

Energetic Performance of Service-oriented Multi-radio Networks: Issues and Perspectives

Mauro Caporuscio¹, Damien Charlet¹, Valerie Issarny¹, and Alfredo Navarra²

¹ INRIA-Rocquencourt

Domaine de Voluceau - 78153 Le Chesnay, France

{mauro.caporuscio, damian.charlet, valerie.issarny}@inria.fr

² LaBRI

Cours de la Libération - 33405 Talence, France

navarra@labri.fr

Abstract. Wireless devices now hold multiple radio interfaces, allowing to switch from one network to another according to required connectivity and related quality. Still, the selection of the best radio interface for a specific connection is under the responsibility of the end-user in most cases. Integrated multi-radio network management so as to improve the overall performance of the network(s) has led to a number of research efforts over the last few years. However, several challenges remain due to the inherent complexity of the problem. This paper specifically concentrates on the comprehensive analysis of energy-efficient multi-radio networking for pervasive computing. Building upon the service oriented architectural style, we consider pervasive networks of services, which are deployed on the various networked nodes. The issue is then to optimize the energetic performance of the pervasive network through careful selection of the radio link over which service access should be realized for each such access. This leads us to examine first the energetic performance of service access for most common wireless interfaces in use today (Bluetooth, WiFi and GPRS) and then introduce a formal model of service-oriented multi-radio networks. The proposed model enables characterizing the optimal network configuration in terms of energetic performance, which is shown to be a NP-hard problem and thus requires adequate approximation.

1 Introduction

As in particular targeted by Beyond 3G (B3G) networks, the recent evolution of mobile networks introduces the convergence of wireless technologies, where several radio interfaces may be used concurrently. B3G-capable devices then hold several radio interfaces, and allow switching from one radio interface to another (e.g., upon disconnection, to save money, to save energy, ...). Still, taking benefit of such a rich networking environment is raising tremendous challenges. Indeed, while operating systems start embedding support for integrated management of multi-radio networks (e.g., Windows Mobile 5 deals with transparent switching from GPRS to WiFi during Internet access upon WiFi connectivity), the end-user is in charge of explicitly choosing and possibly switching networks in most cases. A key challenge for integrated management of multi-radio networks is then to effectively contribute to improving the performance of the networking environment. Multiple performance criteria must further be taken into account. These in particular include network throughput and energy consumption, which are antagonistic performance factors. Additionally, multi-radio network management shall be made as far as possible transparent to the end-users and possibly to application developers, simply requiring them to abstractly specify base connection profile/constraints.

The above concern has led researchers to investigate multi-radio network management below the application layer. The connectivity management middleware presented in [21] provides a connection manager from which the application may get accurate knowledge about connectivity through the various links and related performance, and then decide upon the specific network connection(s) to establish. Network connections may also be adapted by the middleware according to application-defined policy and changes in the networking environment. Adaptive network connection may alternatively be realized in the lower network layer [23], leading to fully transparent solutions towards optimizing multi-radio networking on the end-user

devices. One approach is then to favor the use of the low power radio as long as related bandwidth meets system requirements (e.g., to exchange control messages and even data messages for which the bandwidth requirement is low [5]). Coordinated usage of the various radio networks may also be customized for specific networking functions, like resource discovery [17, 7].

Independently of the system layer in which multi-radio network management is implemented, most solutions we are aware of concentrate on optimizing network performance locally, i.e., the decision of which network channel to use for a specific connection is based on local information about network quality. As such, it cannot be guaranteed that this will lead to a globally optimal network usage in terms of performance. However, it seems to be the only tractable way to multi-radio network management [1]. Still, network connection on the terminal should not only be established according to local network quality and bandwidth requirement for the specific connection. Performance of network usage in terms of energy consumption depends on the set of local connections already open due to the non-negligible base energy cost associated with network interfaces [7]. Obviously, the energetic performance of the multi-radio network evolves as network connections are established and closed on the terminal, possibly requiring adaptation over time. Also this does not solely concern the energetic performance on the given terminal but also the energetic performance of the wireless nodes in communication range. Furthermore, networking capabilities vary among nodes, as it cannot be assumed that all networked nodes embed the very same set of network interfaces and even if they do, they must agree on the networking mode and possibly network channel [1]. As a result, energy-efficient multi-radio network management is still in its infancy, with several issues to be solved before it can be effectively deployed.

