

Invited Article

Minimalist Self-Organization in Wireless Networks

Sébastien Roy, Philippe Leroux

Department of Electrical and Computer Engineering, Université Laval, Québec, Canada

Correspondence: Sébastien Roy, sebasroy@gel.ulaval.ca, s.roy@ieee.org

Manuscript communication: received 23 October 2010, accepted 12 January 2011

Abstract— Many fields of human endeavour, such as biology and the theory of complex systems, are now embracing the concept of self-organization based on local actions leading to a desirable global emergent behavior. While many examples, both natural and artificial, can be found of such self-organized systems, the relationship between the local rules and the global behavior remains elusive and no systematic procedure is known to engineer a specific global result. Given the increasing pervasiveness of wireless networks of all sorts, including ad hoc networks competing within narrow unlicensed bands and wireless sensor networks, self-organization could constitute the next defining paradigm in wireless communications. It can be shown that a set of heuristic principles can be leveraged to engineer a self-organized connection-oriented wireless network with minimal complexity. Such a system requires no centralization of information, yet achieves a nearly optimal global state with only a modest amount of local signaling. It will naturally and jointly balance the many parameters related to radio resource management, exhibiting great adaptability, fault tolerance and scalability.

Keywords— Multi-agents, distributed base stations, emergent behavior, channel allocation, channel segregation, interference.

This work was presented in part as a keynote speech at the 2008 International Conference on Advanced Technologies for Communications (ATC), Hanoi, Vietnam.

1 INTRODUCTION

1.1 Wireless Networks as Distributed Software Systems

Modern wireless networks, whether they are random access local-area networks (WiFi), cellular networks, or wireless sensor networks, are conceptually similar to distributed computers. Terminals and access points are typically capable of sophisticated processing and collaborate among themselves so that the network as a whole provides one or more useful functions. Therefore, the entire network is in essence a computer whose main function is to ensure that its parts can communicate effectively among themselves in spite of spectral congestion and the inherent difficulties associated with the wireless channel. It would typically achieve this by 1- collecting information on the environment, such as channel quality, degree of interference, and communication requirements (bit rate, quality of service, etc.), and 2- computing a ‘good’ overall solution according to some fitness criterion determining how the resources (space, time, channels, access radio interfaces) will be allocated.

Clearly, from a global point of view, the task of the network is to dynamically and continuously compute the solution of a multidimensional constrained optimization problem. This problem is of the “traveling salesman” variety and does not scale well, since it is NP-complete. Therefore, current practical systems are typically content to split the problem into smaller parts and obtain a suboptimal solution.

We can think of the higher reasoning functions of the network associated with resource allocation as software-based, whether or not they are actually implemented in software. This allows us to view a wireless network essentially as a platform supporting a parallel and distributed software system. With this point of view, we can draw insights from the computer science community, which has long been concerned with the complexities of distributed systems, and see how they apply to the wireless context.

1.2 Increasing Software Localization

In [1], Parunak observed that the history of software is one of “increasing localization and encapsulation.” Initially, there was the monolithic program where the control flow would jump arbitrarily from one point to the next, making it difficult to define and manipulate any element other than a single line of code or the entire program. With the advent of structured programming came subroutines, which were tightly written, independent and reusable portions of code performing a well-delimited function. While a subroutine’s behavior (code) is encapsulated, its state must be supplied externally (through arguments), and it becomes active only when invoked by a call. Then came object-oriented programming where the code and its state were both encapsulated in the form of an *object*. While this provided many advantages, objects are still passive and gain control only when they receive an external message.

The natural evolution of objects is to cut the last

link to centralized control by allowing them to invoke themselves and to have their own internal goals. Hence, behavior, state and invocation are all localized, leading to the concept of a *software agent*. The agent is therefore a small piece of localized, independent software that acts “on behalf of” the system as a whole, and thus has full authority to determine which action is appropriate on its own. Relevant variants of this concept include *intelligent agents* (which exhibit some form of reasoning or learning), *autonomous agent* (which can invoke itself and take appropriate action depending on circumstances), *distributed agents* (a set of agents implemented on several physically distinct machines), and *multi-agent systems* (distributed agents which need to communicate among themselves and collaborate to accomplish their objective).

Visionary economist and writer Ernst Friedrich Schumaker said

“The system of nature, of which man is a part, tends to be self-balancing, self-adjusting, self-cleansing. Not so with technology.”

This quote emphasizes the embarrassing failure of man-made technology to adapt to even the slightest change that was not foreseen by its creators, something that nature accomplishes with ease. The ultimate vision behind software localization is to avoid the complexities associated with traditional top-down design and to provide some measure of *dynamic self-organization*, in response to changing circumstances. This capability automatically brings many benefits, including robustness, fault tolerance, flexibility, ease of maintenance, and others.

1.3 Emergent Behavior

In essence, this vision of multi-agent systems falls under the umbrella of *nature-inspired computing*, and is related to cellular automata, swarm intelligence, ant colonies, and amorphous computing. In a living organism, each cell has its own internal goals, yet alters its behavior in response to external information from its environment. It also alters the environment by liberating chemicals. A collection of such cells, each acting of its own accord and being influenced by, and influencing its environment, can achieve a high degree of structure and perform a high-level function which is significantly more sophisticated than the simple tasks performed by the individual cells.

This is emergent behavior, whereby a complex system or pattern emerges out of a great number of simple actions and interactions. Conway’s game of life [2] constitutes the classic example. When emergent behavior is observed, dynamic self-organization is normally implicit, with all the advantages that it brings. However, it is challenging to *engineer* a specific emergent behavior, since there is no known recipe or precise systematic method of deriving local behaviors and interactions backwards from the desired global behavior.

An important goal in agent architectures is to provide some degree of self-organization and dynamic adaptation capability without resorting to top-down

control from a higher system authority. There are researchers who implement these features with sophisticated agents that explicitly reason about their interactions, much as a human would. The drawback of this approach is that it reintroduces many problems associated with complex system design that motivated increased software localization in the first place [1].

Thus, engineering emergent behavior from multiple simple, minimalist agents is a more promising approach, in spite of its engineering challenges. The latter are not, however, insurmountable; there might not be any systematic procedure to derive the agent’s behaviors, but there are heuristic guidelines which can be applied in practice.

1.4 Application to Wireless Networks

The deployment and maintenance of a conventional cellular wireless system are complex, demanding tasks. Rigid organization is favored in the classical cellular paradigm, with specific subsets of channels being assigned to each cell. This limits the co-channel interference, but also restricts the solution space, thus leading to suboptimal resource allocations. The size and shape of a cell, as well as the height of the base station tower are important parameters which depend on the propagation environment, obstacles, the density of wireless traffic, the presence of interference sources, etc. Thus, deployment typically requires extensive planning and coverage measurements. As the environment changes (new interference sources, higher density of wireless traffic, changes in the propagation environment due to new buildings, etc.), the design choices gradually become invalid and the system must be periodically upgraded / re-engineered.

On the other hand, 802.11 wireless local area networks require little in the way of planning and can be readily deployed in an ad-hoc manner. However, their random-access approach makes it difficult to guarantee any type of quality of service, and the network collapses under the weight of collisions as the traffic offering goes above a certain fraction of the theoretical capacity. They are also plagued by the so-called “hidden terminal” problem.

Despite these shortcomings, 802.11 networks enjoy considerable commercial success and this is no doubt linked to their low cost and ease of deployment. The 802.11 “distributed coordination function” (DCF) which embodies the basic channel random-access mechanism (based on carrier sense multiple access with collision avoidance – CSMA/CA), is a primitive example of distributed self-organization.

Self-organization through emergent behavior in the context of wireless systems has the potential to simply and efficiently arrive at better resource allocation solutions in a robust, dynamically adaptive manner, while providing the requested quality of service if required. Indeed, quality of service just needs to be factored into the global resource allocation problem. If such a problem is tackled in a distributed manner as described, it matters little how many dimensions and

constraints are present. It becomes possible to include all variables in a single optimization problem, something which is not feasible in conventional systems for various reasons, including the difficulty in centralizing all relevant information at the decision point. Such a centralization is unnecessary in multi-agent systems based on swarming (MASS).

The increasingly relevant usage of multiple antennas for beamforming and spatial multiplexing (MIMO processing) also adds degrees of freedom to the optimization problem, making the MASS approach even more appealing.

