts"). This is an advantage over neural networks, which also can be seen are a directed graph of different variables: the conformation of the edges in an FCM reflects the real dependencies between the variables, whilst that of a neural network does not necessarily show it [15].

Bayesian networks, which are as well capable of representing cause–effect relationships among variables of a given problem, and Markov networks share a common major limitation: the difficulty to deal with causality loops. Bayesian networks are based on directed acyclic graphs and cannot be applied to such problems at all. Markov networks can employ approximate algorithms when causality loops are present, yet are not guaranteed to converge. On the contrary, the inference process in FCMs works independently from the presence of loops in the structure.

A more detailed description of FCMs will be provided in Section 4.

3. The proposed energy-efficient architecture

This section describes the architecture proposed by the authors. For clarity, Section 3.1 describes the main modules of the proposed cognitive architecture, while Section 3.2 explores into more detail the employed energy saving mechanisms.

3.1. The proposed architecture

Though the approach we propose can in principle be applied to any network, we focus on the Universal Terrestrial Radio Access Network (UTRAN), and use parameters that refer to the High Speed Downlink Packet Access (HSDPA) protocol.

A UTRAN is composed of two types of element: the Radio Network Controller (RNC) and the Node B. The RNC is in charge of controlling the Nodes B that are linked to it and managing the available radio resources. Examples of functions carried out by the RNC are handover control, admission control and power control. The Node B, also referred to as base station,[1] acts as a transceiver and allows the terminals (User Equipments, technically) to access the core network. It performs low level functions such as signal processing, modulation, and diversity combination.

The cognitive cycle Fig. 1 is an abstraction that represents the fundamental activities to be carried out by a cognitive entity [16]. The environment is sensed, and the information collected enables the cognitive entity to reason, i.e. to assess the possible actions that can be taken (plan) and finalize a decision (decide). Finally, action takes place, and the environment is again sensed so as to evaluate what effects are produced.

Fig. 2 analyzes how such steps can be embedded in the UTRAN. As Node B acts as an interface between the terminals and the network, it is the appropriate device to perform environmental monitoring. Reasoning and learning, however, are best allocated at the RNC level. The main reason is that a single RNC can potentially drive multiple base stations and can thus aim at global optimization, which would otherwise be not feasible. Ultimately, actions can be undertaken only by Node B, properly instructed by the corresponding RNC. According to our model a new instance of the reasoning formalism is created in the RNC for each Node B driven by it. Although it would be more correct to refer to the pair Node B/RNC, we will mention almost exclusively Node B as the cognitive system in the remainder of the paper.

Actions, in the problem posed in this paper, deal with energy saving and are analyzed in detail in Section 3.2, whereas the definition of a reasoning entity suitable for installation on the RNC is described in Section 4. In particular, Section 4.1 describes how the information collected by sensing modules and the actions of acting modules can be transformed into concepts that can be used in the reasoning formalism.

3.2. Energy saving mechanisms

The energy saving techniques considered in this work are applicable to scenarios in which networks operate in different frequency bands covering the same geographical area. Such scenarios are almost routinely the case in many areas around the world, and their proliferation is likely to increase in the future with new technologies and addi-

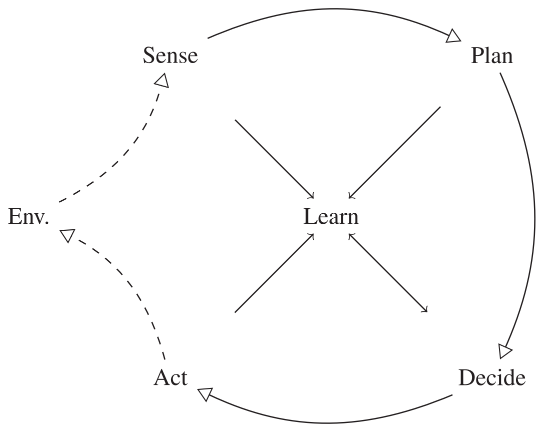

Fig. 1. The cognition loop (adapted from [16]). Highlighted, the direct connection that the acting and sensing stages share with the surrounding environment.

[1] The terms "Node B" and "base station" will be used interchangeably throughout the remainder of the text.

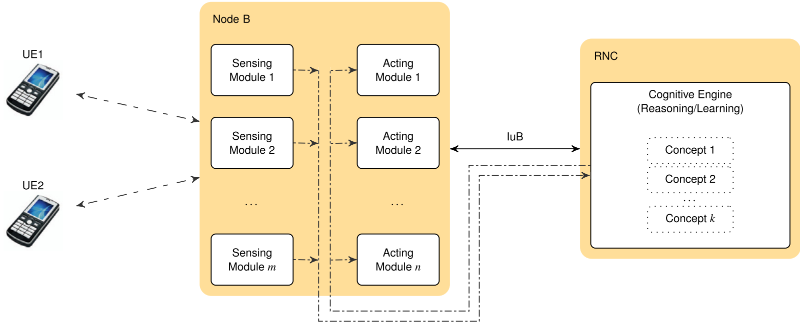

Fig. 2. How the proposed architecture fits in a UTRAN. Node B monitors the environment and sends the data to the Radio Network Controller (RNC), which upon reasoning and learning, drives the acting modules in Node B, thereby closing the cognitive loop. For every Node B an instance of the cognitive formalism is present in the RNC.

tional bands, such as the earmarked IMT-Advanced bands, coming into operation [30].

In this context, two energy saving techniques are investigated. Both are based on the dynamic redistribution of traffic load or users between bands, but aim at achieving different effects [17]. The first aims to allow radio network equipment in the bands that the users originated from to be switched off or put into stand-by mode. The second, aims to optimize propagation characteristics for the given scenario, dependent on traffic load, shadowing, and other factors.

Said objectives are not conflicting with each other. As the following subsections show, the techniques presented operate on slightly different aspects – one aims at using space in a single band or sector and the other aims at employing lower frequencies – and for this reason they can run in parallel.

3.2.1. Opportunistic reallocation to power down radio network equipment

Two scenarios are considered: (i) the opportunistic reallocation of all traffic load or users from a band to another band, in order to allow the band that the users or traffic load originated from to be entirely switched off and (ii) the opportunistic reallocation of a sufficient number of users or traffic load to allow the cell to operate in omnidirectional mode instead of sectorized mode, while still adequately carrying all offered traffic. Power consumption dependence on transmission power is limited in most radio base stations, so any increase in required transmission power due to, e.g., a reduction in antenna gain in omnidirectional mode, is not likely to have a very significant effect on overall power consumption. Given this, switching off the largest number of radio chains possible, for the sectorization switching example through reducing sectors from, e.g., three (sectorized) to one (omnidirectional), is usually considered a reasonable approach to achieve energy saving. Besides, though not investigated in this paper, it is inherently clear that entirely opportunistic switching off the cell is also a good solution, should appropriate measures be taken to avoid any negative consequences of such an action: in fact, switching off a base station leads to the maximum possible energy saving but, at the same time, does not allow any terminal to connect to the network. Connectivity, in such case, may be granted by neighboring cells, coordinating with one another to provide coverage yet minimizing interference.

This particular solution is clearly suited to locations and times where the deployed networks' capacity is higher than traffic load. In many cases, it is necessary to deploy a considerably large network capacity to cover the peak hour load for example, but at other times all such capacity might not be needed. In such cases, the extra capacity could be switched off in order to save energy. This paper considers such capacity being provided across multiple bands, and theorizes the ability to re-allocate users between bands to allow network equipment to be switched off.