This paper specifically concentrates on the comprehensive analysis of the energetic performance of multi-radio networking in the context of pervasive computing. Indeed, multi-radio networking is a key enabler of the pervasive computing vision, as it promotes anytime, anywhere access to the digital world, whether proximity-based or not. Improving energetic performance of the multi-radio network will allow enhancing network connectivity, while optimizing autonomy of the wireless nodes.

The approach to pervasive computing that we promote lies in service orientation with autonomous pervasive services being deployed on the networked nodes, whether mobile or stationary, wired or wireless, resource-constrained or resource-rich [10]. Pervasive services abstract the various digital resources, from base sensor/actuator (e.g., display) to advanced applications (e.g., nomadic collaborative gaming). Then, services network together according to required and provided functionalities, composing their functions to realize rich distributed services [13, 14]. The actual networking of services depends on the technologies embedded on the hosts of the services' clients and providers since access to a service requires the service's client and provider to use a common network interface to communicate, a common service discovery protocol to find each other, and a common interaction protocol to understand each other (Section 2). Further, energy-efficient networking of services shall account for the respective energetic performance of the underlying radio interfaces (Section 3). This leads us to introduce a formal model of service-oriented multi-radio networks, from which to derive networking configurations that optimize the networks' energetic performance (Section 4). As shown, the problem is NP-hard and optimal energetic performance can only be approximated. As part of the PLASTIC project [10], we are currently developing a middleware for pervasive service-oriented computing, embedding support for energy-efficient multi-radio networking (Section 5).

2 Service-oriented Multi-radio Systems

The Service-Oriented Architectural (SOA) style is structured around 3 key architectural components: (*i*) service provider, (*ii*) service client, and (*iii*) service discovery. Base architectural connectors then relate to protocols for: (1) service discovery, and (2) service interaction.

The SOA paradigm is the right paradigm to engineer pervasive applications. Functionalities provided by networked resources may conveniently be abstracted as services. Specifically, a pervasive service represents an autonomous networked entity that provides a set of functionalities to its environment, which is continuously changing in particular due to changes in network connectivity. Thanks to the service discovery protocol, networked pervasive services may be dynamically discovered for consumption by service clients, as

hosts of clients and services move and enter in communication range of each other (either via the network infrastructure or directly in an ad hoc way). However, both discovery and consumption of services require service clients and providers to agree on the semantics of service functionalities at the application layer, and to use matching connectors, from the middleware to the link layers. These are basic requirements for services clients and providers to actually meet and interact in a way that guarantees dependable service provisioning. One way to satisfy such requirements is for clients and providers to use both common service descriptions, and networking protocols at all layers. This assumption is made by most software platforms for pervasive computing (e.g., Gaia [22], Aura [8], WSAMI [9]). Those platforms introduce advanced middleware to ease the development of pervasive applications composed out of networked resources. However, they are too restrictive regarding the networked resources that may be integrated since resources have to host the specific middleware to be known by pervasive applications. Furthermore, proposed middleware do not account for the multi-radio networking capabilities of hosts that is becoming quite common feature for today's advanced handhelds/Smartphone.

As part of the research of the INRIA ARLES³ project-team, we are currently investigating solutions to actually enable pervasive services in B3G networks, i.e., allowing services to be deployed on the various B3G nodes and to compose, independently of their underlying technologies. Specifically, we concentrate on the development of pervasive connectors so that services available in the B3G network may be discovered and accessed although:

- Services clients and providers may be developed independently without a priori knowledge of each other and using different software technologies, thus preventing interaction based on syntactic matching of provided and required service functionalities;
- Services clients and providers may be developed using heterogeneous middleware technologies and thus built upon distinct service discovery and access protocols, and
- Service clients and providers may communicate via heterogeneous network technologies, and have access to multiple radio technologies at the same time.

Our approach to the first issue lies in semantic-rich services, exploiting the semantic Web for the semantic modeling of service capabilities. This allows powerful and unambiguous reasoning about the behavior of pervasive services. Related service discovery and composition for the pervasive computing environment are presented in [13,14]. In a complementary way, actual interactions between services clients and providers rely on online event-based middleware protocol translation [6]. Actual communication in the context of heterogeneous networking then subdivides into two issues depending on whether the client and service embed compatible network technologies or not. In the latter case, we consider usage of network bridges (e.g., [20]). In this paper, we concentrate more specifically on the former issue, further assuming availability of multi-radio networks on hosts, but not necessarily.