2 THE DISTRIBUTED BASE STATION PARADIGM

In the distributed base station (DBS) paradigm, large central base stations are replaced by a plurality of smaller access points more or less evenly distributed in space so that their areas of coverage overlap significantly. This implies that any given terminal can be served by multiple access points simultaneously, thus leveraging the benefits of macrodiversity. This collection of access points, which provides access to the network for the terminal, plays the role of a single base station in a conventional system.

In the DBS concept, we assume that each access point is linked to the other access points via a set of links which is independent from the channels used to communicate with terminals. This “infrastructure network” can be wired or wireless (on a different band), and is used by the access points to coordinate their work, as well as process and forward data to and from terminals. On the uplink for example, the infrastructure network will route data from many access points to a concentrator node which will perform diversity combining and possibly other signal processing tasks to estimate the transmitted packet from a given terminal.

Thus, there are three tiers of devices arranged in a bottom-up hierarchy: terminals, access points, and concentrator nodes, the latter being connected to a larger wired network such as the Internet. This constitutes a conceptual framework which can be reduced to practice in many different ways. It will be discussed herein as a means of providing connection-oriented communications, as in a cellular system, but could equally apply, with appropriate changes, in the context of wireless LANs or wireless sensor networks. Also, the infrastructure network could be implemented as a form of wireless mesh, thus enabling rapid deployment, but many other forms are possible.

Most importantly, the DBS paradigm is designed to capitalize on self-organization to offer some of the best features of both cellular (robust quality of service, effective management of multiple channels, efficient spectrum usage) and wireless LANs (rapid easy deployment and maintenance, low cost), while leveraging macrodiversity and offering a high degree of robustness against changes in the environment (including access point failures).

2.1 DBS with Centralized Management

Unlike what is proposed here, it is noteworthy that there is a body of literature on so-called distributed base stations or distributed arrays which tie in to a centralizing “brain” via fiber optics. While this is related to the DBS paradigm herein and is known to have many advantages (see [3] and references therein), it is also very different since it consists in a centrally located conventional base station to which remote radio heads (RRHs) are connected. Thus, the physical antennas might be separated physically, but the decision power is centralized. This has the drawback of requiring more complex scheduling and management algorithms, as in traditional cellular systems, of requiring huge bandwidth to centralize the information (hence the fiber optic connections), and of being very difficult to scale up.

One recent paper [4] brings this concept closer to what is presented herein by comparing distributed versus centralized scheduling algorithms in the context of RRHs linked to a central point by fiber optics for the uplink of 3G-LTE (Long Term Evolution) cellular systems. It is therein assumed that RRHs are imbued with local decisional power, and it is shown that the distributed approach outperforms the centralized approach in terms of sum-throughput. Thus, the superiority of a distributed approach seems to be demonstrated in the context of an existing, practical cellular standard. However, the approach described is deterministic, semi-centralized and is not self-organized. What is proposed here goes much further in terms of exploiting synergy and emergent behavior to derive unprecedented benefits in the areas of robustness, scalability, self-organization, and implementation simplicity.

2.2 Connection Management

In such a system, a mobile establishes its link to the network via multiple stations simultaneously. Not all DBS may relay all mobiles, and minimalist strategies such as relaying the closest mobiles prevent the system from adapting towards more globally optimal solutions. It was shown in [8] that connections can be easily managed by agents corresponding to the DBS behaving autonomously, with their local actions leading to an emergent behavior which optimizes connections and QoS for the overall network. This agent system is herein used with no modification. Cooperating mobiles are seen as a single connection for a DBS (since they use the same channel).

2.3 Interference Management

Interference in a DBS scheme is complex. There exists a macrodiversity versus reuseability compromise which comes from the fact that a mobile using macrodiversity will have a larger spatial footprint on the channel it is using (due to the many links). On the other hand, the more compact such footprints are, the more mobiles it is possible to accommodate since the reuse distance is smaller. At first glance, the two aspects seem to be

in opposition. However, macrodiversity should enhance the overall link quality of mobiles thus leading to resource savings by lowering the transmit power (with respect to what would be required to achieve the same link quality with a single access point) and hence the spatial footprint.

The challenge here is to deal with many connections (relaying connections for macro-diversity) in assigning channels to mobiles while optimizing the overall network's performance.

2.4 Cooperation

In the context of DBS, we see that all aspects - macrodiversity, connection and interference management - are entangled in such a way that the behavior of each distributed agent has an effect on the overall network in a non-linear and intricate manner. Adding cooperation will necessarily add to these complex relations. However, and because the whole system has been designed to adapt to what it senses, the design equation boils down to the design of an additional agent which will work fairly independently and handle cooperation. Connection and interference agents should then be able to adapt to the new conditions. Especially, the channel allocation agent should be able to efficiently make use of the channels freed by cooperating mobiles (since two mobiles will then share a channel).

3 ENGINEERING SWARMING

Cellular systems achieve spectral efficiency and robust quality of service through precise, deterministic planning and resource allocation. However, this is costly and suboptimal (in terms of channel usage, for example, in the case of a rigid per-cell allocation) and adaptation to changes (such as major environment changes or scaling to meet increased traffic demands) is difficult. Wireless LANs on the other hand are cheap, easy to deploy and maintain, but have difficulty with quality of service and are spectrally inefficient. Indeed, the efficiency is typically of one third in loaded networks and can fall to as low as one tenth in certain cases (because of collisions, transmission errors, idle periods and preamble overhead) [5].

We aim for an intermediate solution between these two extremes, capitalizing on self-organization. Access points are disseminated in such a way that their coverage areas overlap significantly. Such access points would have higher hardware complexity than 802.11 counterparts, but still be considerably simpler and cheaper than cellular base stations. They would not need the same power or height as standard base stations, nor the same coverage planning effort, because of the macrodiversity gain. On the other hand, each access point is capable of operating on multiple channels simultaneously (unlike 802.11 access points). Said channels are considered at an abstract level, and can in practice consist of time slots, frequency bands, orthogonal codes, or combinations thereof. It will simply be assumed that these channels are orthogonal among

themselves, and that each access point can address a maximum number N of such channels at any given time (where N could be construed to indicate the number of effective radio interfaces).

In this context, many interrelated optimization problems crop up, such as how to assign a multiplicity of access points to any given terminal, given the limit N , in such a way that all terminals desiring their connection can obtain it with the desired quality of service. If we only take into account for the time being the limit of N connections per access point and ignore the issue of actual channel allocation and the effect of self-interference due to channel reuse, we already have an NP-complete problem, which, for a large number of access points, cannot practically be tackled in an optimal, centralized, deterministic manner. We will therefore resort to multi-agent systems with swarming (MASS) as a distributed computing platform to address this *connection problem*.

The environment, through which the agents communicate, is an integral part of a MASS system. In our case, the environment is the wireless medium and it is noteworthy that any multi-agent system essentially reduces to three elements: the agents themselves, the environment, and a coupling relationship between them. The addition of the swarming component is, according to Parunak, obtained by incorporating three principles:

1. the *coupling* principle, whereby processes within the agents must continually exchange information via the environment;
2. the *autocatalysis* principle, whereby this interaction is self-maintained and self-perpetuating;
3. the *function* principle, whereby the organization thus induced provides a useful emergent function.

In the following, we will provide some indications on the reduction to practice of Parunak's swarming principle within the DBS context.

3.1 Coupling

We've seen that the coupling is realized via an environment where sources of information are distributed, and organized behavior can only emerge through continuous exchanges. In the DBS paradigm, the access points and terminals constitute information sources distributed within the wireless medium, the latter being the environment. For example, terminals constitute sources of information on connection requests and QoS requirements.

In ant colonies, the coupling leading to emergent behavior is realized by the ants leaving behind a trail of pheromones. In artificial systems, the literature refers to this as "volatile markers," which are bits of information left in the environment which decay in time (as their relevance diminishes) and eventually disappear. Terminals send requests for connections with a given level of QoS, and such requests can be seen as messages (markers) impressed on the environment (the wireless medium). The request disappears with distance, since its relevance diminishes as we move further away from its source (given that the terminal cannot be served by a

remote access points). One or more access point will respond to the request and provide a diversity link to the terminal. This action is in turn sensed by other access points in the vicinity and will impact their behavior. In this manner, local actions ripple through the system, impacting on an increasingly larger neighborhood, thus leading to a globally emergent behavior.

3.2 Autocatalysis

The autocatalysis principle implies that agents retrieve, process, and send out information in a continuous fashion. The agents must not wait for some event or discrete state transition, but must favor continuous flow. In this context, both *amplification* and *limitation* mechanisms are called for. Amplification refers to positive feedback such that convergence towards a desirable solution is accelerated. Any action by an agent which is deemed to move the system in a positive direction should therefore trigger similar actions in nearby agents, thus accelerating convergence through a form of viral contagion.