3.2.2. Opportunistic reallocation to improve propagation

This solution is based on the ability to opportunistically re-allocate users or traffic loads to lower frequency spectrum to improve propagation, hence reducing necessary transmission power. This, however, will often have the negative effect of increasing inter-cell interference in frequency reuse scenarios, if the density of base stations is high.

Such a possibility can nevertheless be mitigated by considering some specific factors when finding the optimal spectrum available to perform the re-allocation. Such factors are, for instance, the traffic area-density, required base station density and propagation distance, and the frequency-dependent propagation characteristics in the locality. It is noted that in a number of cases the final result may still be the same, i.e., that it is preferable to allocate to the lower frequency spectrum opportunistically, when possible.

This solution is applicable like in the previous case at times when spectral capacity is in excess. Such cases might be operable in non-busy periods or office hours in a business district scenario, or at vacation times.

4. Embedding fuzzy cognitive maps in radio network controllers

Devised in the 1980s as a mathematical tool to help experts discover causal implications in social science problems [18], Fuzzy Cognitive Maps (FCMs) have been applied to different domains, from the simulation of virtual worlds [19] to the creation of medical decision support systems [20].

Nodes in such a graph are called concepts and can potentially represent any variable that characterizes the system that is being modeled, such as for instance "the amount of users in a cell at a given time" or "the energy consumed by the base station". The labeled edges express the strength of the causal relationship binding together two concepts. As an example, with reference to Fig. 3, the causal relationship between concepts b and c is strong and positive, with b being the cause and c the effect. To make a more concrete example, suppose that b and c represent the already mentioned concepts, i.e. "the amount of users in a cell at a given time" and "the consumed energy consumed by the base station", respectively. According to the model, if the number of users increases, the consumed energy is expected to increase as well. In other words, the variation of the number of users in the cell is said to *cause* the amount of energy consumed to vary *a lot*.

In their simplest form, concepts take values in the binary discrete set $\{0, 1\}$. Zero conveys the idea that the concept in question is *inactive* (or *low* or *off*, depending on the concept nature). Conversely, a value of one means that the concept is *active*. Examples of concepts that fit this mapping include "the use of packet fragmentation", for, in a network it can be said that there is congestion (1) or not (0).

It should be noted, though, that other sets, such as $\{-1, 1\}$ and $\{-1, 0, 1\}$, are commonly chosen. Whereas a null value practically absorbs any causal implication between two concepts, the effect of a negative value is to invert causality – which is what actually happens in some situations. An example for the former category is the classic Transmission Control Protocol (TCP) congestion window mechanism, which traditionally either increases or decreases, depending on network conditions. Clearly, this behavior can be well mapped on the domain $\{-1, 1\}$. However, if we suppose that we create an algorithm for which the congestion window can also maintain its size, then a more apt domain would be $\{-1, 0, 1\}$.

Finally, continuous intervals are valid concept domains, as well. However, they are not a popular choice, for the inference process in such case can lead to chaotic behavior, thus not converging to any solution.

As for edge labels, typical values lie in the continuous real interval $[-1; +1]$. The closer to the boundaries the value is, the stronger the causal implication is: positive or negative, depending on the boundary it approaches, whether the right one or the left one.

The inference process is computationally lightweight and, most importantly, guaranteed to converge in a finite number of steps, provided that concepts are mapped on discrete sets. For instance, an FCM with n distinct, binary concepts converges to a solution in no more than 2^n steps [21].

The process involves repeated multiplications between the vector of all concepts and the adjacency matrix of the FCM studied. The vector used in the multiplication is referred to as the system state and represents the current state of the system. In a system having n distinct concepts, the state is a vector of dimension n. The system state determines which concepts are active/inactive—or increasing/constant/decreasing, depending on the choice made for the mapping domain. The multiplication operation uses the state vector as a stimulus, and makes it propagate throughout the network, flowing from causes to effects. Clearly, the vector-by-matrix multiplication might result in values lying outside the original domain. For this reason, a non-linear operation, such as a threshold, is usually needed to map the result onto the mapping domain chosen[19]. An example of the inference process is shown in Fig. 4 where the FCM is the same given in Fig. 3.

We propose to equip base stations with such cognitive capabilities, in order to save energy while adapting to the changing environment.

According to the procedure originally described by the authors in [22], we will define in the following sections a proper FCM for the problem of energy saving in cellular networks.

4.1. Identification of the concepts characterizing the problem

As outlined in [22], the first step towards the definition of a FCM involves the identification of the concepts that will compose the system state.

With reference to Fig. 5, we can define a set C of all the concepts that characterize the system under study. We can think of such concepts as belonging to different sets:

- Set A comprises all concepts on which the reasoning entity has direct control.
- Set Q collects all concepts that the reasoning entity cannot control directly but that are interesting because they give feedback on the achieved performance.
- Set E collects all concepts on which the reasoning entity has no direct control nor carry relevant information regarding the performance.

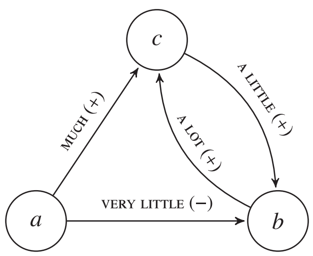

(a) Graphical representation

$$F = \begin{pmatrix} 0 & -0.1 & 0.6 \\ 0 & 0 & 0.9 \\ 0 & 0.3 & 0 \end{pmatrix}$$

(b) Mathematical representation (adjacency matrix)

Fig. 3. Fuzzy cognitive map example.

Multiplication Threshold operation

$$\begin{pmatrix}1 & 1 & 1\end{pmatrix} \cdot \begin{pmatrix}0 & -0.1 & 0.6\\ 0 & 0 & 0.9\\ 0 & 0.3 & \end{pmatrix} = \begin{pmatrix}0 & 0.2 & 1.5\end{pmatrix} \rightarrow \begin{pmatrix}0 & 1 & 1\end{pmatrix}$$

$$\begin{pmatrix}0 & 1 & 1\end{pmatrix} \cdot \begin{pmatrix}0 & -0.1 & 0.6\\ 0 & 0 & 0.9\\ 0 & 0.3 & \end{pmatrix} = \begin{pmatrix}0.0 & 0.3 & 0.9\end{pmatrix} \rightarrow \begin{pmatrix}0 & 1 & 1\end{pmatrix}$$

Fig. 4. FCM inference process example. As can be seen, the reasoning process results into the fixed point (0, 1, 1) in one iteration. Concepts mapped to the {0, 1} domain and threshold set to 0.

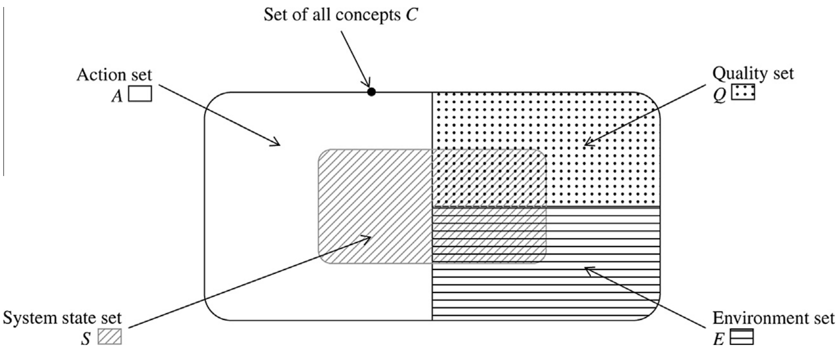

Fig. 5. Relations among the concept sets and the FCM.

