

Multi-Agent Design for the Physical Layer of a Distributed Base Station Network

Philippe Leroux and Sébastien Roy
Université Laval
Canada

1. Introduction

As wireless networks are becoming more omnipresent and pervasive, appropriate resource allocation and organization becomes an increasingly pressing challenge. There exist on the consumer market two important types of wireless network technologies. On the one hand, cellular mobile networks are highly centralized and hierarchical. By contrast, wireless local area networks (WLANs) are deployed in an ad-hoc unstructured manner, thus avoiding the need for elaborate and costly planning. However, WLANs such as those falling under the highly successful 802.11 standard do not manage interference effectively and tend to collapse at high offered traffic loads. It can be seen that cellular and WLAN represent two radically different approaches in radio resource management, characterized by different sets of advantages and drawbacks.

The purpose of this chapter is to demonstrate the feasibility of a connection-oriented self-organized wireless system which offers efficient radio resource management and provides the best aspects of both cellular (reliable, connection-oriented operation even at high offered loads) and WLANs (ad-hoc deployment and distributed intelligence). This is achieved based on the multi-agent concept and local synergistic micro interaction (between neighboring transceivers) from which a global organization emerges.

The notion of Multiple Agent (MA) considered is of the “ant” variety, whereby small minimalist agents sense their environment and react to it in an interdependent manner. Social insects and mostly ants or bees are the most cited biological examples. In the literature such approaches have already been used to solve many combinatorial/optimization problems (Beongku et al., 2003; Brueckner & Parunak, 2003; Muraleedharan & Osadciw, 2003).

This design philosophy differs from more traditional approaches which consist in postulating criteria expressed by equations and models in order to formulate the problem in such a way that an optimal solution is derived within the defined context. In the agent approach, precise mathematical formulation of the problem is neither required nor very useful. The approach thus becomes attractive for tackling complex multidimensional problems which would otherwise be intractable. Therefore, our goal is not to demonstrate an optimal design, but to illustrate how a Multi Agent System (MAS) can be empirically designed and fine-tuned to fit a specific application. Moreover, it will be seen that such a dynamically adaptive solution, in spite of its empirical nature, offers many advantages over a rigid analytically-derived counterpart.

We will focus herein on Parunak's methodology (Parunak, 1997) because it offers an intuitive modeling framework, which is well suited to the empirical design approach.

Considering wireless networks, this chapter describes a flexible distributed base station (DBS) framework which removes many limitations of current networks in order to augment the solution space. For example, a plurality of DBS can simultaneously provide a network link to the same mobile, thus leveraging macrodiversity to improve link quality and/or achieve power savings. These DBS are designed with auto-organization in mind, such that the network structures itself autonomously. This is where MAS come in, offering the desired distributed intelligence, adaptability, scalability and auto configuration properties. However, the DBS architecture is challenging in at least three aspects:

1. It requires the continuously-updated solving of a large combinatorial problem, namely finding a good allocation of DBS resources to mobiles requiring service.
2. Interference must be handled in a transparent way so that mobiles can gain the best benefit of macrodiversity without being restrained by interfering mobiles.
3. Power control is an important aspect for both energy consumption and network capacity given that it is tightly-coupled with interference patterns.

These three aspects are entangled together such that an optimal allocation is a complex combinatorial problem. It is NP hard unless some heavy simplifying assumptions are made (on the geometry, on propagation, or other aspects). Moreover, in the context of mobility, an optimal solution at one point in time is not optimal if it cannot easily adapt to changing parameters (mobiles' positions, fading, etc.).

Yet, this complex context is well suited for a MAS design. Indeed, MA need an active environment in which to generate interaction. And each event of allocating power, channel or connections to mobiles, that a DBS generate, has consequences on other mobiles' links. This creates the required active environment in which agents can sense parameters such as the received power, interference and link quality, and where decisions can be made locally to generate new actions. In turn, the effect of these actions are sensed by other agents. The next section describes the challenges of the proposed DBS architecture. Then, MA design concepts used in this study are described. The fourth section details the proposed design of three categories of interacting agents respectively for :

1. macrodiversity connection management,
2. channel allocation, and
3. power level control.

Finally, the system is emulated. Results, including simulation of complex cases with randomly distributed DBS and mobile traffic, show first the resource allocation quality that can be obtained, and second the effectiveness of MAS design in terms of auto-configuration/scalability and dynamic adaptation properties.

A final brief discussion will extend Parunak's agent design principles to summarize the lessons learned from designing MAS for the application at hand.

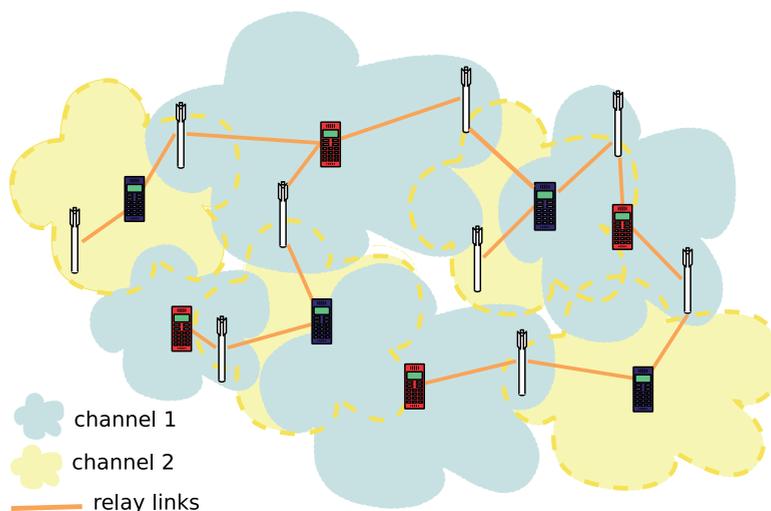


Fig. 1. Illustration of a DBS architecture exploiting 2 channels, and showing macrodiversity relay connections.