However, if amplification is allowed to function without bounds, the system could easily accelerate towards an undesirable state or suboptimal solution. So a limitation mechanism is also required to achieve balance. In our wireless DBS context, a poor signal quality will favor multiple access point connections. If this statement was allowed to take effect without bounds, a very poor signal would soon have connections to hundreds of access points. But this is naturally limited by the maximum number N of connections per AP, signal attenuation through distance, eventually the obtention of a “good enough” connection, and other factors.

Also, the system must keep priming itself to explore the solution space even if it reaches a so-called “sweet spot,” since the latter might be suboptimal, and the environment is likely to change rapidly. This continuous dynamic optimization is guaranteed by maintaining a *flow of entropy*. In the DBS context, connections are continuously created and destroyed to create and maintain such a flow. Furthermore, the random activation times of the agents ensures that the system does not fall into any periodically-repeating pattern (such as observed with Conway’s classic game of life, where all cells are activated simultaneously).

3.3 Function

Autocatalysis ensures that the interactions discussed under coupling are self-sustaining, but not necessarily that they are useful. There is no known precise, systematic procedure to derive local behaviors based on the desired useful global function. There are, however, heuristic approaches and one such approach consists in finding a utility function which can be evaluated locally by agents and that maps to the global desired behavior. Finding such a function is pretty much a matter of trial and error, again guided by heuristics. The utility function in effect translates the flow of entropy into desirable decisions. It can be seen as embodying a

multi-dimensional problem within a single dimension (scalar) quantity.

In our DBS system, link quality is rated according to such a continuous function to determine the best links to establish and / or sever at each agent activation.

4 CONNECTION AGENTS

The goal is to design minimalist, low-complexity agents, which will address global optimization problems in a totally distributed manner, avoiding any explicit centralization of information (which would consume significant bandwidth and represent a considerable overhead in storage, network activity, and computing power). It might at first seem that a wireless communication system in which most major functions are totally distributed would almost collapse under the overhead of control signaling. However, that is not the case with the MASS approach because decisions are taken at the local level.

Agents must operate entirely based on information collected continuously from the environment, not on explicit signaling from a system operator. The connection agents [6], implemented at the access points, have access to a minimal amount of information:

1. the set of terminals present and within range, either requesting connection or already connected to the network;
2. the quality of service requested by each;
3. the level of QoS each connected terminal currently enjoys;
4. the incremental diversity gain the access point can bring to a currently connected terminal.

For the sake of simplicity and without loss of generality, error probability is used as an index to measure both QoS and diversity gain. Quality of service can mean latency, bit rate, and / or bit error probability (BER). Since we are dealing with a connection-oriented network, latency is guaranteed. Furthermore, there is a direct tradeoff between BER and bit rate using adaptive modulation. For these reasons, BER is chosen as the sole QoS index and the modulation will be assumed fixed.

Furthermore, it was shown in [7] that when employing maximal-ratio combining (MRC) in a white Gaussian noise environment, the BER could be approximated by a product form. This implies that, on a logarithmic scale, the BER figures observed on individual links add up when the said links are combined using MRC.

In this context, the following notation is used:

- $P_e(i, m)$ is the logarithm of the BER of the link from terminal m to access point i and also the incremental diversity gain access point i can bring to mobile m .
- $P_{\text{tot}}(m) = \sum_i P_e(i, m)$ is the total error probability for terminal m . This quantity is set to 0 if the terminal is not connected to the network.
- $P_{\text{req}}(m)$ is the logarithm of the error probability requested by terminal m , and constitutes its required QoS.

- P_{thresh} is a logarithmic BER threshold below which the terminal is not considered in range and will not be relayed by the access point.

While it is more convenient in analysis and simulation to work exclusively with the BER, it might clearly be easier in an actual physical implementation to measure the signal-to-noise ratio and deduce the BER from there (since for a given modulation type, there is a direct monotonic relation between the two).

We have devised the following utility function to rate links for connection and disconnection:

$$D(i, m) = \left(\frac{P_{\text{tot}}(m)}{P_{\text{req}}(m)} \right) \times \log \left(\frac{P_{\text{tot}}(m)}{P_e(i, m)} \right), \quad (1)$$

where the first factor translates the QoS actually provided with respect to the requested QoS and the second factor (logarithmically scaled) measures the relative contribution of access point i to terminal m with respect to its overall connection quality.

When evaluating for disconnection, the access point will disconnect from the best-served terminal, which is the terminal with the maximum $D(i, m)$. According to the first factor, this implies that the terminal has a good overall connection quality with respect to its requested QoS, and, according to the second factor, the relative contribution of AP i is low.

Likewise, when evaluating for connection, the goal is to establish a link to the worst-served terminal, i.e. the one with the lowest value of $D(i, m)$. A low value implies that the terminal has a poor overall connection quality and / or the relative contribution (incremental diversity gain) of AP i would be important.

The simple act of normalizing the first factor of the utility function by the requested QoS is sufficient to implement the support of QoS classes within the system.

The agent is activated at random intervals according to a Poisson distribution. It has three distinct behaviors, depending on the number of active links the AP currently has:

1. Initialization: If there are fewer than $N - n - 1$ connections, several connections are established to the closest terminals (with the lowest $P_e(i, m)$) until the limit N is reached.
2. Connection: If there are fewer than N connections, the utility function is evaluated and a link is established to the terminal in range with the lowest $D(i, m)$.
3. Disconnection: If there are more than $N - n$ connections, the mobile with the highest $D(i, m)$ is disconnected.

The parameter n is normally set to 1.

Extensive simulations have been performed to characterize the emergent behavior of the connection agents. One scenario is based on 200 terminals randomly and uniformly disseminated over a $5 \text{ km} \times 5 \text{ km}$ area. A set of 19 access points are deployed to cover this area according to a predetermined, uniform, "honeycomb" pattern, i.e. the APs are in approximately the same positions as if each was in the center of a traditional

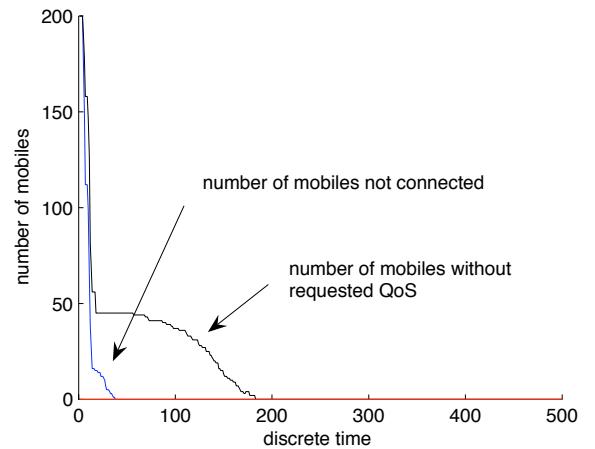


Figure 1. Transient and steady-state behavior when $N = 25$, maximum total number of connections=475.

hexagonal cell. The mean trigger rate of the agents is set to $\lambda = 3$ time units.

As far as propagation parameters go, the propagation exponent is 4, i.e. signal strength decays proportionally to $1/d^4$. Furthermore, Rayleigh fading is assumed except when the link length is less than 100 meters, in which case a line of sight is assumed to exist, thus leading to Rice fading with $K = 5$ dB. This propagation model is assumed throughout this paper.

For best results, the system must have sufficient degrees of freedom to favor macrodiversity, i.e. $19 \times N \gg 200$. Figure 1 shows simulation results when $N = 25$. Three behavioral phases can be clearly distinguished. In the first phase, terminals are rapidly connected to the nearest access points. During this phase, which ends at time index 8, most terminals are connected, but a quarter of them still don't have their requested QoS. In the second phase, links are broken and rearranged to improve the global solution until, at time index 190, nearly all terminals achieve their QoS requirements. In the steady-state, the MASS system achieves the same performance as an optimal, omniscient, centralized algorithm used for comparison purposes. It was assumed that the QoS requirement for all terminals corresponded to a BER level of 10^{-3} or better.

For the same simulation scenario, Figure 2 shows a sorted profile of the terminals according to their achieved QoS (P_{tot}) when two classes of terminals are present. Half of the terminals request a BER of $10^{-2.5}$ (class 1), and half of the terminals request a BER of 10^{-5} (class 2). It can be seen that the MASS tends to distribute resources so that terminals are more evenly served than with the centralized algorithm.