As an example, let us consider a generic wireless network. Transmitting stations might enable packet fragmentation to hinder data corruption due to channel noise, which otherwise would reduce the throughput. The key concepts in such scenario can be "fragmentation", "noise", and "throughput". According to our framework, fragmentation can be directly controlled and belongs to set A. Throughput and noise cannot be controlled, but while the former is a relevant performance metric and would belong to set Q, the latter is not and would therefore be put in set E.

This toy example clearly shows that not necessarily all the variables have to be taken into account in the definition of the system state. Indeed, as can be seen in Fig. 5, $S \subseteq C$, meaning that some variables may not be considered, depending on the problem formulation. With respect to the toy example we devised, we do not consider the jitter experienced by stations ($\in Q$), as we are not interested in it.

Once concepts are found and classified, it is possible to create the system state vector, $\boldsymbol{s} = (\boldsymbol{a}, \boldsymbol{q}, \boldsymbol{e})$, where:

$$\begin{aligned}\boldsymbol{v} &= (v_i, \ldots, v_{n_V}) \quad v_i \in V \qquad \forall (\boldsymbol{v}, V)\\ &\in \{(\boldsymbol{a}, A), (\boldsymbol{q}, Q), (\boldsymbol{e}, E)\} \end{aligned} \tag{1}$$

The FCM needs to converge to a solution state $\boldsymbol{s}^* = (\boldsymbol{a}^*, \boldsymbol{q}, \boldsymbol{e})$ by finding a vector $\boldsymbol{a}^*$ such that the constraints expressed by $\boldsymbol{q}$ are satisfied before environmental conditions $\boldsymbol{e}$ change.

It should be noted that elements in set E are important although they cannot be directly controlled by the system nor do they provide any information directly related to the performance of the system. In fact, they might have relationships with concepts in the other sets, which must be taken into account to be, eventually, exploited.

For the problem considered in the paper of energy saving in cellular base stations, we selected the following concepts:

- Concepts in $S \cap A$: the use of higher frequencies (hi), the use of tri-sectorized operational mode (tri).
- Concepts in $S \cap Q$: the energy consumption (en), the blocking rate (br), the Signal to Interference-plus-Noise Ratio (snr).
- Concepts in $S \cap E$: the amount of voice users (v), the amount of users browsing the web (h), the amount of users that transfer data (f).

4.2. Definition of concept domains

The second step in defining the FCM involves the identification of concept domains and the pre-processing operations needed to perform the mapping operation.

The number of steps needed to make the reasoning process converge depends on the domains on which the concepts are mapped and the number of concepts themselves [21]. More precisely, the inference process reaches a solution within l^c steps, l being the number of levels of concept domains and c being the number of concepts. For this reason, we considered appropriate to adopt binary sets as the domains for all concepts.

In particular, we have opted for the $\{-1, 1\}$ domain as the domain for all concepts except the blocking rate, for which we used the binary set $\{0, 1\}$. This choice has been made to highlight the fact that a low blocking rate (0) should not be able to influence or, better, *cause* any other concept. Avoiding the use of zero in all other cases means that any variation can, at least in principle, entail some change in the other concepts.

It should be noted that only concepts in $S \cap A$ can be naturally mapped on such interval. Pre-processing operations are needed in order to map all other concepts. The identification of optimal values for the pre-processing stage depends on the problem under analysis and is a step that must be taken into account for the correct outcome of the operations. The following paragraphs are devoted to explain such operations in detail. When applicable, parameters representing variable names used in the simulating program are reported in parenthesis, using a `mono-spaced font`. Values used are reported in Table 1.

4.2.1. Raw data averages

A collection period (`DELTA_PERIOD`) is defined during which raw measurements for each of the variables in $S \cap (E \cup C)$ are acquired. After this period, if the number of samples collected is greater or equal than a target value, first and second order statistics are computed.

If the number of samples is lower, the sensor is polled at small regular intervals (`GRACE_PERIOD`), until the minimum number of samples is reached. This procedure ensures that a minimum number of collected samples, so as to reduce the chance of spurious variations.

4.2.2. Control chart

To distinguish significant changes from random variations, an approach based on exponential weighted moving average (EWMA) control charts is used [23].

Each sensor draws a control chart in which the Center Line is set to $\hat{\mu}$, i.e. an EWMA-based prediction (weighting coefficient β) of the actual mean value μ:

$$\hat{\mu} = \beta\mu + (1-\beta)\hat{\mu} \qquad \beta \in (0;1] \tag{2}$$

Upper and Lower Control Limits (UCLs and LCLs, respectively) are updated only when the controlled value ($\hat{\mu}$) crosses either boundary. They are bounded to the standard deviation (σ) by the following formulae:

$$\mathrm{UCL} = \hat{\mu} + C\sigma \cdot \sqrt{\frac{\beta}{2-\beta}} \quad \text{and} \quad \mathrm{LCL} = \hat{\mu} - C\sigma \cdot \sqrt{\frac{\beta}{2-\beta}} \tag{3}$$

where β is the same value used as a weighting coefficient in the EWMA-based prediction and C is a real parameter in $(0, +\infty]$.

Table 1
Parameters utilized for the validation. Unity of measurement appears next to the value, surrounded by square brackets.

Parameter name	Value
`DELTA_PERIOD`	180 [s]
`GRACE_PERIOD`	5 [s]
`HYSTERESIS`	threshold ±5%

4.2.3. Discretization

When a variation in a variable has been confirmed by the control chart, the predicted mean is mapped on the continuous set $[-1, 1]$ by means of a linear transformation. Given x_i the input value of a generic variable, its output x_o will be:

$$x_o = \frac{2 \cdot x_i - x_{min} - x_{max}}{(x_{min} - x_{max})} \tag{4}$$

In Eq. (4) x_{min} and x_{max} denote the minimum and maximum value that x_i can reach. As these values may not be known beforehand, if the input value is less than the current minimum or greater than the maximum, the output value will be set to the -1 or +1 and the relative limit updated.

4.2.4. Adaptive threshold and hysteresis

The resulting value is ultimately controlled against a threshold, in order to convert the continuous value to either element of the discrete set $\{-1, 1\}$. Such a threshold is set to the immediately older value recorded.

Choosing a static threshold is not an optimal solution: if the variable varies around a specific value which is not close to the chosen threshold, no variation will ever be registered and the system could possible settle to a sub-optimal point. The implementation of an adaptive threshold reduces the likelihood of such an event.

The drawback, however, is that even small variations around the threshold are perceived as radical changes. To counteract this side-effect we have introduced hysteresis (`HYSTERESIS`).

4.3. Definition and update of the fuzzy cognitive map

The third and last step is about embedding of any a priori knowledge of the problem to the FCM. Let us denote by f_{ij} the edge of the FCM that departs from i and arrives to j, i, and j being generic concepts in S.

1. We assume that concepts in the same set are *causally* independent from one another. Considering the set $S \cup Q$ as an example, this means that, for instance, the variation of the number of users that browse the web has no causal implication to (and from) the variation of the number of the users that place voice calls. This means that no edges arrive or depart from concepts that belong to the same class:
$$f_{ij} = f_{ji} = 0 \qquad \forall i,\ j \in \{S \cap V\}, \quad V \in \{A, Q, E\} \tag{5}$$
2. Concepts in the action set are not directly caused by any other concept. Instead they are triggered by the reasoning process. This translates into the fact that no edges point to any action concept, that is:
$$f_{ia} = 0 \qquad \forall a \in \{S \cap A\}, \quad \forall i \in \{S \cap V\}, \quad V \in \{Q, E\} \tag{6}$$
3. Similarly, we state that concepts related to quality metrics do not cause any variation in the concepts related to the environment. As an example, users will decide to call, browse the web and download files ignoring channel conditions (Signal to Interference and Noise Ratio (SINR), blocking rate and energy consumed by the base station). Mathematically:

$$f_{qe} = 0 \qquad \forall q \in \{S \cap Q\}, \quad \forall e \in \{S \cap E\} \tag{7}$$

4. We also know that both actions increase the number of frequency slots available, and, as a direct consequence, reduce the blocking rate. Therefore we may want to embed such information, by properly setting $f_{hi,br}$ and $f_{tri,br}$.