Considering pervasive services clients and providers with access to multiple radio networks, the key issue is the selection of the optimal radio link for each network connection so as to optimize the overall performance of the pervasive network. However, the pervasive network is actually an aggregation of independent, ad hoc and infrastructure-based, networks and the aggregation is specific to each node according to the node's connectivity and services that are (to be) accessed. The pervasive network shall thus be defined at the service level so that for any node v in the pervasive network, there is at least another node w in the network with v being a client of a service provided by w ⁴.

Our goal is then to optimize the performance of service-oriented pervasive networks that dynamically compose, by realizing the selection of the most appropriate radio network for each service session, at the middleware layer. Various performance factors are eligible, with selection criteria being antagonistic for most of them, e.g.:

- Optimizing response time by selecting the radio network with the highest bandwidth possibly at the expense of power consumption.

³ <http://www-rocq.inria.fr/arles/>

⁴ Notice that by considering service discovery as a service, related networking is also assumed by this definition.

- Optimizing energy consumption by selecting the radio network with the lowest power consumption.
- Optimizing dependability by selecting the radio network that has the largest coverage and guaranteed quality of service, at the expense of financial cost.
- Optimizing financial cost by selecting free of charge radio networks, often at the expense of quality.

Since the primary requirement of pervasive computing is anytime, anywhere access, we favor the energy criterion to maximize autonomy of nodes, while meeting bandwidth requirements for the connection but not enhancing response time. Dependability and financial costs are further taken as additional discriminatory criteria, although we no longer consider them in the remainder of this paper. The next section discusses issues raised by energy-efficient networking, in light of the energy consumption of the wireless networks in use today.

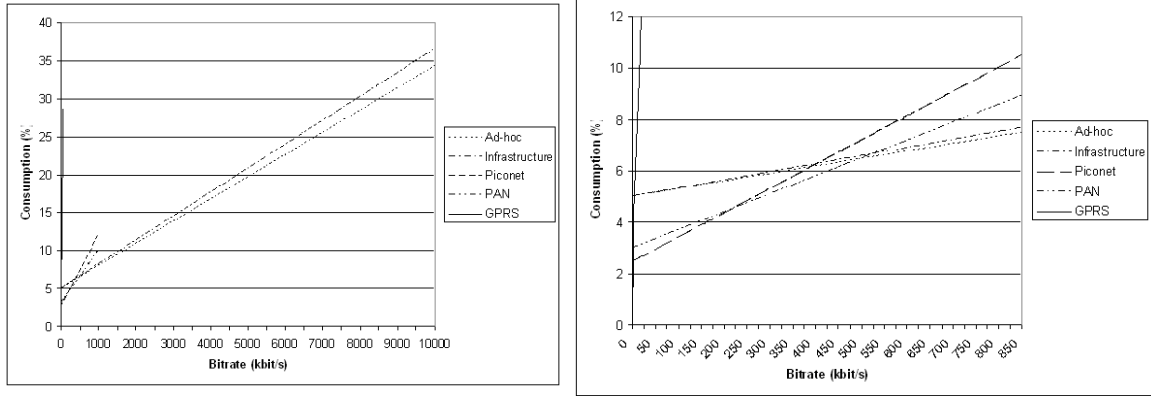
3 Energetic Performance of Multi-Radio Networks

Several studies of the power consumption of radio interfaces may be found in the literature (e.g., [12, 18, 7]). These studies show that energy consumption differs from one radio interface to another. In this section, we report on the energetic performance of radio interfaces available on today's handhelds and smartphones, according to the amount of data transferred, hence simulating access to various types of services (from discrete to continuous). Those measures build upon our previous work [7], which we enrich here with the analysis of GPRS.

3.1 Energetic Performance of Radio Networks

This study reports on the measures of the power consumed on handheld devices (namely HP Ipaq HX4700 and HP Ipaq H6340) during data transfer through each of the device's radio interfaces, i.e., Bluetooth, WiFi and GPRS. The WiFi and Bluetooth interfaces are further studied according to their various operating modes. Specifically, the WiFi technology has been designed to operate in two different modes: ad hoc and infrastructure. Bluetooth may also be configured into two different operating modes: the "master-slave" piconet mode (the default Bluetooth configuration belongs to this mode), and the IP-based PAN (personal area network) mode. Our performance measures thus treat 5 cases of radio interfaces: WiFi infrastructure, WiFi ad hoc, Bluetooth piconet, Bluetooth PAN, and GPRS. In all the tests, we measured the power consumption after half an hour of operation. The curves of discharge of the devices not being linear, all measurements are started when the battery is fully loaded. We discriminate the consumption of the radio interfaces with that of the other peripheral elements by first performing a base measurement with all the radio interfaces down. Results are given in percentage of battery consumed by the radio interface. All the measurements are reiterated several times and the mean value is given as final result (a low coefficient of variation was observed).