5 CHANNEL ALLOCATION AGENTS

It was noted that the connection agent was developed by ignoring the effect of co-channel interference due to channel reuse. However, the nature of MASS is such that, without changing the connection agent, the simple addition of further agents to handle channel allocation is sufficient to effectively manage interference. Indeed,

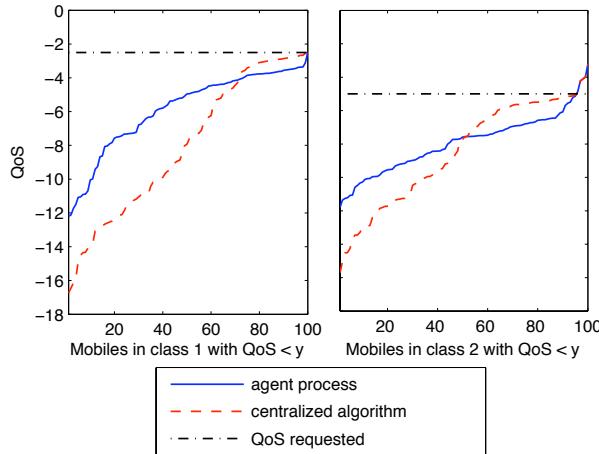


Figure 2. Sorted profiles with 100 mobiles in class 1, 100 mobiles in class 2, and $N = 25$.

the connection agent is robust to changes and will naturally interact with the additional agents in such a way that the emergent behavior tends towards global multidimensional optimization.

To effectively address the channel allocation problem, it was found useful to introduce the concept of *master link*. Hence, any given terminal being connected to the network by linking to multiple APs, one of the said links (and, by extension, the associated AP) is designated as *master*. A terminal's master AP is responsible, jointly with the terminal itself, for selecting the terminal's channel.

The channel allocation function emerges from interactions between 4 types of agents [8], of which two are located at the terminal, and two at the AP:

- **Agents A** are located at the AP and their job is to make channel suggestions to the terminals.
- **Agents B** are located at the terminal; they receive suggestions from agents A, evaluate the quality of the suggested channel from the terminal's point of view, and accept or refuse the said channel.
- **Agents C** are located at the AP; they select, among the terminals who share a master link with the AP, one terminal for a channel change, said channel change being initiated afterwards by agent A.
- **Agents D** are located at the terminal and are responsible for choosing one AP as master among the set of APs linking the terminal to the network.

Figure 3 outlines the interactions among all five agent types, including the connection agent. Agent A chooses a channel to suggest to the terminal, and the strategy employed to select the said channel can greatly influence the system. Three selection behaviors have been considered:

1. **Segregation:** a strategy based on Akaiwa's *channel segregation* algorithm, where channels are chosen according to how frequently they are used.
2. **Clairvoyant:** a strategy where it is assumed that the AP possesses perfect information on the interference-plus-noise power on all channels. This strategy will first select the channel with the least

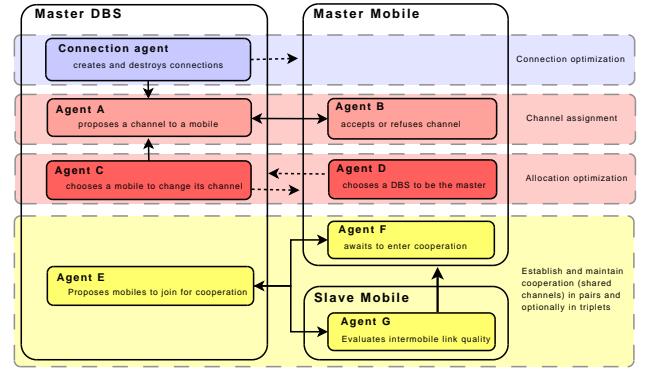


Figure 3. Diagram depicting the various interactions among the different agent types in the DBS architecture.

interference. While it is possible to approach this strategy in practice, it is considered less practical since it would require continuous monitoring of signal power on all channels.

3. **Random:** a strategy where channels are simply selected at random, until a suitable channel is found.

Agent A's suggested channel may seem adequate from the AP's point-of-view, but it must be validated at the terminal as well. Even though the AP and the terminal are very often relatively close, the interference pattern perceived by one and the other on a given channel is not the same. For this reason, agent B at the terminal receives agent A's suggestion and evaluates the SINR on that channel. If it is found to be below a threshold T , the suggestion is rejected and agent A must try again at the next activation.

Agent D looks after the selection of the master AP. At each activation, it selects the AP offering the best link quality (from the terminal's point of view) as the new master. Agent C selects at each activation one terminal among the set of master links, and this terminal is then targeted for a channel change. This implements a form of frequency-hopping which establishes a flow of entropy and forces the system to keep exploring the solution space. The idea is to select a terminal who is currently under-served. Three different criteria have been tested in simulation:

1. according to the weakest link quality to the master AP;
2. according to the weakest overall *connection* quality, including macrodiversity;
3. according to the strongest interference power, as measured by the master AP.

Then, agent A will propose a new channel to the targeted terminal. This is done as described above, with the additional constraint that the proposed channel must present less interference power than the currently used channel. Also, agent A will only examine a maximum of $\text{ChTry}_{\text{MAX}}$ channels before giving up (thus aborting the channel change). This scheme establishes the flow of entropy allowing for continuous and iterative improvement of the global mutual interference pattern. Hence, this is also the aspect of the system which allows it to react favorably to changes, e.g. due

to mobility, failures, new terminals, etc.

Like the connection agents, agents C and D are self-activated at random intervals according to a Poisson distribution with a mean interactivation time of λ . On the other hand, agents A and B must be explicitly activated either by agent C or by the connection agent.

Extensive simulations have shown that the joint work of all five agents is effective at managing both connections and channel allocations, regardless of the particular behaviors chosen (although performance varies slightly). Again, the addition of an additional aspect and additional agents integrates naturally with the existing ones and the entire system remains stable and robust. Furthermore, the proposed MASS approach scales effortlessly to any system size, so long as there are enough resources available to support some degree of macrodiversity.

6 COOPERATION AGENTS

6.1 Assumptions and Motivation

Additional agents (agents E, F, and G) have been devised for pairwise and triplet-wise cooperation (described in detail in [9]).

Technical and economical constraints restrict mobiles to embed a single half-duplex air interface. However, the ability to rapidly switch from receive to transmit and vice-versa is assumed.

Cooperation is considered for two or three mobiles. The three-mobile case can be considered as a linear superposition of mobiles cooperating two at a time.

The proposed method is not meant to improve range, or capacity, it is meant to maintain link quality while freeing radio resources. The benefit of cooperation in this case is synergistic and in line with the underpinning principles of the proposed DBS paradigm. Mobiles which will cooperate should be close enough to simplify synchronization and minimize the effect of intermobile communication on diversity, but also to keep the spatial interference footprint as compact as possible.

6.2 Cooperation Code

The following code, adapted from [10], is used:

M_1	μx_1	■	$\nu \tilde{x}_2^*$
M_2	■	μx_2	\tilde{x}_1^*

Both mobiles share the same channel. First M_1 transmits block x_1 while M_2 listens. Then, M_2 transmits x_2 while M_1 listens. In the last phase, both mobiles transmit modified versions of the symbols observed in the listening phase such that a "half-duplex" version of Alamouti's code is realized. The scaling factors $\mu = \sqrt{2}$ and $\nu = 1$ result in a total transmit energy of 3 times the symbol power over 3 time slots. Moreover, the excess amount of energy used during the broadcasting slots enables higher SINR, thus lowering the probability of a detection error at the opposite mobile.

The estimated symbol at the mobile is

$$\tilde{x}_i = x_i + \frac{z'_i}{b}, \quad (2)$$

where b is the channel coefficient between mobile 1 and 2. During each time slot, the received signal at the DBS is given by

$$y_1 = a_1 x_1 + z_1, \quad (3)$$

$$y_2 = a_2 x_2 + z_2, \quad (4)$$

$$y_3 = a_1 \tilde{x}_1^* - a_2 \tilde{x}_2^* + z_3, \quad (5)$$

and the signal decoding is performed (in Alamouti-like fashion [11]) according to

$$\tilde{x}_1 = a_1^* \frac{y_1 + y_2}{\mu} - a_2 \frac{y_3^*}{\nu} \quad (6)$$

$$= \left(|a_1|^2 + |a_2|^2 \right) x_1 + a_1^* \frac{z_1 + z_2}{\mu} - a_2 a_1^* \frac{z'_1}{b} + \\ |a_2|^2 \frac{z'_2}{b} - a_2 z_3^* \quad (7)$$

A similar rule applies for decoded signal \tilde{x}_2 [9].