The resulting FCM is as follows:

$$F = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & f_{v,en} & f_{v,br} & f_{v,snr} \\ 0 & 0 & 0 & 0 & 0 & f_{h,en} & f_{h,br} & f_{h,snr} \\ 0 & 0 & 0 & 0 & 0 & f_{f,en} & f_{f,br} & f_{f,snr} \\ f_{hi,v} & f_{hi,h} & f_{hi,f} & 0 & 0 & f_{hi,en} & f_{hi,br} & f_{hi,snr} \\ f_{tri,v} & f_{tri,h} & f_{tri,f} & 0 & 0 & f_{tri,en} & f_{tri,br} & f_{tri,snr} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \tag{8}$$

In order to keep the FCM updated a learning algorithm has to be employed. Learning algorithms emulate human learning: just like human beings infer causality between two events when they perceive concomitant variations that respect a chronological order, learning rules for FCMs modify the labels of the edges that connect two concepts that experienced some change in a relatively short period of time.

If a concept experiences a positive variation and soon after another concept experiences a positive variation as well, it can be assumed that the two concepts are bounded by a positive causal relationship. The same holds if both variations experienced are negative. Conversely, if they undergo alternate variations (positive the first and negative the second, or viceversa), it can be inferred that there is a negative causal relationship that connects the two concepts.

A popular learning rule is known as Differential Hebbian Learning (DHL) which updates the edges in an FCM proportionally to the value of variations of the concepts [24]. Mathematically, if we denote by f_{ij}^t a generic edge at time t and by $\dot{C}_i$ the variation of concept i at time t, the DHL rule states that:

$$f_{ij}^t = f_{ij}^{t-1} + \eta\left(-f_{ij}^{t-1} + \dot{C}_i^t \dot{C}_j^t\right) \tag{9}$$

The parameter $\eta \in (0; 1]$ is known as "learning rate" [19] and its purpose is to lower the responsiveness of the algorithm, which otherwise could produce too abrupt updates.

A peculiar property of DHL is that it accounts less for the *causal history* of an edge. As can be seen in Fig. 6a, variations are assigned different levels of importance depending on how many variations of the same type happened immediately before. Even if many variations in a row take place, all sharing the same polarity, it takes only a few steps for the edge to change value (and sign).

Instead, in order to consider all variations of the same importance, we modified the DHL rule to devise a Linear Learning (LL) rule. According to such rule, edges are updated based only on the polarity of the variation, that is:

$$f_{ij}^t = f_{ij}^{t-1} + \eta \cdot \mathrm{sgn}\left(\dot{C}_i^t \dot{C}_j^t\right) \tag{10}$$

where sgn denotes the sign operator. Clipping is done to prevent edge values fall off of the $[-1; 1]$ interval, as also shown in Fig. 6b.

5. Simulation scenario

The simulation scenario is populated by seven base stations arranged according to the traditional honeycomb structure, as shown in Fig. 7. In the simulator, however, cells are represented by circular areas, which obviously overlap to some extent to one another. This means that on average the number of terminals per unit of space will be twice as much as in non-overlapping areas. Due to the symmetry of the structure, the central base station experiences a more uniform load, when compared to the outer base stations. In order to lessen the likelihood of biasing the results, we focus on the central base station.

The cell in the center is served by a base station that is equipped with cognitive capabilities. Such cognitive capabilities are based on the FCM designed in Section 4, and allow the base station to reason about the environment to reduce energy consumption while monitoring the blocking rate. All other base stations in the network do not employ any cognitive scheme and maintain all radio modules enabled at all times.

The simulating platform focuses on the periods of active communications between the base stations and the user terminals associated to it, so that it is possible to monitor energy consumption.

Terminals are static during their communications and are distributed over the coverage area of a base station following a uniform random distribution. Associations and deassociations to/from a base station follow a Poisson process, with parameters λ and μ, respectively. Both λ and μ depend on the type of communication occurring: voice call, web browsing, data transfer. Therefore, we define both parameters for each traffic category, resulting in six parameters: λ_v, μ_v, λ_h, μ_h, and λ_f, μ_f.

Each base station has a peak busy load of 50 users, weighted by the data shown in Fig. 8 in order to reflect real-world situations. Weights reflect the actual hourly load measured in a Vodafone 3G cell in London and were obtained via internal communication within the UK's Mobile VCE Core 5 Green Radio research program.

Voice traffic is modeled after the well-known Brady six-state model [25]. In our case, we restricted our attention to a subset of the six states of the original model, focusing only on the states for which there is a transmission from the base station to the terminal, i.e. when the other end of the communication is active. We assumed an average duration of calls of one minute, resulting in a μ_v of 1/60.

Web traffic has been modeled as a continuous repetition of two states: a downloading period to retrieve a page from the web, and a waiting period, to parse and read the page [26,13]. The download time depends on the size of the web page and eventual embedded objects. Object sizes follow a truncated lognormal distribution, while the number of the embedded objects in a web page follows a

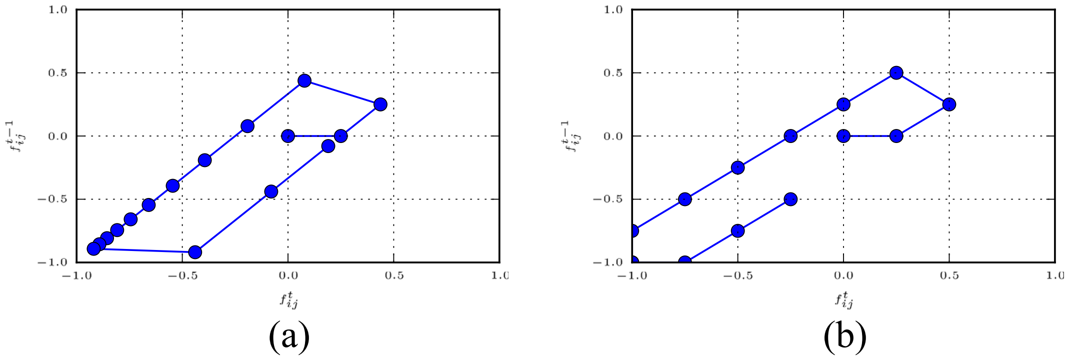

Fig. 6. Lagged-coordinated plot of the evolution of an FCM edge when updated by (a) Differential Hebbian Learning and (b) Linear Learning. Starting point is (0, 0). The sequence of variations is as follows: 2 positive, 10 negative, and 3 positive. Learning rate η set to 0.25.

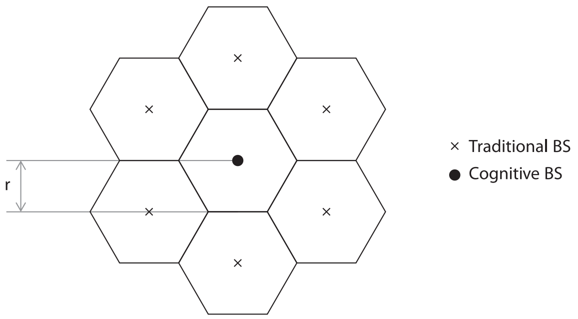

Fig. 7. Layout of the simulated scenario. The coverage radius r is set to 600 m in the simulations.