2. Challenges of the Distributed Base Station Network

2.1 Macrodiversity Potential

In a perfectly geometrical network with homogeneous traffic and symmetric propagation conditions, each DBS needs only to connect to the closest mobiles to maximize the provided quality of service. However, traffic is never homogeneous and varies across time and space in accordance with the users' schedules and patterns of usage. Moreover, propagation conditions are highly dependent on location, with varying availability of lines of sight and saturation of the frequency band due to heavy traffic. In such a context, there is a need for a simple, scalable, dynamic system to allocate relay links to mobiles and to continuously adapt the allocation pattern to changing conditions.

DBS can choose to relay mobiles far away from themselves in order to provide them with more macrodiversity, and thus better balance resource allocation. However, this choice involves a trade off. The exponentially-decaying link quality with the mobile-DBS distance could lead a remote mobile to consume many valuable relay links while deriving only marginal benefits, whereas closer mobiles would obtain much higher macrodiversity benefits from those same resources. Also, macrodiversity links provide not only enhanced overall quality links, but also reliability against network disconnection when undergoing severe fading, and it facilitates handover for mobiles moving outside from the range of some DBS to others. Therefore, a single criterion such as maximizing the minimum QoS for all mobiles would fail in certain conditions where enough resources would be available to provide the majority of mobiles with decent QoS, because of a few mobiles consuming much of these resources while deriving marginal benefits.

Such situations reveal the perils of pursuing a global solution based on a single perhaps overly simplistic quality criterion. In fact, many Pareto equilibrium solutions exist, in which no mobile can gain quality of service without stranding another user. And all these possible so-

lutions present multiple compromises on connection reliability and distribution of QoS. e.g. some solutions could favor maximizing overall signal strength for high transfer rate, others by distributing relaying links differently could prevent disconnections due to sudden strong fading or interference because mobiles would in general enjoy higher probability of being assigned multiple relay connections.

As such, it is not necessarily meaningful to define *a priori* goals for the search of a solution, as it is not known beforehand what are the benefits and drawbacks of each possible Pareto solution. This solution space is moreover hardly tractable due to the discrete nature of the problem, with a finite but large number of link resources to attribute. It is limited by physical conditions where some links may not be feasible due to the weakness of the considered signals. Also each link brings an increment of additional quality to the mobile's overall link quality, whose importance heavily depends on local propagation conditions that can vary continuously (with slow fading and changing mobile-DBS distances), or abruptly given the arrival of new connections in frequency allocation or strong fading situations. In this context, no analytically tractable mathematical framework exists leading to an optimum solution taking into account all the dimensions of the problem. More specifically, one must consider that an optimal solution, at a given state of the network and a given frame in time, could be too heavily specialized to that particular situation, such that a sudden change (strong fading, new mobiles joining the network) would make it ineffective. By analogy, it is known in biology that a species too well adapted to its environment is heavily endangered due to its limited capacity to adapt to environmental changes. Hence, a good solution is not an optimal one, but a good enough one that provides margins for adaptation in time to face changes.

There is a strong need for distributed techniques which are flexible enough to be tuned to the desired compromises while being able to handle unexpected events.

2.2 Channel allocation

Channel allocation faces the same propagation issues as connection management. To understand the implications of channel management, we introduce the concept of channel footprint. In a given situation (mobiles and DBS positions and relative densities, available channels, power allocations, etc.), a mobile's channel footprint can be understood as the space it occupies to maintain all its relaying links to DBS at a sufficient quality level. Hence, a second mobile, if emitting on the same channel inside this space, would affect some or all of the first mobile's connections.

In the cellular context, it is assumed that each mobile enjoys the same channel footprint which is controlled by the cell division of the space and an appropriate interference level threshold to allow or prevent the reuse of a channel across cells. In the 802.11 protocol, it is a handshake mechanism (the RTS/CTS exchange) which alerts neighboring transceivers that the channel will be in use, in order to control, to some extent, this channel footprint by preventing neighbors from reusing the channel in the vicinity, thus minimizing the hidden terminal effect (Ware et al., 2001).

In the DBS architecture, it would be appropriate that mobiles be offered varying channel footprints to adjust availability of channel resources and support the various needs of mobiles for macrodiversity. Indeed, mobiles needing more macrodiversity would require a larger footprint. Moreover, mobiles close to all of their relaying DBS should allow other mobiles to reuse the same channel at a closer range, compared to mobiles far from all DBS. This holds since these mobiles can support higher interference power and still maintain a good signal to interference plus noise ratio (SINR). This aspect of channel allocation was taken into consideration

in the dynamic allocation scheme known as the *umbrella cell system* (Furukawa & Akaiwa, 1994). However, channel reuse should not necessarily always be maximized, since it could lead to locally unnecessarily compacted channel allocation, even in low traffic conditions. Hence, channel allocation should be able to adapt to always balance resources to sustain the high throughput rate of modern communication services.

Another aspect to consider is the irregular geometry of the DBS network. The DBS density will necessarily change across space given some maximum or mean local load. And also, local mobile traffic varies across time and space (e.g. from residential neighborhood to office centers). This leads to constantly changing disparities in local loads of the network to which it must adapt the channel allocation. In the cellular world, solutions exist in the form of Dynamic Channel Allocation schemes (DCA), such as the well-known segregation scheme by Akaiwa & Andoh (1993), allowing cells with higher loads to use more channels.