Three-way cooperation in the same spirit is also possible according to the following code structure:

M_1	μx_1	■	■	$\nu \tilde{x}_2^*$	■	- $\nu \tilde{x}_3^*$
M_2	■	μx_2	■	- $\nu \tilde{x}_1^*$	$\nu \tilde{x}_3^*$	■
M_3	■	■	μx_3	■	- $\nu \tilde{x}_2^*$	$\nu \tilde{x}_1^*$

with $\mu = \sqrt{\frac{18}{4}}$ and $\nu = \sqrt{\frac{6}{5}}$.

Cooperation being a kind of orthogonal repetition code, it is compared to an equivalent — in terms of code rate — (3,1) Hamming code for non-cooperating mobiles. 4-QAM (QPSK) modulation is employed for both two-way cooperation and no cooperation (Hamming code). In both cases, 2 bits are transmitted per 3 time units per mobile, and each mobile consumes the same transmit energy per bit. However, only one channel is occupied for 2 mobiles in pairwise cooperating mode, leading to higher spectral efficiency at the system level.

For three mobiles, six time slots are used but only 3 symbols transmitted, such that to maintain a constant number of bits per mobile per time unit with respect to the other modes, 16QAM modulation should be used (4 bits per mobile in 6 time slots = 2/3 bit per time unit).

6.3 Cooperation Agent Protocol

Consider that a mobile is connected to the network via a plurality of DBS, but that one of them constitutes its "master" connection. That particular DBS controls the mobile's channel, while the other DBS provide secondary, or what is referred to as "relay" connections. Cooperation is organized in a similar master-slave arrangement, where one mobile within a cooperating set is a master while the rest are slaves.

One agent located at the DBS will at each activation propose mobiles from its master connection list for cooperation. For some possible combinations of mobiles chosen randomly, where the mobiles' SINR at the DBS is above a predetermined threshold, i.e.

$$\text{SINR}_{M_1, \text{DBS}} > \text{Th}_B, \quad (8)$$

the DBS sends a request to a potential slave mobile to test the reception of a master mobile. If the slave senses that the intermobile link is good enough, i.e.

$$\text{SINR}_{M_1, M_2} > \text{Th}_{\text{coop}}, \quad (9)$$

it sends an acknowledgment to the DBS and the two mobiles start cooperating on the master mobile's channel. Therefore, the slave mobile is connected to the network via the connections allocated to the master mobile. Also, the slave mobile's channel is freed as well as the master and relay connections it had. This offers more relaying resources for macrodiversity for other mobiles.

Three-way cooperation occurs when a DBS proposes a mobile for association with two mobiles already cooperating. Then, the new candidate evaluates the two intermobile links to acknowledge whether it can cooperate or not. A second agent is situated at the mobile and activated when the mobile is in slave mode. The agent evaluates the intermobile link quality, and if it falls under a certain SNR, the mobile requests that cooperation be terminated. The master DBS then has to find a free channel to reconnect the mobile. Also, the mobile then momentarily loses all of its relaying connections, since the relay DBS must first evaluate if they can reconnect to the mobile on its new channel. This will necessarily be done at the next activation of the connection agents of each DBS. If the DBS fails to allocate a channel, the mobile is blocked.

Hence cooperation starts and ends, in an opportunistic fashion, throughout the network given the required sufficient conditions for cooperation between a set of mobiles are met. The agent protocol requires very little processing overhead, the complexity of the algorithms being minimal. It does, however, lead to more pressure on the channel allocation agent as mobiles join or end cooperation when moving randomly.

Simulation of the system is performed over a square field of 25 km². The field is wrapped around both ways to form a torus in order to avoid "edge" artifacts. One thousand mobiles evolve, moving in random directions, with individual speeds being drawn from a uniform distribution between 0 and 18 km/h at the beginning of the simulation.

Agents activate randomly with a mean interval of 3 seconds between activations, coordinating connections (see [6]), channel assignments [8] and managing cooperation links [9]. Mobiles are randomly positioned. One hundred DBS are also randomly scattered throughout the area, thus avoiding any form of regular topological pattern. Some regions of the field therefore enjoy very good coverage while others do not. The purpose of the configuration is to observe how well the agents are able to compensate the inequities and balance the resources across all mobiles. Each DBS can relay 25 mobiles. There are 50 shared channels, Th_{coop} is set to 25dB and Th_B to 10dB. The scenario is executed for 200 seconds, and the last 100 seconds are used to measure the means of the observed parameters (the first part is left aside to not account for the transitional state).

Figure 4 shows the sorted mean symbol error rate of

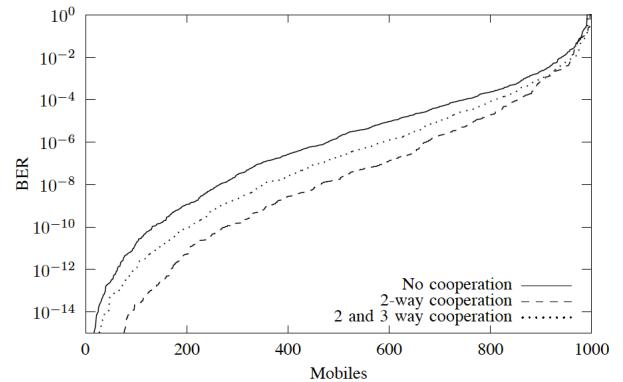


Figure 4. BER performance compared with and without cooperation (2 way, and both 2 and 3 way enabled).

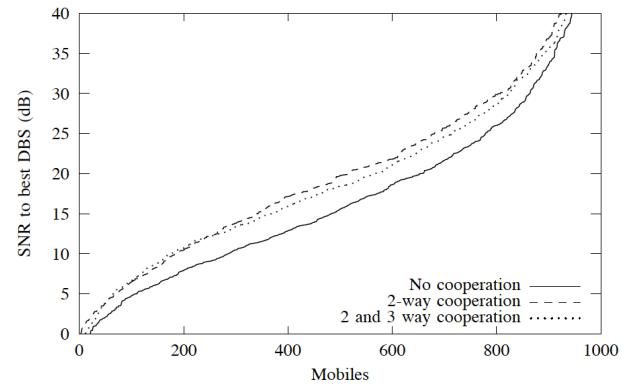


Figure 5. Effect of cooperation on mobiles' SNR.

each mobile (in effect approximating the density distribution over the network) in order to observe how well the multiple agents handle the various load balancing challenges. Therefore, the x axis represents the mobiles ordered according to the quality of their overall link (including the effect of macrodiversity). The y axis is the BER (on a \log_{10} scale). Curves are presented for the three cases of no cooperation, two-way cooperation only, and both 2 and 3 way cooperation.

Despite the Hamming code being more efficient than the two way cooperation in terms of BER at the same SNR, cooperation proves to be a better strategy. With three-way cooperation enabled, the advantage is even more obvious. Two aspects are involved. First, cooperation allows for better spectral efficiency such that the granularity of DBS and the ability of the channel allocation agents to adapt to changing interference patterns (while mobiles join and leave cooperation) lead to better SNRs for all mobiles even though not all of them engage in cooperation. Figure 5 illustrates the enhanced SNR by plotting the sorted mobiles' SNR (to their best relaying DBS) (mobiles order is not comparable to the previous figure). Mobiles often obtain an additional 5dB of SNR.

This factor alone does not fully characterize the benefits of cooperation. In the DBS model, macrodiversity is an important factor, and cooperation leverages the macrodiversity gain in a synergistic way. Indeed, the diversity obtained from cooperation multiplies the

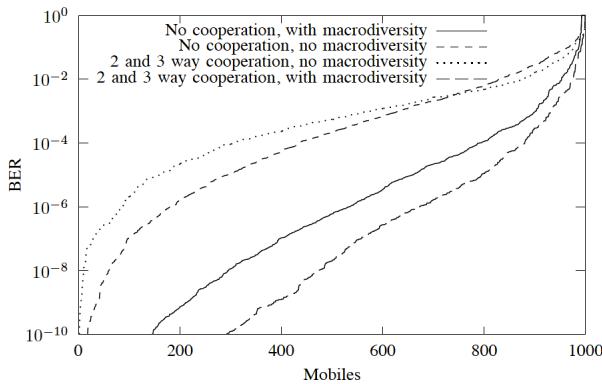


Figure 6. BER performance compared with and without cooperation (2 way, and both 2 and 3 way enabled).

diversity obtained from the multiple relaying links. Hence, 2-way cooperation with 2 relaying DBS would be better for high SNR than no cooperation with a (3,1) Hamming code. This is illustrated in Figure 6, where without macrodiversity, the gain of cooperation allows for poorly connected mobiles to enjoy higher SINR, at the expense of mobiles enjoying higher SINR which loose on BER. Although this allows for more evenly balanced resources, the benefits are more marginal. However with macrodiversity, cooperation enables improved BER for all mobiles with respect to a non-cooperating scheme.