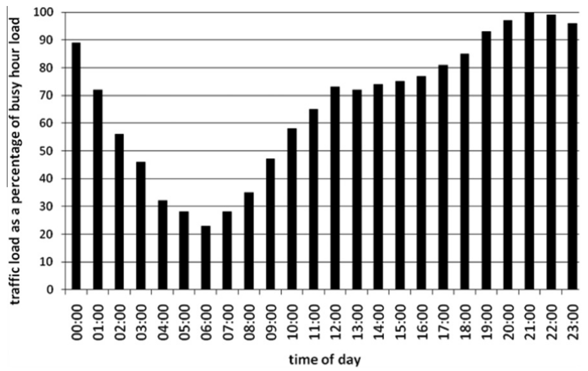

Fig. 8. Hourly variation of traffic load as a percentage of busy hour load over a typical day for a mobile network operator in London, UK.

truncated Pareto distribution. Reading and parsing times can be found by sampling an exponential distribution. Assuming that a web session for a mobile user lasts, on average, five minutes, we fixed μ_w to a value of 1/300.

The model for FTP sessions is similar to that for web sessions, except for the fact that there is no parsing time [27,13]. The download time exclusively depends on the size of the object to be transferred, which follows a Pareto distribution. Reading time is again modeled by an exponential random variable. Considering that a mobile user is not likely to make extensive use of FTP, we hypothesize an average session duration of two minutes, equivalent to a μ_f of 1/120.

In the simulations we fix the composition of traffic as follows: 50% voice traffic (α_v), 40% web browsing (α_h) and 10% FTP traffic (α_f). Traffic categories are independent from one another. By approximating each type of traffic as an $M/M/\infty$ it is possible to find the relative birth rates as:

$$\lambda = N_t \alpha_i \mu_i \qquad i \in v, h, f \tag{11}$$

where N_t is the number of users in a cell at a given hour t.

The environment broadly reflects that of an HSDPA network. The main parameters that characterize the system are reported in Table 2.

Base stations all operate on two bands, centered at 2 GHz and 5 GHz, and are characterized by a coverage radius of 600 m. The from-the-socket power P_M is computed according to a well known linear function of the transmission power P_{tx} [29], namely:

$$P_M = a \cdot P_{tx} + b \tag{12}$$

According to internal documentation within the Mobile VCE Green Radio research program, it has been shown that an HSDPA base stations consumes 857 W at 100% transmission power and 561 W at 20% transmission power. Constants a and b have been computed by regression and are equal to 9.25 from-the-socket Watts per transmission Watt and 487 Watts, respectively.

Regarding user capacity, a base station can accommodate at most 22 users per band when operating in omnidirectional mode and up to 15 users per band per sector when operating in tri-sectorized mode.

Table 2
Simulation configuration parameters. d is the distance in km.

Parameter	Value
System configuration	Broadly reflecting HSDPA Rel. 5
Spectral efficiency	0.8 b/s/Hz
Bandwidth per HSDPA band	5 MHz
Channel path loss models [28]	2 GHz: $128.1 + 37.6 \cdot \log(d)$ 5 GHz: $141.52 + 28 \cdot \log(d)$
HSDPA pilot power	20% of cell power budget

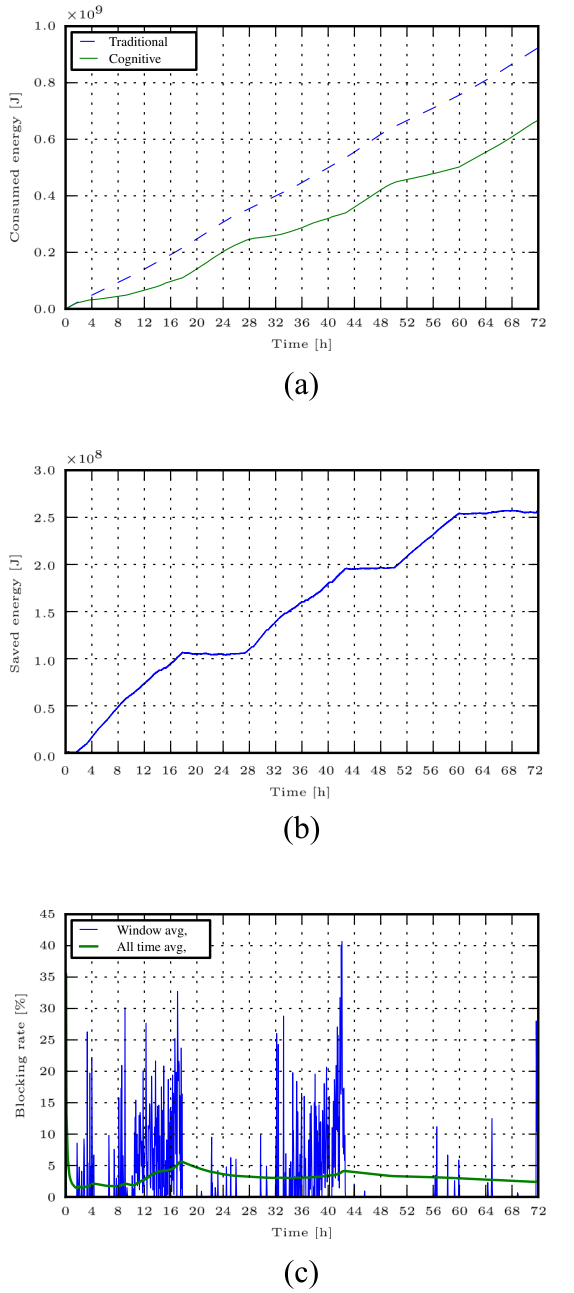

Fig. 9. (a) Consumed energy in traditional and cognitive base stations, and (b) energy saved by employing the cognitive scheme over a period of 72 h. (c) Example of the evolution of the blocking rate.

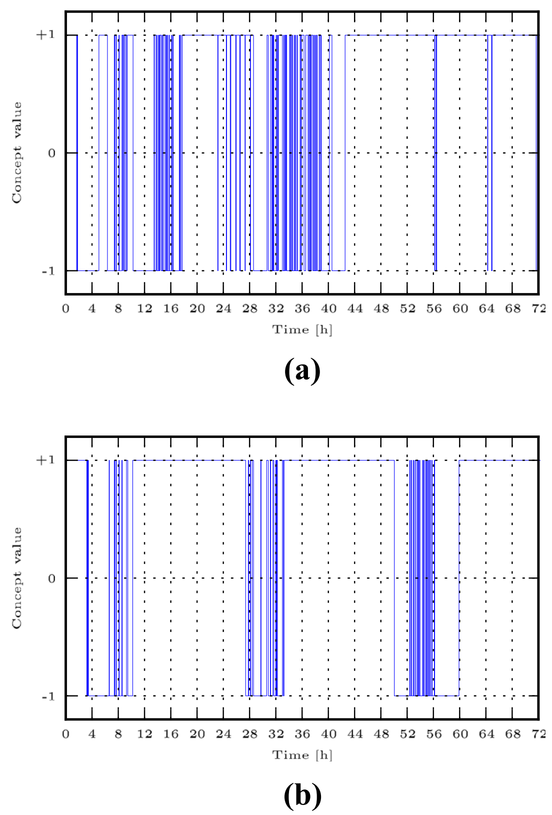

Fig. 10. Evolution of the action concepts: (a) use of higher frequencies (*hi*) and (b) use of tri-sectorized mode (*tri*).

6. Results

Simulation time covers three days, starting from midnight. As we intend to simulate three consecutive working days, the hourly variation reported in Fig. 8 does not change from day to day.