Finally, faster channel allocation adaptation is required with high mobility, as mobiles move in and out of DBS' ranges. In cellular systems, this aspect is only tackled via a handover mechanism. However, in the DBS concept, DBS share channels via the macrodiversity relaying links. Hence, mobiles need not change their channel simply because they change a relaying link. This additional complexity implies channel segregation algorithms are not as efficient in the DBS concept.

2.3 Power allocation

Power level management can strongly affect the previously described aspects of connections and channel management. It can also offer a powerful means to leverage the possibilities in terms of resource management. Indeed, in the context of DBS, a higher power level implies modifying the mobile's channel footprint, and hence allowing it to reach perhaps more DBS for more macrodiversity. Or, on the contrary, a lower power level will leverage the channel's reuse possibilities by generating less interference. Power control can therefore help to provide much higher resource availability, and can provide synergistic behavior with the channel and connection allocation to leverage the possibilities that macrodiversity can offer to balance the resources across mobiles, throughout the network.

However, there potentially exists a maximum power level dynamic range which limits the effectiveness of power control. Indeed, decentralized DCA implies sensing the availability of channels. For example, the 802.11 protocol implements CSMA (Carrier-Sense Multiple Access) to prevent improper reuse of channels. In the cellular world with DCA, a maximum sensed interference power level threshold is considered to decide if the channel is available. Therefore, the dynamic range of power-level adjustments cannot exceed a certain range such that mobiles with low emission power needs are not disrupted by the interference caused by a new connection from another mobile, perhaps much further away, but with a much higher power level. It is noteworthy that this power level range is dependent on the space distribution of mobiles (and channel allocation). The range itself may not be constant at all scales such that it could be high throughout the entire network when compared to smaller areas of the network. Indeed, the SINR is a relative quantity which depends on the spatial distribution of access points and traffic.

Power level allocation is also a contention process where each mobile strives to maximize its QoS. Considering the viewpoint of one mobile, in order to maximize its QoS, it wants to be within reach of a maximum number of DBS to maximize macrodiversity, and it also wants to maximize its SINR for each of its links. Therefore, it wants to maximize its power level. However, if all mobiles did so, none would get any benefit. And since no mobile can obtain

any gains by reducing its power level, a non cooperative strategy is not a good choice for a game-theoretic approach to power level adaptation.

Necessarily, some mobiles will have to “accept” to reduce their power level in order to allow other mobiles in need to enjoy better QoS by reducing interference and enabling them to reach more DBS for macrodiversity. Yet, and due to the non linearity of the propagation environment, there necessarily is a point of diminishing returns for mobiles to reduce their power level. While any reduction necessarily implies a reduction in interference, the potential gain for other mobiles does not necessarily offset or compensate (given a compromise choice at a global scale) the loss in QoS for this mobile. It is to be understood here that there exist trade-offs for an infinity of Pareto solutions. Therefore, and again, postulating one global unidimensional criteria (e.g. as is done in traditional algorithms (Grandhi et al., 1993)) to derive a power allocation method would not allow assessment of the potential benefits of different trade-offs. Indeed, the results of the proposed design will show how the traditional approach to power control (which consists in maximizing the minimum SINR for all mobiles) in spite of offering interesting capabilities in some situations, also prevents most mobiles from achieving their QoS potential.

2.4 Complexity, Dynamics and Scalability

2.4.1 Complexity

In existing types of networks, the complexity is constrained by simplifying the hypotheses. For example, in cellular networks, channel allocation is simplified by segregating channels given an interference power threshold in order to guarantee a minimum SINR for all mobiles in a cell. This assumption simplifies the evaluation of provided QoS, as it guarantees a minimum QoS for connected mobiles, and avoids the hidden terminal effect, such that there only remains to evaluate the probability of a connection being blocked (when all channels in a cell or sector are occupied).

In the considered architecture, such assumptions are not made a priori as the purpose of the DBS architecture is to maximize flexibility. And considering the number of possible combinations of connections, or channels or even power levels, it is obvious that an exhaustive search to find all Pareto solutions is pointless. Even considering an exhaustive search in the case of a very simple scenario with only a few mobiles is pointless, since in such cases, the non-linear effects and interactions of large networks would not apply and the obtained results would be too limited to draw meaningful conclusions.

Also, postulating a unidimensional criterion and over-simplifying the non-linear effects involved, in order to provide a tractable mathematical framework would limit the solution space and therefore restrict the possibilities of such an architecture.

MA offer interesting properties to cope with complexity. The approach involves segmenting the problem into multiple subproblems where each is tackled by its own agent class. Heavy calculations for evaluating and selecting combinations are also avoided. Instead, specific combinations are attempted and modified by agents’ actions through local interactions.

2.4.2 Dynamics and Scalability

One particular aspect to consider is the fact that a given resource allocation solution must necessarily adapt to changes in a mobile wireless network. Such a solution must also adapt to unexpected events, such as the failure of a DBS. And finally it must scale, such that adding DBS locally will seamlessly, without any need for configuration, increase the capacity of the network in terms of either provided QoS or number of provided connections.

3. Multi-Agent Design

To solve the resource allocation problem, with the previously described considerations, multiple agents or bio-inspired optimization seems appropriate, as such approaches provide the most important sought-after characteristics, namely

- scalability ;
- dynamic adaptation ;
- auto-configuration ;
- reliability facing unexpected events.