Moreover, cooperation allows for DBS receivers to be freed, since only one is needed to receive two or three communications such that DBS can relay more mobiles. It also lowers the interference allowing for remote DBS to reach the signal and enable more macrodiversity links. This synergistic effect is captured by Table I, which gives the mean number D of macrodiversity links. It would take about 40 receivers per DBS, compared to 25 to obtain as many macrodiversity links without cooperation.

The channel management agents are also helped by the introduction of cooperation. Table I provides the number of blocked channel allocation attempts per simulation time step (B). This value is clearly reduced since more channels are made available by virtue of cooperation. Hence, fewer attempts are required to connect a mobile. This happens even though additional mobiles need to be assigned a new channel when cooperation links end and slave cooperating mobile need a new channel to remain connected.

Table I
PERFORMANCE PARAMETERS ARE GIVEN FOR THE THREE COOPERATION SCENARIOS (NO COOPERATION, 2-WAY COOPERATION, AND 3-WAY COOPERATION) INCLUDING MEAN NUMBER OF MACRODIVERSITY LINKS (D), MEAN NUMBER OF BLOCKED CHANNEL ALLOCATION ATTEMPTS (B), AND MEAN NUMBER OF DISCONNECTED OF COMPLETELY DISCONNECTED MOBILES (L).

	D	2Coop	3Coop	Join/leave	B	L
No coop	2.34				7.09	0.64
2 coop	3.18	566		7.8	6.23	0.75
2-3 coop	3.75	330	330	11.2	4.72	0.77

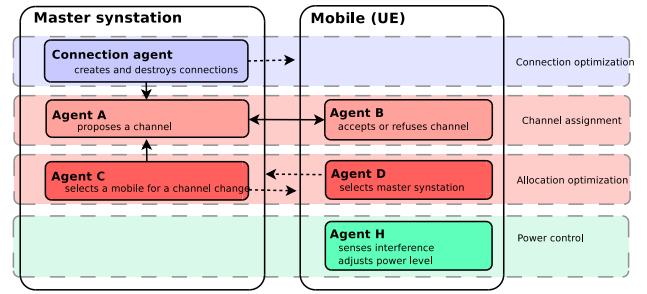


Figure 7. Diagram depicting the various interactions among the different agent types with the inclusion of agent class H for power control.

Table I also lists the number of mobiles cooperating in pairs, the number of mobiles cooperating in triplets, the mean number of mobiles joining and leaving cooperation per second, and the mean number L of mobiles which are completely disconnected per simulation time step. This value rises slightly in cooperation scenarios because of the effect of suddenly orphaned slave mobiles when a cooperation pair / triplet is disbanded.

7 POWER CONTROL AGENTS

The idea of introducing an additional agent at the mobile for power control purposes is a natural extension of the multi-agent platform described so far. Power control in various forms helps limit interference footprints in existing systems, thus saving energy and reducing further the interference levels affecting link quality. The approach taken here is to introduce an additional agent at the mobile which simply senses the energy on the currently used channel and adjusts the transmit power according to a simple technique described below. However, this approach is not used in conjunction with the cooperation agents at this stage. It turns out that the coupling between the power control and the cooperation functions is intimate and difficult to manage / predict, such that they cannot be combined in a trivial manner, unlike the previous classes of agents. This question therefore requires further investigation.

Thus, the power control functionality was studied in the absence of the cooperation agents and the applicable agent interaction scheme is depicted in Figure 7. The question naturally arises whether the interaction between power control and the other classes of agents could also be problematic. There is in fact a potential risk of instability if actions taken by one type of agent, i.e. power level adjustments, excessively impede or exacerbate another agent's action (connection / disconnection or channel allocation) in a manner that disrupts the overall balance of the system. Hence, the limitation aspect is important and it should ensure that reactions / counter-reactions are not excessive in any direction. For example, a sudden change in interference should not trigger wild oscillations in other agents within nearby access points and mobiles, but rather smooth, progressive adjustments.

Given that mobile m has a power level ratio of p_m , an associated overall link quality index is defined as

follows:

$$Q_m = \frac{P_{\text{tot}}(m)}{P_{\text{req}}(m)} = \frac{\log_{10} \text{BER}_m}{\log_{10} \text{BER}_m^R} = \frac{\sum_{k=1}^{M_m} \log_{10} \text{BER}(m, k)}{\log_{10} \text{BER}_m^R}. \quad (10)$$

The sum on k iterates through the M_m DBS relaying mobile m . As with the other agent classes, normalizing with respect to the logarithm of the requested BER ($P_{\text{req}}(m)$) naturally implements classes of QoS.

If Q_m is above 1, the mobile is considered *overserved* since it benefits from a better BER level than that requested. Hence, we could conclude that in the interest of preserving resources, such a mobile should reduce its power level. If Q is below 1, then a natural solution would be to increase the power level until $Q = 1$ is reached. However, this must be done with care and progressively since increasing the power level also increases the interference footprint and affects multiple nearby mobiles. The challenge here is to induce interactions that create an appropriate balance between all these considerations. Furthermore, automatically forcing overserved mobiles to lower their power level is not necessarily optimal, since there are cases where this excess link quality margin can be enjoyed without significant negative impact on neighbors, e.g. when the local traffic level is low compared to the available radio resources.

There are other subtle implications if the power level is naively reduced in the overserved case. In very good channel conditions, the power level could be reduced to the point where the signal level is close to the noise floor, making that mobile very vulnerable to a deep fade and the so-called hidden terminal effect. Indeed, if a mobile does not transmit with sufficient strength, nearby DBS will not detect it and will choose to use its channel, thus making it prone to a sudden interference flooding. Using the same logic, mobiles that cannot reach their requested BER / QoS would systematically reach their maximum transmit power, a solution point which may not be optimal from a systemic standpoint.

The above discussion highlights the need for a more subtle approach. To achieve this, we introduce the Need variable. It corresponds to the power level a given mobile m should converge to, determined by its current quality Q_m , provided that nothing changes in the mean time. It follows that the point of equilibrium (homeostasis) is reached when the current Need_m value equals the current power level p_m for all mobiles.

To ensure proper dynamic behavior, the Need variable incorporates an appropriate shaping function and is given by

$$\text{Need}_m(Q_m) = 2e^{-SQ_m} (SQ_m + 1) - 1, \quad (11)$$

where S is a scaling parameter applicable to Q_m which can be adjusted to allow mobiles to enjoy QoS higher than their requested BER. This implies, depending on the value of S , that the Need_m is not necessarily smaller than p_m when $Q_m > 1$. Simulations have shown that the optimal value of S to maximize QoS across the entire system is 0.8, and this holds across wide ranges

of conditions in terms of traffic, mobile speeds, and available resources.

The shape of the Need function is largely conditioned by the exponential in (11). As before, it is not the only choice, but experiment has shown it to be viable. The choice may in fact seem arbitrary since this function gives us the p_m the mobile should converge to irrespective of its distance from serving access points, propagation or interference conditions. It should be understood to make sense only in the context of a dynamic balancing act with other mobiles which also work according to the same guidelines.

In practice, this requires translating the Need_m value into a power level increment step to be applied at the current power control agent activation. The conversion to this delta / step value must be performed in such a way as to incorporate the capability to converge towards and attain a homeostasis point. It must also incorporate an inherent limitation mechanism to prevent overshoot and ensuing instabilities.