The energy consumption of a radio interface is made up of a constant consumption, and of a per-packet consumption relating to the bitrate of the transmission. We call the former value base consumption, and infer the per-packet communication consumption from several measurements at different bitrates. Details about performed measurements may be found in [7] for Bluetooth and WiFi, while we used similar testbed for GPRS. Consumption curves are given in Figure 1(a). As shown, GPRS is the least energy-efficient radio network with the lowest bandwidth capability. Figure 1(b) further shows that the Bluetooth interface is less consuming than the WiFi interface for bitrates smaller than 400 kbit/s. Nevertheless, the Bluetooth's high per-packet cost quickly renders the WiFi interface more attractive. Moreover, the lengths of the curves on Figure 1(a) clearly show the huge difference in available bandwidth (i.e., 1 Mbit/s for Bluetooth versus 10 Mbit/s for WiFi). Thus, if while using the Bluetooth interface, the required bitrate increases over time, one may quickly face hardware non feasibility whereas WiFi could sustain a much larger bandwidth.



a) Power consumption of the radio interfaces after 30 min. according to bitrate. Scale : WiFi bandwidth
 b) Power consumption of the radio interfaces zoomed around Bluetooth bandwidth

Fig. 1. Power consumption of Bluetooth and WiFi interfaces according to bitrate

3.2 Issues in Energy-efficient Network Configuration

In B3G networks, devices may discover and access services using several radio interfaces. Thus, each device must decide which interface to use, for each requested service, with respect to the following criteria: (i) the connectivity must be available and exploitable; (ii) the link must provide enough bandwidth; and (iii) energy consumption, which is crucial for battery-powered devices, must be minimized. Meanwhile, in complex situations, the third criterion is not sufficient by itself. Indeed, local energy optimization (i.e., minimizing the energy consumption on each client) may lead to a non-optimal global situation, as illustrated below.

Consider a client application running on a B3G-capable device that wants to discover and access a service. The interface to use is chosen among the available radio interfaces that provide a sufficient bandwidth with respect to the required throughput. The most energy efficient interface, according to the measurements described above, shall be used. It should be noticed that, in this situation, service discovery and service access are independent and done sequentially. If needed, the interface chosen for service discovery may be switched off before the one chosen for service access is switched on. Thus, the radio interface(s) chosen for service discovery and service access may be different. It is also conceivable to discover the service using an interface that is not available on the provider in the case some networked intermediary contributes to the protocol (e.g., directory/cache node). Since the needed throughput for service discovery is very low (and almost null in the case of passive discovery), all the available interfaces on the client are eligible to process this operation. Thus, the interface with the lowest energy consumption in the idle state (i.e., Bluetooth according to previous results) shall be used. Service access shall further consider the connectivity on the provider side. Then, the interface having the lowest consumption at mean service throughput, among those offering sufficient bandwidth, shall be chosen.

When several services are to be accessed concurrently from the same device, repeating the above method as is, upon each new service access results in configuring the first service accesses on the same radio interface (i.e., on the most energy efficient radio interface), until the remaining bandwidth on this interface becomes too low to accept a new service access. When a new service access request arrives, configuration takes place on the second most energy efficient radio interface. Having two interfaces switched on concurrently significantly increases the global energy consumption, whereas accessing the services on another radio interface with a larger bandwidth could be more efficient. Figure 2 depicts such a situation. The first column represents the initial situation, where 3 services are accessed using the most energy efficient interface (e.g., Bluetooth), filling its bandwidth. When a new service is added, it must be configured on another interface (e.g., WiFi). Thus, the WiFi interface has to be switched on (Figure 2, second column). Reconfiguring the local multi-radio network so that all the services are accessed over WiFi would allow to save energy (Figure 2, right-most column).