Following Parunak's (Parunak, 1997) design principles, three main characteristics need to be provided in a MA design: coupling, auto-catalysis, and function. Coupling implies that each MA process is coupled directly or not to the others and their environment (e.g indirectly using pheromones via an environment). Auto-catalysis implies that the agents' actions taken in the right direction¹, by the nature of the agents' processes, favor similar actions leading the system to converge to a desirable state (positive feedback reinforcing the convergence towards the solution). And finally, the system must be such that a useful global function emerges out of the induced local interactions.

3.1 Coupling

To achieve coupling, Parunak explains that we first need an *active environment*. The radio propagation medium constitutes just such an environment, as each mobile emitting on a given channel influences the others due to interference. Hence, a mobile's movement changes the interference patterns for all others in its immediate vicinity. Additionally, mobiles are entities which strive to acquire connections and in so doing, they necessarily broadcast information to inform neighboring DBS of their presence and of their link quality. This forms an active environment in which information is exchanged to sustain coupled processes.

We emphasize the fact that DBS and mobiles do form appropriate entities to host agents that are *small in size and scope*. In particular, DBS, compared to central cellular base stations, are specifically meant to be small, and will necessarily have small scope as they can only relay a (smaller) limited number of mobiles in their vicinity.

As a final criterion related to coupling, agents should be mapped as *entities*, not *functions* since an agent does not implement a complete function. That is, the function optimizing resource allocation should be the result of the interaction of the agents and not be implemented as the output of one agent. Indeed, an ant (in ant colonies) does not find a shortest path alone.

In the proposed system, the agents are mapped to either mobiles or DBS. Their actions will then be to either allocate or deallocate a channel, or a connection, or modify a mobile's power level. Necessarily, all processes which modify resource allocation are all coupled since each agent's actions will not only influence the concerned mobile (changing channel, obtaining a new relay connection or changing its power level), but also influence the neighboring mobiles, modifying their own channel footprint, their QoS, hence influencing other agents, and coupling each agent's processes together indirectly.

¹ Since a priori goals are not explicitly defined, neither is the concept of a "right direction". Rather, a behavior is designed, tuned and retained because its auto-catalysis properties happen to converge to a solution which satisfies the needs of the system. Therefore, such a design allows wide exploration of the solution space rather than restricting to predefined goals by not including all the effects involved in the multidimensional problem. The design represents a certain creative process.

3.2 Auto-catalysis

3.2.1 Flows

For agents to maintain their interactions, they must be designed to let the process evolve continuously. Therefore, agents should not be designed based on discrete state transitions, leading to pauses in the processes because of unverified conditions. That is why we must favor *flows instead of transitions*. One way to achieve this is for agents to use volatile markers (i.e. permanent and non-obstructive source of information which dissipate in time as they become irrelevant –e.g. pheromones in ant colonies) to inform other agents on their particular state, so that the agents' processes continuously evolve rather than stop and wait for specific conditions.

It is a design choice that *no explicit information exchange* is performed concerning the positions of mobiles and DBS, available resources, etc. As mentioned, the available information stems from what DBS and mobiles can sense locally (mobiles' needs and QoS), which represents our volatile markers. These bits of information are by nature volatile, as they only stay in the environment as long as they are broadcasted by the mobiles, and hence are necessarily current.

Since agents should not wait for predefined conditions to take actions, it is a *comparative* basis that will trigger a corresponding action of:

1. allocating/deallocating a macrodiversity connection;
2. changing a mobile's channel (frequency hopping);
3. increasing/decreasing a mobile's power level to a certain amount.

3.2.2 Homeostasis

The notion stems directly from biology in which systems always strive to maintain an equilibrium or *homeostasis point*, e.g. the blood sugar concentration is maintained (mainly) by two different hormones which have opposite effects to balance the concentration.

This point of equilibrium must be sustained by an ongoing *flow* to ensure the system continuously explores the solution space and does not get stuck in a deadend. This flow is analogous to the variations of a stock market title whose value is influenced (at a macro level) by the traders' actions of selling and buying. In turn, at the micro level, the variations of the values influences the traders' decisions.

The corresponding aspect of our system is created by forcing DBS to continuously create and destroy connections, continuously change channels (via channel hopping), and continuously adjust power levels. Each of these actions — at the macro level of agents — influences the status of mobiles, and these changes are in turn sensed by surrounding mobiles and DBS.

In effect, the flow of actions makes the system converge to a homeostasis point. This point will be dependent on the the state of the network (traffic, available resources, etc.) due to the comparative basis that triggers actions. As long as there exists a bias observed by the agents that will trigger an action, the system will converge or oscillate to its homeostasis point. These variations are important, since without them, and if there is no other change in the system (e.g., induced by mobile motion), the sensed QoS of mobiles would never change, never trigger actions, and the system might simply stop short of an optimal state.

3.2.3 Amplification and limitation

Together, amplification and limitation constitute an other important aspect to generate the convergence to a homeostasis point. Amplification implies a positive feedback mechanism

such that convergence (to a solution) is favored. In other words, the actions of an agent which lead the system in a desirable global direction should be favored and should also influence the surrounding agents to act in the same direction.

In effect, an MA system is comparable to a Genetic Algorithm (Goldberg, 1989) preserving "genes" that seem to provide the best fitness and hence are part of an optimal solution. The difference is that there are no external observing entities that measure via a metric the fitness of candidate solutions. Rather, it is the interactions between agents and their environment — the propagation medium — that must provide the natural selection function.