The convergence to an equilibrium point or homeostasis state is induced by making the delta value linearly or non-linearly proportional to $\text{Need}_m - p_m$. Hence, if we are far from this equilibrium (defined as the point where $\text{Need}_m = p_m$), we move quickly towards it; if we are close, we slow down, thus emulating classic steepest descent behavior. A relationship between linear and quadratic was found to be favorable to convergence and stability. In light of this, we initially define the step

$$\Delta_m^{(1)} = \beta \text{sign}(\text{Need}_m - p_m) |\text{Need}_m - p_m|^{1.5}, \quad (12)$$

where the constant β can be adjusted to obtain the desired compromise between convergence speed and stability. This relationship will be refined below to incorporate the limitation aspect.

It has been demonstrated empirically that our initial expression for Δ_m above leads to unacceptable instability for high values of β . However, it is possible to modify it so that high values of β can be used (resulting in fast convergence time) while stability is obtained through additional parameters. Indeed, a respectable convergence speed, obtained with values of β equal to or in excess of 5, is important to support mobility, a condition which led to instability with our initial definition (12). The desired result is achievable by scaling (12) by the current power level p_m and the desired power level Need_m . This ensures that changes in power levels remain very subtle when the values of p_m and / or Need_m are small. These are indeed the conditions where rapid changes can suddenly alter the interference patterns causing multiple broken links and sending exaggerated ripples throughout the system. Thus, our new definition is

$$\Delta_m^{(2)} = p_m |\text{Need}_m| \beta \text{sign}(\text{Need}_m - p_m) |\text{Need}_m - p_m|^{1.5}. \quad (13)$$

Finally, the delta value is additionally constrained to

avoid exceeding the valid power level range, i.e.

$$\begin{aligned}\Delta_m^{(2)} < 0 \rightarrow \Delta_m^{(3)} &= \max \left\{ \Delta_m^{(2)}, -\frac{p_m}{2} \right\}, \\ \Delta_m^{(2)} > 0 \rightarrow \Delta_m^{(3)} &= \min \left\{ \Delta_m^{(2)}, \frac{1}{2}(1-p_m) \right\}.\end{aligned}\quad (14)$$

At each activation of the mobile's PC agent, the power level is adjusted according to

$$p_m^{(\nu+1)} = p_m^{(\nu)} + \Delta_m^{(3)}, \quad (15)$$

where ν is the activation index.

To properly characterize the proposed power control scheme, two existing power control algorithms from the literature were adapted to the DBS context and used as benchmarks within the simulations. These algorithms were selected because they are both representative of the state of the art, yet are also well-suited to the characteristics of the DBS paradigm, including high-order macrodiversity. The first one is Grandhi's centralized power control algorithm [12], which seeks to maximize the lowest SIR in the system through an approach based on eigenanalysis. The algorithm, as applied to the DBS context, is described in [13]. The second benchmark algorithm is Yanikomeroglu's SIR-Balanced Macro Power Control (SBMPC) scheme [14]. Formulated for the context of distributed antennas in CDMA systems, it requires very little change to apply to the DBS paradigm. Again, this is detailed in [13].

As before, a square area of 25 km² is considered, containing a population of 1000 mobiles in motion, and 100 fixed DBS uniformly and randomly scattered throughout. Each DBS can relay a maximum of 25 mobiles simultaneously, making the average number of relay links per mobile in this scenario equal to 2.5. Mobiles move in random directions with random speed drawn from a uniform distribution between 0 and V_{\max} . A mobile's maximum transmit power is 1 W measured at 1 meter from the antenna. Thermal noise at the receiver affects a signal bandwidth of 30 kHz at an assumed temperature of 20°C, yielding $N_0 = -129$ dBW. The total number of available channels in the system is denoted Ch.

Simulations are run for 1000 seconds and repeated 10 times with different initializations of the geometry (DBS positions and mobile initial positions, directions and speeds). Time is discretized with a time step of 1 second. At each time step, physical parameters are evaluated (mobile positions, propagation, interference, BER, connection outage). Agents activate randomly according to a Poisson distribution with parameter $\lambda = 3$ time steps. At each time step, agents which activate evaluate their local state and take actions accordingly. Depending on the agent class, this can be adjusting the power level, creating / destroying a connection, or performing a channel assignment action).

At each time step, the set of QoS (total BER level given on a logarithmic scale) for each mobile is sorted, thus providing a snapshot in time of the distribution of the network resources across all mobiles. These sorted distributions are then averaged for all the time steps of the simulation. Given this information, it is then

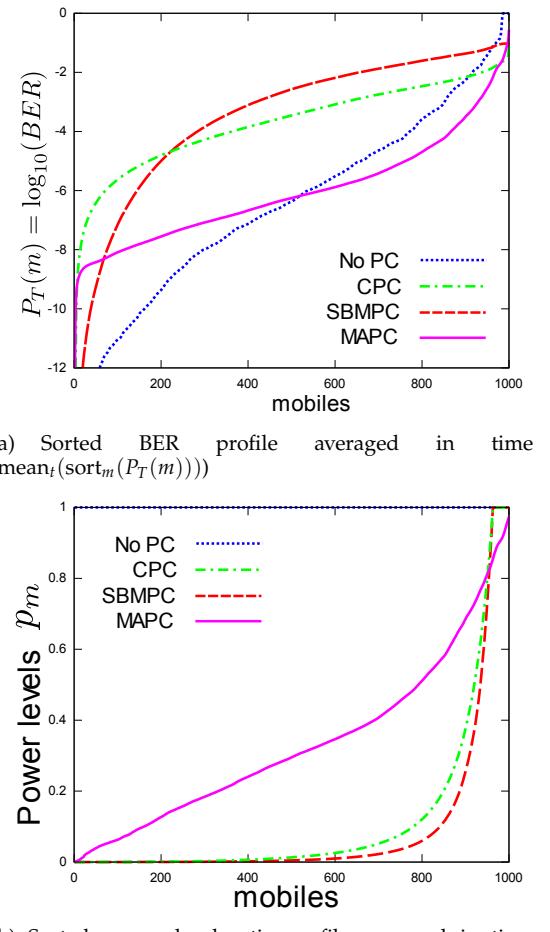


Figure 8. Simulation results with Ch = 40, N = 0, V_{max} = 0.

possible to compare how each algorithm distributes resources. The same is done for the power level allocation. Also, to verify the stability over time (considering the dynamic properties) of the algorithm, two parameters are interesting to observe to understand how the system handles outage:

- (1) the mean number \bar{N}_d of mobiles that loose all connections to the network per second, and
- (2) the mean time \bar{t}_r it takes for the network to reconnect a mobile after it has been disconnected.

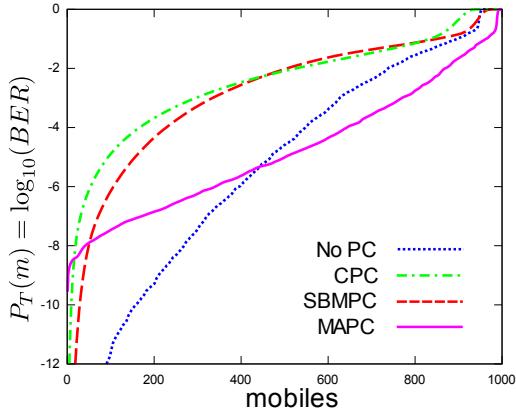
Figure 8(a) shows the base results, that is, with a static simulation where fading is considered as if mobiles were moving, but mobility is not considered (to observe a nominal capacity without taking into account dynamic adaptation of the algorithms). Noise is also not considered in this case.

The graphic reveals different aspects. First, with no PC, the QoS is clearly not balanced, but more importantly, not all mobiles can be connected as is seen from the right hand side of the graph. With the centralized algorithms, we can clearly see that QoS is balanced, and all mobiles are connected. MAPC on the other hand is able to provide much more QoS to almost all mobiles, while only impeding (compared to SBMPC) very few mobiles.

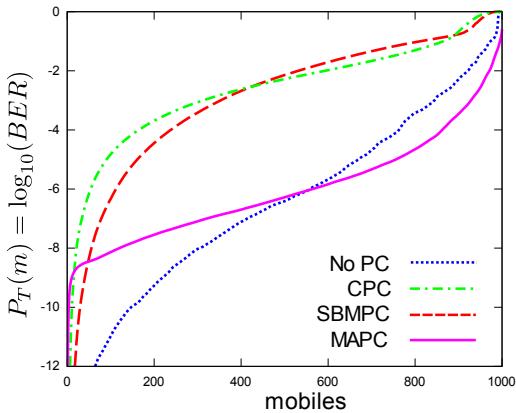
Values of the parameters \bar{N}_d and \bar{t}_r are given in Table II for this simulation scenario. One immediate

Table II
OUTAGE BEHAVIOR WITHOUT MOBILITY AND NOISE Ch = 40.