Therefore, when such a situation happens, the client should reconfigure its local multi-radio network, by taking into account all the concurrent service accesses as a whole.




Fig. 2. Reconfiguration upon new service access

However, if the services requested by the client are located on several providers, the most energy-efficient network configuration from the standpoint of the client device may imply a higher energy consumption on the providers and thus a higher global consumption. As an example scenario, consider two devices providing four services to one client (e.g., device *A* provides services 1 and 2; device *B* provides services 3 and 4). Because




Fig. 3. An early example

of the high bandwidth requirement of service 1 (e.g., video streaming), the only usable radio interface is WiFi. According to the aforementioned method, the client decides to fill WiFi's bandwidth by also using it to access services 2 and 3, and switches onto the Bluetooth interface to access the last service (as depicted in Figure 3.a). Meanwhile, device *B* has to switch on both its WiFi and Bluetooth interfaces to provide services 3 and 4. To overcome this issue, a global decision among the client and the providers shall be taken. The client would be informed of the poor situation its choice involves for the provider. Then, it could decide to use its Bluetooth interface to access both services 3 and 4, resulting in a better global energy efficiency, at the expense of a small increase in the client's consumption (as depicted in Figure 3.b). The same kind of reasoning may be used, the other way round, when configuring several clients requiring service access on one provider.

As a conclusion, a decision about the most energy-efficient configuration for a real-scale system, composed of several clients discovering and concurrently accessing services on several providers, has to face all the aforementioned issues at once. That is, when a client wants to access a new service, its decision shall take

into account both its concurrent service access and the provider’s preferences. The decision may result in a complete reconfiguration of either or both the client and the provider, which may result in the need of a system-scale reconfiguration. The next section formalizes the issue of energy-efficient multi-radio networking, which will serve as a basis to study supporting configuration algorithms.

4 Energy-efficient Multi-radio Networks

In this section we give the formal model of the problem previously discussed. We call such a problem *Minimum Energy Communication Interfaces (MERC I)*. We also provide an Integer Linear Programming formulation (ILP) [16, 25] that optimally solves the problem for small instances.

ILP is a well known optimization technique that aims to solve problems in which one seeks to minimize or maximize a linear function by systematically choosing the values of variables, both integer and non-integer, from within an allowed set. Differently to Linear Programming (i.e., variable values are not integer), which can be solved efficiently in the worst case, ILP problems are in the worst case undecidable. Then, we give some hints of the fact that *MERC I* is a *NP*-hard problem hence an approximated solution might be required. Such an approximation can be obtained by means of the relaxed version of the ILP formulation. This means removing the integrity constraints (see [16, 25] for further details).

4.1 Model

Figure 4 illustrates a multi-radio system composed of an arbitrary number of nodes. The nodes communicate with each other by means of five different wireless network interfaces, namely Bluetooth PAN mode (BT/i), Bluetooth Piconet mode (BT/a), GPRS, WiFi infrastructure mode (WF/i) and WiFi ad-hoc mode (WF/a).

We define a connection graph $G = (V, E)$ where vertices in V are the network nodes, and where an edge $e = \{v, w\} \in E$ exists between two nodes $v, w \in V$ if and only if v and w are neighbors and share at least one network interface (i.e., communicate). Given a node v , in the following we will denote by $N_G(v)$ the neighborhood of v in G . In particular, given $G = (V, E)$ we associate to each node $v \in V$ one variable for each possible interface $v_{bt/i}, v_{bt/a}, v_{gprs}, v_{wf/i}, v_{wf/a}$ and a further one to express its current available energy v_{pow} . The activation of one interface is translated in setting the corresponding variable to 1 (represented




Fig. 4. Graph representation

as black circles in Figure 4), otherwise its value is 0 (represented as white circles in Figure 4). Note that a variable is 0 whenever the corresponding node does not provide the related interface.

As introduced in the previous section, the energy consumption induced by transmitting data on a given interface j is specified by the function $cost_j$. This is linear with respect to the amount of bandwidth used for

the communication and considers a fixed cost of activation⁵. Let $a_{bt/i}$, $a_{bt/a}$, a_{gprs} , $a_{wf/i}$, $a_{wf/a}$ (resp. $p_{bt/i}$, $p_{bt/a}$, p_{gprs} , $p_{wf/i}$, $p_{wf/a}$) be the activation costs (resp. the energy consumption per bandwidth expressed by means of Kbps) of the corresponding interfaces. The cost function is defined as follows:

Definition 1 (Cost). *Given a node $v \in V$, the energy consumption for the usage of its interface $j \in \{bt/i, bt/a, gprs, wf/i, wf/a\}$ is:*

$$cost_j^v(k_j^v) = a_j^v + p_j^v \cdot k_j^v$$

where:

- a_j^v is the energy needed by node v to switch on the interface j .
- p_j^v is the energy needed by node v to provide 1 Kbps of bandwidth.
- k_j^v is the bandwidth provided by the interface j of node v .