Limitation also implies preventing the whole system from focusing on one point (exacerbating the convergence of actions to a local minimum) and thus miss a better solution. Moreover, limitation can favor convergence by dampening the effect of amplification to prevent the system from going past a solution or oscillating around it without converging.

3.3 Function

Coupling may be trivial to obtain and auto-catalysis somewhat more involved, but if the process as a whole does not realize a useful function, then it is irrelevant. Function implies that the homeostasis point described previously is useful for the system, e.g. in biology the homeostasis point for the blood sugar concentration is such that enough sugar is available to fuel the cells, but not too much to avoid excessive sugar loss through the kidneys.

In our system, the sought-after function consists in

- **maximizing the potential usage of the resources;**
- and **balancing them** to offer a good compromise of quality across all mobiles, while not hindering the overall system performance.

Most often, function is obtained through a *utility function* which translates the flow of variations (of QoS) sensed into rational decisions. That is, it converts a multi-dimensional problem into a one-dimensional quantity upon which decisions for actions are based.

In spite of the fact that many frameworks attempt to provide mathematical support to derive such utility functions (such as game theory (Mackenzie & Wicker, 2001) or COIN theory (Tumer & Wolpert, 2004)), these frameworks mostly consider intelligent agents having the ability to learn (eventually using reinforcement learning techniques), which is not the nature of the proposed design. Ultimately, defining simple agent behavior to obtain an intended global behavior still relies on intuition and art such as in Conway's "game of life" (Elwyn R. Berlekamp et al., 1982), or with Wolfram's cellular automata (Wolfram, 2002). Therefore, no systematic procedure is known which derives the locally-applicable utility function from the desired global behavior.

Function can also be sustained (especially if a utility function is not found) with

- *behavior diversity* and
- *randomness*.

Randomness can be helpful to introduce alternative solutions, that will or not be kept in time given how effective they are. Behavior diversity can be obtained by forcing neighboring agents to act differently so as to provide different reactions and experiments given identical stimulus. These properties support the function property by *breaking the symmetry* so as to prevent the system from entering any deterministic patterns which might hinder convergence. In the following section it is described how auto-catalysis and function are obtained for each class of agents.

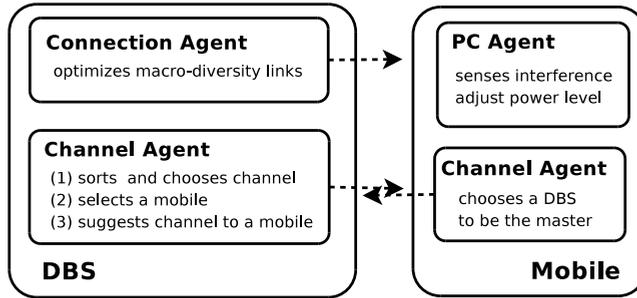


Fig. 2. Agents actions and mapping to entities.

4. Agents' Design

4.1 Agents Sensing Abilities

A mobile broadcasts its actual QoS designated $P_T(m)$, so that DBS in its vicinity can sense its needs. This QoS corresponds to the total BER after macrodiversity combining the mobile's received signals. Moreover, a DBS senses the potential additional link quality it can provide (or already provides) to a mobile. Considering mobile m and DBS B , this additional link quality is named $P_e(m, b)$, and represents the BER of the link from m to DBS b (Leroux et al., 2006). Also, to incorporate the notion of classes of QoS, the DBS knows that mobile m requires an overall quality of $P_d(m)$, i.e. $P_T(m) < P_d(m)$.

For convenience, these values ($P_e(m, b)$, $P_d(m)$ and $P_T(m)$) are expressed in a logarithmic scale of base 10 of the BER. It is shown in (Leroux et al., 2006) that the following holds in a Rice fading environment with different Rice fading parameters ($SINR, K$ factor) values for each link:

$$P_T(m) \approx a + \sum_b P_e(m, b), \tag{1}$$

where the relation (which implies that the combined BER is the sum in the logarithmic domain of the individual link BERs) is exact for certain types of modulation (such as DPSK) and approximate for other types (such as coherent QPSK), and a is a constant related to the modulation.

By definition, if mobile m is not relayed by DBS b , we have

$$P_e(m, b) \triangleq 0. \tag{2}$$

4.2 Connection Allocation

4.2.1 Coupling

Connection agents are mapped to DBS. Indeed, each DBS senses local information on mobiles' needs and can take decisions with regards to the allocation of connections. Mobiles can be relayed by many DBS offering more or less QoS. Therefore, a DBS decision (whether to relay or not a mobile) will influence what its neighboring DBS senses and therefore influence their actions. Hence, the whole network is interdependent and linked or coupled via macrodiversity links.

4.2.2 Flow of action

While attempting to maintain connectivity, DBS should *continuously* change their links to gradually converge to an optimal configuration which maximizes the benefits of macrodiversity. It is this flow of changes in connection allocation which sustains the properties of auto-catalysis and function.

In order to do so, two types of actions are defined : disconnection and reconnection. Each action is taken alternatively, given the number of active connections the DBS has. Suppose a DBS can provide at most N connections and has m active connections, it will perform

1. $N - n$ disconnections, if it has $m > N - n$ connections or,
2. $N - m$ connections, if it has $m \leq N - m$ connections.,

where n is a system parameter (typically $n = 1$) under the designer's control. For large values of N , increasing n helps accelerate the convergence of the system, yet high values of n will suppress the iterative selection mechanism such that the system may not converge any longer. Hence two opposing "forces" must be designed to link the information sensed to the choice that must be taken : which mobiles should be connected or disconnected. Finally, the balance between these two forces should lead connections to a state that represents an homeostasis point.