The dashed line in Fig. 9a represents the energy consumption by a traditional base station, i.e. when all six modules are always on. As can be expected, the curve is a linear function of the time, showing no change in the behavior of the base station. The behavior of the cognitive base station is represented by the solid line, which resembles a piecewise linear function. The curve reveals that at times when the user load is low, it is possible to save energy by switching off part of the radio modules. Conversely, when there is a high user load all radio modules must be kept active and no energy saving is possible. Such periods are the intervals in which the solid curve runs in parallel with the dashed curve, i.e. from hour 18 to 28, from hour 42 to 50, and from hour 60 to 72.

Fig. 9b shows the difference between the energy consumed by a traditional base station and the cognitive base station. The periods in which energy saving is not possible can be recognized more clearly. By analyzing the slopes, it can be observed that our scheme saves slightly more than 250 MJ over the investigated 72-h duration (3.5 MJ/h or almost 1 kW), equivalent to approximately 27% saving over the traditional case. In the busy periods in which all radio modules must remain active, energy saving is negligible. However, in quiet periods in which it is possible to employ power saving modes, this saving can be as much as 50% of the total energy consumption of the traditional system (6.25 MJ/h or 1.7 kW).

Inevitably, there is a trade-off between energy saving and blocking rate. As the thin line in Fig. 9c indicates, the blocking rate tends to increase when a subset of the transceivers is turned off. High peaks of the blocking rate can be registered by averaging the value using a window equal to

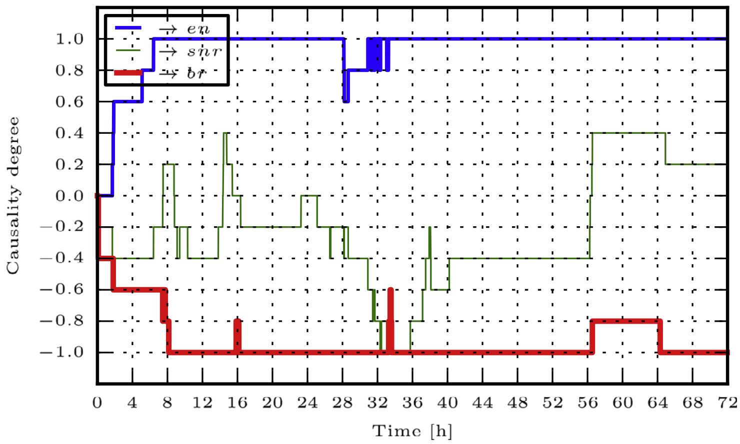

(a)

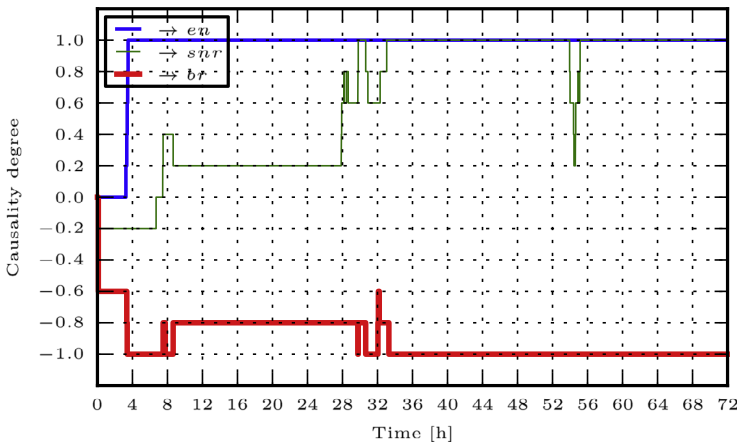

(b)

Fig. 11. Example of evolution of the causal relationships: (a) between *hi* (use of higher frequencies) and the quality-related concepts and (b) between *tri* (use of tri-sectorized mode) and the quality-related concepts.

the collection period (180 s). However, it should be noted that it may be that only a few users try and join the network during such intervals, thereby the difference between local maxima and following minima is often pronounced. The thick curve represents the all-time average, in which the blocking rate remains below 5% throughout the simulation.

Fig. 10 reproduces the behavior of the radio modules in the cognitive base station. In such graphs, the step curves are either high or low, symbolizing the use or not, in order, of the extra capacity. Fig. 10a shows the use of the higher band, whereas Fig. 10b shows the use of the tri-sectorized mode. The oscillations during the low-activity periods hint at the fact that an even greater saving could be, in principle, possible.

For the sake of clarity, we also report the causal relationships between the action concepts and the quality-related concepts (Fig. 11). It can be noticed that the cause-effect relation between any action concept and the blocking rate is negative and approaches the lower bound (-1). This means that using the higher band in conjunction with the lower band and using the tri-sectorized mode causes the blocking rate to decrease. A similar effect happens with energy: turning on the radio modules causes the consumption of energy to rise (positive cause–effect relation). The relation with the SINR is less clear. Whereas the employment of tri-sectorized mode seems to increase the SINR, when turning on the higher band it an increase in the SINR may not follow. This could be explained by the fact that, although it is true that switching on more radio modules increases the available capacity and should lower the interference, it is likewise true that interference depends also on the activity in the neighboring cells, which is neither modeled in our system, nor it can be controlled.

7. Discussion

For a real-world implementation of the architecture proposed we should focus on two main aspects, namely

(i) the cost of collecting all the cognitive information and (ii) the extension of the algorithm to include multiple cells.

To better analyze the first point, let us investigate the concept categories as we have defined in Section 4.1. It should be noted that all action-related concepts are private to a cognitive base station and, therefore, they are inherently known by it. Clearly, there is no cost associated with their retrieval. The same applies to the concepts linked to the quality of service. At any time, the cognitive base station has access to all measurements related to energy consumption, blocking rate, and SINR. In particular, the latter concept refers to the ratio measured at the base station level. Finally, it is sensible to assume that the cognitive base stations know (or can compute) also the concepts pertaining to the environmental class, i.e., the number of users connected and the type of traffic they generate. For this category, there is a cost to be paid for the information retrieval (statistics computation, traffic analysis, etc.). However, it is minimal and its impact on the overall cost is negligible.

On the other hand, it is possible (and worth investigating) to include in the reasoning process external measurements, such as the quality experienced by the terminals (the so-called Quality of Experience), or in a more futuristic scenario, actions that take place at the terminal. In this case, the presence of a signaling protocol must be envisioned, so that the information (knowledge and commands) can flow between remote entities. Such a protocol clearly increases the complexity, hence the cost, of the overall architecture but can potentially lead to improved performance levels.

As for the second point, the algorithm, as presented in this work, has been simulated on a single cell system. However, in order for it to be implemented in the real world, it is mandatory to consider systems composed of multiple cells.

Interestingly, the devised framework is particularly scalable. In fact, FCMs can easily be merged with one another[24] and lend themselves nicely to create a distributed reasoning system.

However, attention must be paid at the potential pitfalls when developing such a combined scheme. Indeed, it is not obvious that merging the knowledge accumulated by different reasoning systems will lead to a greater performance level.

For instance, let us imagine a scenario in which there are two neighboring cognitive base stations and any terminal can connect to either one, indifferently. In this scenario, it is possible that one base station learns that the blocking probability tends to remain low, though most of its radio modules are turned off. The first issue here is that, apparently the first station that acquires such a knowledge is more likely do develop a selfish behavior (at the expenses of the other). The second issue lies in the knowledge exchange. After merging knowledge, it may happen that the other base station learns that most radio modules can be safely turned off, which is obviously wrong.