In this setting, our objective is to minimize the overall energy cost among the whole graph G while satisfying the required connections. Yet other constraints come from bandwidth requirements of the nodes, which we denote by K_j^v for each node v and interface j . The problem can be formulated as an ILP as follows:

$$\min \sum_{v \in V} \sum_{j \in \{bt/i, bt/a, gprs, wf/i, wf/a\}} cost_j^v(k_j^v) \times v_j$$

Subject to the constraints implied by the graph G ,

- (1) $\forall v \in V$
 $v_{bt/i} + v_{bt/a} \leq 1$
 $v_{wf/i} + v_{wf/a} \leq 1$
- (2) $\forall v \in V$, $\sum_{j \in \{bt/i, bt/a, gprs, wf/i, wf/a\}} cost_j^v(k_j^v) \times v_j < v_{pow}$
- (3) $\forall v, w \in V$ such that $\{v, w\} \in E$:
 $\sum_{j \in \{bt/i, bt/a, gprs, wf/i, wf/a\}} v_j w_j \geq 1$
- (4) $\forall v \in V$ and $\forall j \in \{bt/i, bt/a, gprs, wf/i, wf/a\}$
 $k_j^v \geq \sum_{w \in N_G(v)} K_j^w$
- (5) $\forall v \in V$ and $\forall j \in \{bt/i, bt/a, gprs, wf/i, wf/a\}$
 $v_j \in \{0, 1\}$
 $k_{bt/i}^v, k_{bt/a}^v \in [0, 1Mbps]$
 $k_{gprs}^v \in [0, 172Kbps]$
 $k_{wf/i}^v, k_{wf/a}^v \in [0, 11Mbps]$

Where the constraints expressed by (1) imply that a node cannot activate neither both Bluetooth types bt/i and bt/a nor both WiFi types wf/i and wf/a at the same time. (2) restricts the possibilities of a generic node v to activate some interface if its available energy v_{pow} is not enough. (3) means that there must exist at least one common interface among two nodes v and w that are connected by an edge in G . (4) implies that each node v must provide enough bandwidth to meet the requirements of all its neighbors $N_G(v)$ (expressed by K_j^w). (5) defines the scope of the involved variables.

After those constraints dictated by the structure of the graph G , we also have to add the constraints dictated by the capabilities of the nodes. Namely, if a node v requires some high bandwidth K^v for which for instance Bluetooth is not enough, then in our formulation there will appear the further constraints $v_{bt/i} = 0$ and $v_{bt/a} = 0$. Other constraints can come from the fact that the residual energy of a node v is not enough to activate some interface like for instance wf/i or v does simply not provide such a service. In both cases it is translated in setting $v_{wf/i} = 0$.

⁵ It is worth noting that this extreme simplification has been made for keeping the model as simple as possible. Indeed energy consumption is affected by many other factors such as number of participants, network traffic load, packet collisions, as well as environmental conditions (i.e., obstacles, radio frequency interferences). See for instance [2, 11, 15, 19].

4.2 Analysis

The formalization presented above is an Integer Linear Program (ILP) [16, 25]. Usually such kind of formulations can cope with small instances since the required time to solve them might be exponential in the size of the input. This happens when the problem is *NP*-hard. In what follows we give an intuition about the complexity of *MERCI*, showing that the underlying decisional problem can be as hard as *SATISFABILITY* (*SAT* for brevity), known to be *NP*-complete [4, 24].