4.2.3 Function

Randomness is incorporated by having the connection agents activate (to perform a connection or disconnection action) randomly, following a Poisson law considering a discrete model of time. This should also help to prevent any periodic pattern from taking hold.

A *utility function* is designed to link the information sensed to the choice of actions : links are rated according to a continuous function so the DBS can compare the links and decide which to connect or to disconnect.

Two metrics are designed, providing information on:

1. a measure of how well mobile m is served with respect to its requested QoS, i.e.

$$F_{\text{need}}(m) = \frac{P'_T(m)}{P'_d(m)}, \quad (3)$$

2. and a measure of how much diversity DBS b is providing to m with respect to the mobile's overall link quality, i.e,

$$F_{\text{div}}(m, b) = \frac{P_e(m, b)}{P'_T(m)}. \quad (4)$$

P'_T is understood in these definitions as the mobile's total link quality if DBS b is connected to the mobile (whether it is evaluating for connection or disconnection).

These two simple functions provide sufficient information for the DBS to compare the mobiles' links and take a decision based on:

- how much a mobile *needs* more macrodiversity links;
- and to which extent this DBS is the one which will provide the mobile with an efficient macrodiversity link (relative to the other DBS currently serving the mobile).

These information bits still need to be combined, and this is where an appropriate trade-off is induced by using different combination functions.

To design the combination function, different characteristics must be considered:

1. If a mobile has no connection, it must be favored, since basic connectivity should take precedence.
2. If a mobile has only one connection, the DBS should not disconnect it.

Furthermore, there are two complementary compromises involved in the DBS' decision process:

1. either to remove a link because the mobile already enjoys sufficient QoS,
2. or to maintain it because it is the main DBS providing it;

and,

1. either to connect a mobile because it is in need,
2. versus not connecting it because the additional diversity brought to this mobile would be low (compared to other possible connections).

Finally, the function must provide a natural ordering to classify the compromises in order to take a decision.

The following function addresses all the characteristics discussed above:

$$C(m, b) = F_{\text{need}}(m) \times \log(F_{\text{div}}(m, b)). \quad (5)$$

This function is necessarily positive or null. It is null if the mobile has no connection, since, if it were connected to the DBS, it would have $P_T(m) = P_e(m, b)$ which implies $F_{\text{div}} = 1$ giving a null value of the logarithm. Likewise, it is null if considered for disconnection and the mobile's only link is to the considered DBS. Hence, if this utility function is null, the agent will either privilege this mobile for connection or not disconnect it to keep the mobile's existing connection active.

The evaluation of the compromise is obtained by the multiplication of the two terms. Hence, the more the DBS provides diversity, or the higher is the current QoS enjoyed by the mobile, the higher is the function's value.

The choice of compromise itself comes derives from "shaping function" used prior to the multiplication of the two metrics. Simulation showed that the optimization happens most efficiently if the shaping function of the second term is concave (naturally, it should be strictly increasing), hence the use of the logarithm, which also provides the necessary null value for a mobile with a single link.

4.2.4 Limitation and amplification

Limitation and amplification is naturally obtained with the environment propagation properties. Indeed, a poor signal quality will favor multiple connections (amplification), but distance (mobile to DBS) and the infrastructure link capacity of DBS will restrict excessive connection growth (limitation).

Also, this amplification (or attraction of macrodiversity links) and limitation sustains the homeostatic behavior where mobiles in need get more links up to an equilibrium point where additional links to these mobiles would overwhelmingly affect an otherwise well-served mobile.

4.3 Channel allocation

4.3.1 Flow

The flow of actions in the channel allocation agents naturally consists of the changes in channel allocation, or *channel hopping* which modifies mobiles' QoS and interference patterns which in turn should trigger other changes.

For this flow to be generated properly, appropriate actions are specified in the following.

4.3.2 Coupling

Following Parunak's principles, the sought-after function (optimizing the allocation) is divided into independent actions whose interactions should lead to the other two properties (auto-catalysis and function).

First, given the macrodiversity context, a mobile will choose one of its relaying DBS to be its "master" connection, which implies one type of action and one agent (to select the master) mapped at each mobile.

Second, DBS will choose mobiles (from their master links) and change their channels as is done in cellular systems. Except that here, the change, or channel hopping, will not be triggered by specified conditions (e.g. a mobile SINR falling below a threshold, or a mobile changing cell). Instead, the flow of channel hopping will be sustained by having DBS choose a mobile at each agent activation and change its channel. Channel allocation agents will activate in the same way that the connection agents do. Two types of actions must be defined:

1. choosing a mobile, and
2. choosing a channel.

Mapping these actions at the DBS level, rather than letting the mobile decide when to change channel makes sense in that DBS can gather information most effectively on the different channels in use, thus preventing mobiles from having to continuously scan channels.

4.3.2.1 Sensing

In addition to the mobile's sensed link quality, DBS can sense

1. the received power of surrounding mobiles $p_r(m, b)$;
2. and the interference level on various channels $p_I(b, c)$ (for channel c at DBS b).

4.3.3 Function

Maximizing the channel usage constitutes, in a sense, an effort against the second law of thermodynamics. Indeed, the channel allocation, if optimal at some point in time, will necessarily deteriorate with mobility as two mobiles transmitting on the same channel get closer to a point where the interference will degrade the offered QoS, such that resources are not balanced anymore. Considering this aspect, and rather than trying to solve an NP-complete problem, load balancing is obtained by always attempting to change the channels of mobiles in need such that they enjoy better SINR.