This plain scenario tells us that several are the factors that should be thought of before considering a network-wide implementation of the scheme offered. Solutions may range from the specification of overarching common goals, to the application of game-theoretic models, from reputation-based systems, to the definition of ad hoc rules. No matter the choice, it is worth bearing in mind that every solution might hinder scalability in the long run.

8. Conclusions

As the ICT sector is partly responsible with the increase in the global carbon footprint, methods are being developed to reduce the amount of energy consumed in the area of telecommunications.

We have described how this problem can be tackled by exploiting the cognitive networking vision, according to which networks should be able to self-configure in view of a specific goal, limiting the need of human intervention. Specifically, we have proposed a novel dynamic scheme to perform energy saving in HSDPA base station by driving the selection of the active transceivers.

Results demonstrate that the usage of fuzzy cognitive maps for reasoning and decision-making enables to save a relevant amount of energy in operation under realistic traffic patterns. As a consequence, the proposed architecture seems suitable to support dynamic configuration of cellular base stations. Clearly, the achieved savings correspond to an increase in the blocking probability. The identification of a refined methodology to enable a more flexible trade-off between energy saving and blocking rate will be subject of further investigation on the topic, as well as a generalization of the proposed architecture to enable collaboration among several neighboring base stations.

Acknowledgments

This work was partially supported by the Research Project GREENET (PITN-GA2010264759), the Green Radio Core Research Program of the Virtual Centre of Excellence in Mobile & Personal Communications, Mobile VCE, www.mobilevce.com, the ICT-ACROPOLIS Network of Excellence, FP7 project number 257626, www.ict-acropolis.eu, and COST Actions IC0902 and IC0905 "TERRA".

References

[1] A. Bianzino, C. Chaudet, D. Rossi, J. Rougier, A survey of green networking research, IEEE Communications Surveys & Tutorials, 14(1) (2012).
[2] The Climate Group London, The Global e-Sustainability Initiative (GeSI), Smart 2020: Enabling the Low Carbon Economy in the Information Age, 2008.
[3] L. Saker, S.E. Elayoubi, Sleep mode implementation issues in green base stations, in: IEEE 21st International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC), 2010, pp. 1683–1688.
[4] D. Tipper, A. Rezgui, P. Krishnamurthy, P. Pacharintanakul, Dimming cellular networks, in: IEEE Global Telecommunications Conference (GLOBECOM), 2010, pp. 1–6.
[5] E. Oh, B. Krishnamachari, Energy savings through dynamic base station switching in cellular wireless access networks, in: IEEE Global Telecommunications Conference (GLOBECOM), December 2010, pp. 1–5.
[6] R.W. Thomas, L.A. DaSilva, A.B. MacKenzie, Cognitive networks, in: 1st IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2005, pp. 352–360.
[7] M.A. Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, Optimal energy savings in cellular access networks, in: IEEE International Conference on Communications Workshops (ICC), June 2009, pp. 1–5.

[8] S. Thajchayapong, J.M. Peha, Mobility patterns in microcellular wireless networks, IEEE Transactions on Mobile Computing 5 (2006) 52–63.
[9] L. Chiaraviglio, D. Ciullo, M. Meo, M.A. Marsan, Energy-efficient management of UMTS access networks, in: 21st International Teletraffic Congress (ITC), September 2009, IEEE, 2009, pp. 1–8.
[10] A. Kaloxylos, T. Rosowski, K. Tsagkaris, J. Gebert, E. Bogenfeld, P. Magdalinos, A. Galani, K. Nolte, The e3 architecture for future cognitive mobile networks, in: IEEE 20th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), September 2009, pp. 1601–1605.
[11] M.-G. Di Benedetto, L.D. Nardis, Cognitive routing models in UWB networks, in: 3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CrownCom), 2008, pp. 1–6.
[12] D.D. Clark, C. Partridge, J.C. Ramming, J.T. Wroclawski, A knowledge plane for the internet, in: Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, ACM, Karlsruhe, Germany, 2003, pp. 3–10.
[13] O. Holland, O. Cabral, F. Velez, A. Aijaz, P. Pangalos, H.A. Aghvami, Opportunistic load and spectrum management for mobile communications energy efficiency, in: 2011 IEEE 22nd International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC), IEEE, 2011, pp. 666–670.
[14] C. Facchini, F. Granelli, N.L. da Fonseca, Cognitive service-oriented infrastructures, Journal of Internet Engineering (2010). Special issue on Service-Oriented Architectures.
[15] D. Fletcher, D. Nguyen, K. Cios, Autonomous synthesis of fuzzy cognitive maps from observational data: preliminaries, in: IEEE Aerospace Conference, 2005, pp. 1–9.
[16] C. Fortuna, M. Mohori, Trends in the development of communication networks: cognitive networks, Computer Networks 53 (9) (2009) 1354–1376.
[17] O. Holland, C. Facchini, H. Aghvami, O. Cabral, F. Velez, Opportunistic spectrum and load management for green radio, in: E. Hossain, V.K. Bhargava, G. Fettweis (Eds.), Green Radio Communication Networks, Cambridge University Press, 2012, pp. 56–99.
[18] B. Kosko, Fuzzy cognitive maps, International Journal of Man–Machine Studies 24 (1) (1986) 65–75.
[19] J.A. Dickerson, B. Kosko, Virtual worlds as fuzzy cognitive maps, in: IEEE Virtual Reality Annual International Symposium, Seattle, WA, USA, 1993, pp. 471–477.
[20] C.D. Stylios, V.C. Georgopoulos, G.A. Malandraki, S. Chouliara, Fuzzy cognitive map architectures for medical decision support systems, Applied Soft Computing 8 (3) (2009) 1243–1251.
[21] A.K. Tsadiras, Comparing the inference capabilities of binary, trivalent and sigmoid fuzzy cognitive maps, Information Sciences 178 (20) (2008) 3880–3894. Special Issue on Industrial Applications of Neural Networks, 10th Engineering Applications of Neural Networks.
[22] C. Facchini, F. Granelli, Towards a model for quantitative reasoning in cognitive nodes, in: 3rd IEEE Workshop on Enabling the Future Service-Oriented Internet–Towards Socially-Aware Networks, Honolulu, Hawaii, USA, December 2009, pp. 1–6.
[23] M.G. Kallitsis, G. Michailidis, M. Devetsikiotis, Measurement-based optimal resource allocation for network services with pricing differentiation, Performance Evaluation 66 (2009) 505–523.
[24] B. Kosko, Neural Networks and Fuzzy Systems: A Dynamical Systems Approach to Machine Intelligence, Prentice-Hall, 1992.
[25] P. Brady, A model for generating on–off speech patterns in two-way conversation, The Bell System Technical Journal 48 (1969) 2445–2472.
[26] O. Holland, A. Attar, O. Cabral, F. Velez, A. Aghvami, Intra-operator spectrum sharing concepts for energy efficiency and throughput enhancement, in: 2010 3rd International Symposium on Applied Sciences in Biomedical and Communication Technologies (ISABEL), IEEE, 2010, pp. 1–6.
[27] O. Holland, V. Friderikos, A. Aghvami, Green spectrum management for mobile operators, in: 2010 IEEE GLOBECOM Workshops, IEEE, 2010, pp. 1458–1463.
[28] O. Cabral, F. Meucci, A. Mihovska, F.J. Velez, N.R. Prasad, R. Prasad, Integrated common radio resource management with spectrum aggregation over non-contiguous frequency bands, Wireless Personal Communications 59 (2011) 499–523.
[29] F. Richter, A.J. Fehske, G.P. Fettweis, Energy efficiency aspects of base station deployment strategies for cellular networks, in: 70th Vehicular Technology Conference Fall, IEEE, 2009, pp. 1–5.
[30] ITU-R, Background on IMT-Advanced, 2008. Available at http://www.itu.int/md/R07-IMT.ADV-C-0001/en.