Lemma 1. *MERCI is NP-hard.*

Proof. [Sketch] In order to provide an easy intuition of the complexity of the problem, let us consider a slightly different problem (we denote it by *m-MERCI*) in which each interface, if activated, precludes the activation of another one (like between bt/i and bt/a). We also remove from the original problem any argument about bandwidth or residual energy.⁶ We show that deciding if there exists a solution for such a modified version of *MERCI* is of the same complexity of solving the *SAT* problem, hence incurring in a contradiction, unless $P \equiv NP$. Given an instance of *SAT* we provide a polynomial reduction to *m-MERCI* and we show that a solution for *m-MERCI* would be a solution for *SAT*. An instance of *SAT* is composed by a set of clauses of binary variables that must be satisfied at the same time by means of a single assignment of values to those variables. From an instance of *SAT* we construct an instance of *m-MERCI* as follows. For each variable X in the *SAT* instance we assume two interfaces X/i and X/a in *m-MERCI* being mutual exclusive and corresponding to X and $\neg X$ respectively. For each clause in *SAT* we introduce a node in *m-MERCI* with the interfaces corresponding to the appearing variables in the clause. Two nodes of our instance will be joint by at least one edge if they share an interface X or its complementary one $\neg X$. In order to satisfy the built links, a solution for *m-MERCI* should provide a common interface between the corresponding nodes. Note that this is equivalent to make true the corresponding clauses associated to such nodes by setting to 1 the activated interface. This gives a good intuition about the fact that answering to the *m-MERCI* problem would provide the same answer for *SAT*. \square

From the previous lemma, it follows that for large instances of *MERCI*, approximated solutions would be preferred. Those can be obtained by means of the standard relaxation methodology [16, 25]. This implies modifying the constraints (4) of the previous ILP formulation hence allowing the corresponding variables to assume real values (not only integer) among the given intervals. Note that the relaxed version can be polynomially solved since it is composed by a polynomial number of variables and constraints with respect to the input graph. Once it is solved, the variables must be changed back to integer values. The quality of the obtained solution in terms of approximation factors by means of such a methodology will be material of future works.

Another approach could be related to best effort algorithms. It can happen, in fact, that an input instance does not have a solution but still we would like to build the network satisfying as many connections as possible. This would lead to apply approximation algorithms like the one for *MAX-SAT* that provides a 0.77-approximation factor [3]. This means that if an optimal solution can simultaneously satisfy m clauses (and this is the maximum possible) by means of an assignment for the variables composing the given clauses, the algorithm finds an assignment that guarantees the satisfaction of at least $0.77m$ clauses.

5 Conclusion

B3G networks combine multiple wireless networking technologies in order to benefit from their respective advantages and specifics. The increase in computing and communication capacities of portable devices, as well as their mass marketing, further allow envisaging the widespread deployment of multi-radio pervasive networks that compose the functionalities of the networked nodes, which range from base sensors/actuators to Internet servers. Such pervasive networks open tremendous challenges for the development and deployment of distributed systems. A user having a multi-radio capable device benefits from such a pervasive network

⁶ We can simply assume that each node has always enough bandwidth and energy to satisfy the related requirements.

by increasing the perimeter of reachable service providers, but this is at the expense of a higher network management complexity. This complexity, induced by the heterogeneity of the wireless technologies, should be hidden to the user and, to be effective, to the application (e.g., by a middleware solution). Towards offering such an integrated multi-radio network management, this paper has presented a thorough analysis of the energetic performance of multi-radio networks. Indeed, making the network energy-efficient is a key requirement of the pervasive computing vision since this increases the autonomy of nodes and thus services availability. This paper has more specifically focused on service-oriented multi-radio networks, as we consider service-orientation as a prime enabler of the pervasive computing vision due to the openness of the computing environment that it enables.

The energetic performance of multi-radio networks is dependent upon the respective energetic performance of its constituting radio networks. Then, a trivial approach to making the network energy-efficient is to choose the most energy-efficient radio link for each service session, provided the radio link meets the session's bandwidth requirements. However, such a solution does not account for the global energetic performance of the network, considering both concurrent network access on the terminal and network usage on peers. Formal modeling of the properties of energy-efficient service-oriented multi-radio networks enables us to show that the problem is NP-hard and thus requires a careful approximation. Dynamic configuration of energy-efficient service-oriented multi-radio networks is part of our future work in the context of the PLASTIC project. We specifically aim at developing a middleware-layer solution, closely coupled with the lower network layers for the sake of effectiveness. We further aim at a decentralized solution due to the high dynamics of the networking environment.