Three utility functions need to be designed taking as input what the agents can sense, and yielding a chosen parameter value as output.

- a. Mobile m will choose a master DBS (among its relaying DBS) on activation (where its activation follows a Poisson law) based on the DBS from which it obtains the highest link quality:

$$b = \arg \max_b \{P_e(b, m)\}. \quad (6)$$

- b. DBS b will choose a mobile (among mobiles connected as master to b), that is the most in need, i.e.

$$m = \arg \min_m \{F_{\text{need}}(m)\}. \quad (7)$$

- c. Ideally, the DBS should try to use the channel with the lowest interference power level:

$$c = \arg \min_c \{p_I(b, c)\}. \quad (8)$$

However, it may be overwhelming for a DBS to systematically sense channels to maintain up-to-date information on interference levels on all channels, and this behavior (utility function c.) is therefore only used as a benchmark.

Akaiwa & Andoh (1993) suggested a selection mechanism which is used herein with some modifications. DBS b will scan channels in the order of a given priority list it maintains, and determine if a channel can be assigned according to

- whether the resulting SINR will be above an SINR threshold ;
- and (in addition to Akaiwa's method) whether it will also be above the actual SINR the mobile enjoys.

The SINR threshold represents a mean to control the hidden terminal effect. It is a studied parameter in order to observe to which extent it prevents HTE while not limiting the flexibility of the system.

For Akaiwa's segregation algorithm, the priority list is obtained dynamically given the ratio for each channel of previous assignments versus previous assignment attempts.

Finally, a random priority list is proposed as a simple, yet effective (as we will see) alternative to the segregation algorithm approach.

4.3.4 Limitation

DBS will only test a limited number of channels given by the Ch_{max} parameter, before giving up. Indeed, there is no guarantee that the DBS will find a channel that will suit the chosen mobile. Therefore, and instead of letting it scan all channels, it is forced to limit its search. Eventually, it will try again, or another DBS will, thus providing behavior diversity as well. In effect, the DBS are only trying to maintain channel assignments in a working state by "upgrading" the solution iteratively in an opportunistic fashion given the eventual availability of channels. It is the effect of a new channel allocation that will cause other DBS to also react and change channels for the mobiles that will see their QoS affected by the new neighboring interference. As this flow of action is sustained, the channel allocation remains functional and should adapt to changes.

4.3.5 Channel availability

An additional functionality is provided for channel availability. A few spare channels are reserved for the initial connection or reconnection of stranded mobiles (instead of using channels from the main pool). Then, a master DBS which has mobiles on these spare channels will attempt to change their channels as a priority instead of choosing another mobile. Such spare channels allow rapid network entry, providing higher availability as well as some time margin for the DBS to find free channels in the main pool. It therefore eases the process and the flow of channel hopping.

4.3.6 Alternatives

Some attempts were made in order to add additional functionality to better handle irregular topologies and classes of QoS. One such attempt was not only to change channels of mobiles in need to better ones, but to also change channels of over-served mobiles with worse ones. However, this approach proved fruitless. One reason is that it is not possible for DBS to differentiate between channels having high interference due to over-served mobiles or poorly served mobiles. And the change to a “worse” channel can have more cons (far too much degradation for other mobiles) than pros (compacting channels more efficiently to free resources). Indeed, the main challenge resides in the degradation of interference. And, in effect, better results are obtained by simple channel hopping with the poorly served mobiles until other mobiles are affected, react, and some balance is obtained.

The concept of classes of channels was also explored, where the threshold considered to assign a channel would be modulated given the class of the channel and the mobile’s need, in order to generate classes of channels with less interference and some with more, for mobiles which can sustain it (which is also inspired from the umbrella-cell mechanism). No significant increase of performance has yet been obtained with this approach.

4.4 Power Allocation

4.4.1 Flow

Necessarily, the flow in power allocation is the result of changes in power level, which in turn changes interference, and the QoS of surrounding mobiles, which in turn should trigger other power level changes. This striving for amplified/limited adjustments of power levels aims to converge to a point of equilibrium : an homeostasis point where the system actually exhibits its expected power control behavior.

4.4.2 Function

The previously defined $F_{\text{need}}(m)$ definition is considered. If it is above 1, the mobile is *over served* (it enjoys a BER better than requested) and should reduce its power level, or increase it if lower than 1. Yet, this in itself would be restrictive, as it would force all mobiles to the same mean quality $\overline{F_{\text{need}}}$ value, without taking into account the non linearity of the problem due in part to the limited dynamic range of mobiles’ power levels. Indeed, and as will be shown, the thermal noise at the receivers imposes a limit on the minimum power level a mobile can transmit without the resulting SINR being too low for the connection (however small the interference level is). And, of course, mobiles saturate at a maximal power level. Also, some mobiles might be able to obtain more quality while not restricting others to maximize their own, and if so, they should.

The Need variable is introduced to describe what could be the power level of the mobile given its $F_{\text{need}}(m)$ factor : it is the value to which the mobile’s power level should converge to if nothing else changes (which is not the case as other mobiles will adjust their power level). When, for all mobiles, the current Need_m value equals the current power level $p(m)$, then the homeostasis point is reached. The variable is defined

$$\text{Need}_m(F_{\text{need}}(m)) = 2e^{(-SF_{\text{need}}(m))} \times (SF_{\text{need}}(m) + 1) - 1, \quad (9)$$

where the value S is a scaling parameter for $F_{\text{need}}(m)$ in order to allow mobiles to potentially obtain QoS higher than their requested $P_d(m)$. Therefore, the Need function will not necessarily be smaller than the current power level if $F_{\text{need}}(m) > 1$. And mobiles will not be forced to restrict their obtained QoS to a global mean value. In the current simulations, this QoS is

maximized with $S = 0.8$, and this has been shown to hold in many different conditions of traffic, mobile speeds in (9) and available resources. .