Christian Facchini received the B.Sc. and M.Sc. degrees in Telecommunications Engineering from the University of Trento, Italy, in 2006, and 2008, respectively. He has been awarded the doctoral degree at the Dept. of Information Engineering and Computer Science at the same University. In 2009 he was visiting student at the Institute of Computing at the State University of Campinas (Brazil) and in 2011 he spent a period as a visiting research associate at the Centre for Telecommunications Research at King's College London (UK), where he worked on green mobile communications. His research interests include design, modeling, and performance evaluation of telecommunication networks, with special focus on intelligent and adaptive networks.

Dr. **Oliver Holland** obtained his B.Sc. degree (with First Class Honours) from Cardiff University, and his Ph.D. from King's College London. He is working within the ICT-ACROPOLIS Network of Excellence on Dynamic Spectrum Access (DSA) solutions and is part of the management team of that project, and is also contributing to the Virtual Centre of Excellence in Mobile and Personal Communications (Mobile VCE) on DSA techniques for Green Radio. He is Treasurer of the IEEE DySPAN Standards Committee (DySPAN-SC); he is also a member of the IEEE P1900.6 and IEEE P1900.1 working groups, was a Technical Editor of the IEEE 1900.4 standard, and is Technical Editor for the developing IEEE 1900.1a standard. He is a member of the Management Committee representing the UK within two prestigious collaborative "COST" actions on the topic of cognitive radio, and holds various leadership positions within these COST actions. He has served on the Technical Programme Committees of all major conferences in the area of mobile and wireless communications, has served as Session Chair and Panellist at a number of conferences covering Green Radio and Cognitive Radio, among other topics, and frequently serves as a reviewer for various prestigious international conferences and journals. He was guest editor of the special issue "Achievements and the Road Ahead: The First Decade of Cognitive Radio," which appeared in IEEE Transactions on Vehicular Technology, was co-chair of the "Cognitive Radio and Cooperative Communications" track of IEEE VTC 2010-Fall, and he is an Associate Editor of IEEE Transactions on Vehicular Technology. He is an Officer of the IEEE Technical Committee on Cognitive Networks (TCCN), serving as liaison between TCCN and COST IC0905 "TERRA", and between TCCN and DySPANSC, and he is currently the Chair of the United Kingdom and Republic of Ireland Chapter of the IEEE Vehicular Technology Society. He has jointly authored over 90 publications on a variety of topics, including several book chapters and one patent. The 68 of his publications that Google Scholar references have been cited more than 290 times.

Fabrizio Granelli received the «Laurea» (M.Sc.) degree in Electronic Engineering from the University of Genoa, Italy, in 1997, with a thesis on video coding, awarded with the TELECOM Italy prize, and the Ph.D. in Telecommunications from the same university, in 2001. Since 2000 he is carrying on his activity as Associate Professor in Telecommunications at the Dept. of Information Engineering and Computer Science – University of Trento (Italy) and coordinator of the Networking Laboratory. In August 2004, he was visiting professor at the State University of Campinas (Brazil). He is author or co-author of more than 100 papers published in international journals, books and conferences, and he is member of the Technical Committee of the International Conference on Communications (from 2003 to 2007) and Global Telecommunications Conference (GLOBECOM2003 and GLOBECOM2004). He is guest-editor of ACM Journal on Mobile Networks and Applications, special issues on "WLAN Optimization at the MAC and

Network Levels'' and ''Ultra-Wide Band for Sensor Networks'', and Co-Chair of 10th IEEE Workshop on Computer-Aided Modeling, Analysis, and Design of Communication Links and Networks (CAMAD'04). He is General Vice-Chair of the First International Conference on Wireless Internet (WICON'05) and General Chair of the 11th IEEE Workshop on Computer-Aided Modeling, Analysis, and Design of Communication Links and Networks (CAMAD'06). He is Co-Chair of GLOBECOM 2007–2009 Symposia on ''Communications QoS, Reliability and Performance Modeling''. His main research activities are in the field of networking and signal processing, with particular reference to network performance modeling, cross-layering, medium access control, wireless networks, cognitive radio systems, and video transmission over packet networks. He is Senior Member of IEEE, Elected Chair of IEEE ComSoc Technical Committee on Communication Systems Integration and Modeling (CSIM) and Associate Editor of IEEE Communications Letters, IEEE Communications Surveys and Tutorials and International Journal on Communication Systems.

Nelson L.S. da Fonseca received his Electrical Engineer (1984) and M.Sc. in Computer Science (1987) degrees from The Pontificial Catholic University at Rio de Janeiro, Brazil, and the M.Sc. (1993) and Ph.D. (1994) degrees in Computer Engineering from The University of Southern California, USA. Since 1995, he has been affiliated with the Institute of Computing of the State University of Campinas, Campinas, Brazil where is currently a Full Professor. He has published over 250 refereed papers and supervised over 50 graduate students. He is the Editor-in-Chief of the IEEE Communications Surveys and Tutorials. He served as Editor-in-Chief for the IEEE Communications Society Electronic Newsletter and Editor of the Global Communications Newsletter. He is a member of the editorial board of: Computer Networks, IEEE Communications Magazine, Peer-to-Peer Networking and Applications, International Journal of Communication Systems and Journal of Internet Service and Applications. He served on the editorial board of the IEEE Transactions on Multimedia, Brazilian Journal of Computer Science, and on the board of the Brazilian Journal on Telecommunications. He is the recipient of Elsevier Computer Networks Editor of the Year 2001. He served as ComSoc Director for Latin America. He served as ComSoc Director of On-line Services and served as technical chair for over 15 ComSoc symposia and workshops.

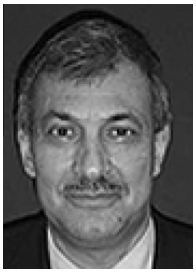

Hamid Aghvami joined King's College London as a member of the academic staff in 1984. In 1989 he was promoted to Reader, and in 1993 was promoted Professor in Telecommunications Engineering. He is presently the Director of the Centre for Telecommunications Research at King's. He carries out consulting work on Digital Radio Communications Systems for British and International companies; he has published over 500 technical papers and given invited talks and courses the world over on various aspects of Personal and Mobile Radio Communications. He was Visiting Professor at NTT Radio Communication Systems Laboratories in 1990, Senior Research Fellow at BT Laboratories in 1998–1999, and was an Executive Advisor to Wireless Facilities Inc., USA, in 1996–2002. He is Managing Director of Wireless Multimedia Communications Ltd., his own consultancy company. He leads an active research team working on numerous mobile and personal communications projects for third and fourth generation systems; these projects are supported both by government and industry. He was a member of the Board of Governors of the IEEE Communications Society in 2001–2003, was a Distinguished Lecturer of the IEEE Communications Society in 2004–2007, and has been member, Chairman, and Vice-Chairman of the technical programme and organizing committees of a large number of international conferences. He is also founder of the International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC), a major yearly conference attracting some 1,000 attendees. He was awarded the IEEE Technical Committee on Personal Communications (TCPC) Recognition Award in 2005 for his outstanding technical contributions to the communications field, and for his service to the scientific and engineering communities. He is a Fellow of the Royal Academy of Engineering, Fellow of the IET, Fellow of the IEEE, and in 2009 was awarded a Fellowship of the Wireless World Research Forum in recognition of his personal contributions to the wireless world, and for his research achievements as Director at the Centre for Telecommunications Research at King's.