References

1. A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Z. A. Multi-radio unification protocol for IEEE 802.11 wireless networks. In *Proceedings of the 1st International Conference on Broadband Networks (BroadNets)*, 2004.
2. J. N. Al-Karaki and A. E. Kamal. Routing techniques in wireless sensor networks: a survey. *IEEE Wireless Communications*, 11(6), 2004.
3. T. Asano, K. Hori, T. Ono, and T. Hirata. A theoretical framework of hybrid approaches to MAX SAT. In *Proc. of the 8th Symposium on Algorithms and Computation (ISAAC)*, volume 1350 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
4. G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer-Verlag, 1999.
5. P. Bahl, A. Adya, J. Padhye, and A. Wolman. Reconsidering wireless systems with multiple radios. *Computer Communication Review*, 34(5), 2004.
6. Y.-D. Bromberg and V. Issarny. Indiss: Interoperable discovery system for networked services. In *Proceedings of Middleware*, 2005.
7. D. Charlet, V. Issarny, and R. Chibout. Service discovery in multi-radio networks. In *Proceedings of the 9th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM)*, 2006. To appear.
8. D. Garlan, D. Siewiorek, A. Smailagic, and P. Steenkiste. Project Aura: Towards distraction-free pervasive computing. *IEEE Pervasive Computing*, 1(2), 2002.
9. V. Issarny, D. Sacchetti, F. Tartanoglu, F. Sailhan, R. Chibout, N. Levy, and A. Talamona. Developing ambient intelligence systems: A solution based on Web services. *Journal of Automated Software Engineering*, 12, 2005.
10. IST PLASTIC Consortium. Providing dependable and adaptive service technology for pervasive information and communication. <http://www.ist-plastic.org/>.
11. C. Jones, K. Sivalingam, P. Argawal, and J. Chen. A survey of energy efficient network protocols for wireless networks. *Wireless Networks*, 7, 2001.
12. J. Lorchat and T. Noel. Power performance comparison of heterogeneous wireless network interfaces. In *Vehicular Technology Conference 2003 (VTC 2003)*, October 2003.
13. S. B. Mokhtar, N. Georgantas, and V. Issarny. Cocoa: Conversation-based service composition in pervasive computing environments. In *Proceedings of the IEEE International Conference on Pervasive Services*, 2006.
14. S. B. Mokhtar, N. Georgantas, and V. Issarny. Efficient semantic service discovery in pervasive computing environments. In *Proceedings of the ACM/IFIP/USENIX 7th International Middleware Conference*, 2006. To appear.

15. A. Navarra. 3-D Minimum Energy Broadcasting. In *Proc. of the 13th Colloquium on Structural Information and Communication Complexity (SIROCCO)*, volume 4056 of *Lecture Notes in Computer Science*. Springer-Verlag, 2006.
16. E. D. Nering and A. W. Tucker. *Linear Programs and related problems*. Computer Science and Scientific Computing Series. Academic Press Inc., 1993.
17. T. Pering, V. Raghunathan, and R. Want. Exploiting radio hierarchies for power-efficient wireless device discovery and connection setup. In *Proceedings of the 18th International Conference on VLSI Design (VLSID)*, 2005.
18. W. Qadeer, T. S. Rosing, J. Ankorn, V. Krishnan, and G. D. Micheli. Heterogeneous wireless network management. In *PACS 2003*, volume 3164. Lecture Notes in Computer Science, 2004.
19. T. S. Rappaport. *Wireless communications: principles and practice*. Prentice-Hall, Englewood Cliffs, NY, 1996.
20. P.-G. Raverdy, V. Issarny, A. de La Chapelle, and R. Chibout. A multi-protocol approach to service discovery and access in pervasive environments. In *Proceedings of Mobiquitous*, 2006.
21. J.-Z. S. Rieki, J. Jurmu, and M. Sauvola. Adaptive connectivity management middleware for heterogeneous wireless networks. *Wireless Communications*, 12(6), 2005.
22. M. Roman, C. K. Hess, R. Cerqueira, A. Ranganat, R. H. Campbell, and K. Nahrstedt. Gaia: A middleware infrastructure for active spaces. *IEEE Pervasive Computing*, 1, 2002.
23. J. Sachs, L. Muoz, R. Aguero, J. Choque, G. Koudouri, R. Karimi, L. J. J. Gebert, F. Meago, and F. Berggren. Future wireless communication based on multi-radio access. In *Future wireless communication based on multi-radio access (WWRf)*, 2004.
24. V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.
25. L. A. Wolsey. *Integer Programming*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., 1998.