The exponential in (9) is a shaping function which also naturally affects the dynamics of the system. In effect, it affects mobiles' convergence speed differently given their needs, and this translates into behavior diversity as no mobile will react in a precisely proportional manner. The proposed function is of course not the only possible choice, but it has proved stable and effective. Again, for MAS, effectiveness does not lie in the mathematical exactness of the function, but in the interactions it will generate.

4.4.3 Homeostasis

Finally, this Need factor must be converted to a delta (step) value to adjust the power level. *Homeostasis* is obtained by comparing the Need value to the current power level the mobile uses to transmit. Hence, the delta value is in the form of $\text{Need}_m - p_m$. The mobile will then try to converge to a Need_m value which depends on local interactions given itself and neighboring mobiles' F_{need} values (given that these are indirectly linked via interference). Eventually, a non-linear concave function helps convergence so that with Need_m and p_m close, the generated delta is kept small to slow down variations and help stabilize the convergence. We postulate

$$\Delta_m = \beta \text{sign}(\text{Need}_m - p_m) (|\text{Need}_m - p_m|)^{1.5}, \quad (10)$$

where the β factor is used to modify the dynamics of the system to attain the proper compromise between convergence speed and stability.

4.4.4 Limitation

Experience shows that this function is too unstable with high values of β . Still, it can be stabilized with additional scaling parameters, while maintaining fast adaptation in time with large values of $\beta \geq 5$, which is important for mobility ($\beta = 5$ is used in the presented simulations). Therefore, it is proposed that Δ be scaled according to the current power level and also the desired power level (the Need value). That way, if these values are small, Δ is also kept small to prevent strong changes in the system that would otherwise suddenly generate exaggerated interference. Indeed, such changes would lead to complications such as breaking existing links or simply propagating exaggerated reactions throughout the system. Building upon (10), the following function is used :

$$\Delta_m = p_m \times |\text{Need}_m| \times \beta \text{sign}(\text{Need}_m - p_m) (|\text{Need}_m - p_m|)^{1.5}. \quad (11)$$

Finally, the delta value is constrained to not exceed the power level range:

$$\Delta_m < 0 \Rightarrow \Delta'_m = \max\left\{\Delta_m, \frac{-p_m}{2}\right\} \quad (12)$$

$$\Delta_m > 0 \Rightarrow \Delta'_m = \min\left\{\Delta_m, \frac{1}{2}(1 - p_m)\right\}. \quad (13)$$

As the mobile's PC agent activates, its power level is adjusted as follows:

$$p_m^{(v+1)} = p_m^{(v)} + \Delta'_m. \quad (14)$$

5. Evaluation

5.1 Simulation platform

For the channel and power agents, simulations are based on the following platform. The results shown for the connection agents are based on a simpler scenario (detailed in the appropriate subsection), in order to isolate the effect of connection management and observe its convergence, while not confusing it with the effect of interference and power-level management.

Physical parameters

A square field of 25 square kilometers is considered, in which 1000 mobiles evolve and 100 DBS are scattered randomly. Hence, the traffic's and network resources' geometry are not uniform, thus generating good and bad coverage of different areas. A mobile moves in a random direction at a random speed taken (at the start of a scenario) out of a uniform distribution over $[0, V_{\max}]$. DBS can relay 25 mobiles each, such that the mean number of macrodiversity links per mobile is 2.5. A mobile's maximum transmit power is 1W at 1 meter of its antenna, and the propagation exponent is 4 ($g_{ij} \sim 1/d^{-4}$). Rayleigh fading is considered, except near a DBS (closer than 100m) where a line of sight component is added with Rice factor $K = 5$ dB. Thermal noise at the receiver is considered for a bandwidth of 30kHz at a temperature of 20°C , hence $N_0 = -129$ dBW. The number of available channels is denoted Ch.

Agents' emulation

Simulations are run for 1000 seconds and repeated 10 times with different initializations of the geometry (DBS positions and mobiles' initial position, directions and speeds). Time is discretized with a time step of 1 second. At each time step, physical parameters are evaluated (mobile's position, propagation, interference, BER, connection outage). Agents activate randomly given a Poisson distribution to estimate the next activation time with parameter $\lambda = 3$ time steps. At each time step, the agents which activate evaluate their local state and take actions accordingly (adjust the power level, hop to a new channel, change connections of the concerned mobile, etc.).

Results

At each time step, the set of QoS indexes (total BER level given on a logarithmic scale $P_T(m) = \log_{10}(BER(m))$) for each mobile are sorted, thus providing a snapshot in time of the distribution of the network's resources across all mobiles. These sorted distributions are then averaged for all the time steps of the simulation. Given this information, it is then possible to compare how each algorithm distributes resources. The same is done for the power level allocation. Also, to verify the stability in time (considering the dynamic properties) of the algorithm, two factors 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.

The latter also provides insight on how well the system is able to provide resources to mobiles with high availability.

Thank You for previewing this eBook

You can read the full version of this eBook in different formats:

- HTML (Free /Available to everyone)
- PDF / TXT (Available to V.I.P. members. Free Standard members can access up to 5 PDF/TXT eBooks per month each month)
- Epub & Mobipocket (Exclusive to V.I.P. members)

To download this full book, simply select the format you desire below