	No PC	CPC	SBMPC	MAPC
$\bar{N}_d (\times 10^{-6})$	570	8.9	130	190
\bar{t}_r (seconds)	24.2	1	2.0	2.65



(a) Low channel resources, Ch = 25, N = 0, V_{max} = 0.



(b) Effect of thermal noise, Ch = 40, N = -129dBW, V_{max} = 0.

Figure 9. Sorted BER profile averaged in time ($\text{mean}_t(\text{sort}_m(P_T(m)))$).

observation based on these values is that the CPC is extremely adept at handling outage. It takes only 1 time step (1 iteration of the simulation) to reconnect a lost mobile, and the probability that a mobile is disconnected is extremely low. Even SBMPC is not as good, but remains excellent compared to no power control. MAPC is doing only slightly worse in this respect, but with a much enhanced QoS provided to all mobiles.

Faced with higher interference levels (same user population but with Ch = 25), it can be observed that the centralized algorithms seem to collapse (Figure 9(a)). Indeed, under high interference levels, maximizing the minimum SIR leads to very poor SIRs for all mobiles. In turn, this generates many disconnections. Figure 9(a) clearly illustrates that the traditional algorithms are inefficient in this case, and lead to an even worse system snapshot than in the absence of PC. However, the MAPC algorithm still manages to provide acceptable levels of QoS, while connecting more mobiles.

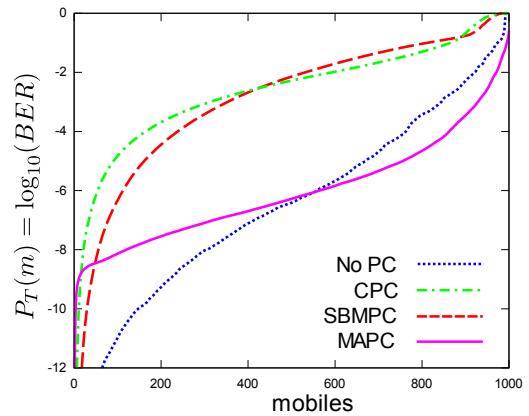


Figure 10. Sorted BER profile averaged in time ($\text{mean}_t(\text{sort}_m(P_T(m)))$), Ch = 40, N = -129dBW, V_{max} = 0.

The performance of the traditional algorithms deteriorates even more when noise is introduced. The fact that these algorithms focus on the SIR can make them unstable in the presence of noise, as Figure 10 attests, thus provoking widespread disconnections.

8 CONCLUSION

While many other aspects still need to be worked out (routing on the infrastructure network, reciprocal downlink operation, etc.), the described work provides a strong case for the feasibility of the MASS approach for radio resource management. The approach is scalable, easily deployed and extended, modular by function yet globally convergent, of disconcerting implementation simplicity, robust, flexible, and tolerant to faults. Multiple resource management aspects have been tackled through the introduction of various classes of agents: connection management, dynamic channel allocation (and inherent scheduling), pair-wise and triplet-wise cooperation, and finally, power control. In general, each new class of agent can be added on top of the previous one and designed independently, yet the different agents interact implicitly to reach an equilibrium. Thus, each new variable added in the system does not compound the optimization problem in the way one would expect with traditional centralized deterministic algorithms. For each class of agents, the performance is seen to be comparable, and in certain respects, better to that of known good centralized /traditional algorithms.

REFERENCES

- [1] H. Van Dyke Parunak, ““Go to the ant”: Engineering principles from natural multi-agent systems,” *Annals of Operations Research*, vol. 75, pp. 69–101, 1997.
- [2] M. Gardner, “The game of life, parts I-III,” in *Wheels, Life, and other Mathematical Amusements*. New York: W. H. Freeman, 1983, ch. 20-22.
- [3] E. Yaacoub and Z. Dawy, “Uplink scheduling in LTE systems using distributed base stations,” *European Transactions on Telecommunications*, vol. 21, no. 6, pp. 532–543, 2010.
- [4] ———, “Centralized and distributed LTE uplink scheduling in a distributed base station scenario,” in *Proc. Int.*

- Conference Advances in Computational Tools for Engineering Applications (ACTEA)*, Beirut, 15–17 July 2009, pp. 11–15.
- [5] Q. Ni, T. Li, T. Turletti, and Y. Xiao, “Saturation throughput analysis of error-prone 802.11 wireless networks,” *Wireless Communications and Mobile Computing*, vol. 5, no. 8, pp. 945–956, 2005.
 - [6] P. Leroux, S. Roy, and J.-Y. Chouinard, “An agent system to manage mobile connections in a distributed base station scheme,” in *Proc. 17th IEEE Int. Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Helsinki, 11–14 September 2006, pp. 1–6.
 - [7] ——, “The performance of soft macrodiversity based on maximal-ratio combining in uncorrelated Rician fading,” in *Proc. 17th IEEE Int. Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Helsinki, 11–14 September 2006, pp. 1–6.
 - [8] ——, “The impact of interference in a distributed base station scheme managed by an agent system,” in *Proc. 16th Int. Conference on Computer Communications and Networks (ICCCN)*, Honolulu, HI, 13–16 August 2007, pp. 961–966.
 - [9] ——, “Synergetic cooperation in a distributed base station system,” in *Proc. 19th IEEE Int. Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Cannes, France, 15–18 September 2008, pp. 1–6.
 - [10] A. Ribeiro, R. Wang, and G. B. Giannakis, “Multi-source cooperation with full-diversity spectral-efficiency and controllable-complexity,” in *Proc. 7th IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, Cannes, France, 2–5 July 2006, pp. 1–5.
 - [11] S. M. Alamouti, “A simple transmit diversity technique for wireless communications,” *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 8, pp. 1451–1458, October 1998.
 - [12] S. A. Grandhi, R. Vijayan, D. J. Goodman, and J. Zander, “Centralized power control in cellular radio systems,” *IEEE Transactions on Vehicular Technology*, vol. 42, no. 4, pp. 466–468, November 1993.
 - [13] P. Leroux and S. Roy, “Distributed power control with multiple agents in a distributed base station scheme using macrodiversity,” in *Proc. Int. Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, Lyon, France, 3–6 November 2009, pp. 500–514.
 - [14] H. Yanikomeroglu and E. S. Sousa, “SIR-balanced macro power control for the reverse link of CDMA sectorized distributed antenna system,” in *Proc. 9th IEEE Int. Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, vol. 2, Boston, MA, 8–11 September 1998, pp. 915–920.



Sébastien Roy received the B.Sc. and M.Sc. degrees in electrical engineering from Laval University, Québec, QC, Canada, in 1991 and 1993, respectively, and the Ph.D. degree from Carleton University, Ottawa, ON, Canada, in 2000. He is currently Full Professor with the Department of Electrical and Computer Engineering, Laval University, where he is pursuing research in the system-level and implementation aspects of signal processing for communications as well as space-time processing and space-time coding. From 2000 to 2002, he was a Natural Sciences and Engineering Research Council of Canada (NSERC) Postdoctoral Fellow at Laval University. He has also been active in industrial consulting with companies such as InterDigital and MacDonald Dettwiler, and was involved in the organization of several international conferences. In 2007 and 2009, he was an invited professor at lcole Nationale Suprieure de Sciences Appliques et de Technologie (ENSSAT), Lannion, France. He received 5 teaching awards and in 2007 received the award for excellence in technology transfer from the strategic network on Systems and Technologies for Advanced Communications (SYTACOM). Dr. Roy was also bestowed the award for Post-Graduate Research Excellence from the Canadian Institute for Telecommunications Research in 2000.



Philippe Leroux received the B.Sc from the electrical engineering school E.S.I.E.E, Paris, France in 2003 and the M.Sc. and Ph.D degree in electrical engineering from Laval University, Québec, QC, Canada in 2004 and 2008, respectively. He is currently a research associate within Prof. Roy's team working on developing the field of swarming agents to automate resource allocation at the physical layer of modern mobile networks. From 2009 to 2010, he was a post-doctorate fellow at LESO - EPFL, Lausanne, Switzerland, and participated in the development of a system-level simulator for urban metabolic models to apply swarming optimization techniques developed in his thesis towards sustainable development at the urban planning level. Throughout the multiple developments of simulators in his Ph.D. in telecommunication and Post-Doctoral work on urban physics, he has acquired expertise in system-level simulators and is now leading the development of a modern cellular communication simulator for the purpose of exploring engineered swarming applications.