Cellular Channel Assignment:
a New Localized and Distributed Algorithm

Roberto Battiti[*][†]   Alan A. Bertossi[*][†]   Mauro Brunato[*][†]

Abstract

As the use of mobile communications systems grows, the need arises for new and more efficient *channel allocation* techniques. The total number of available channels on a real-world network is in fact a scarce resource, and many assignment heuristics suffer from a clear lack of flexibility (FCA), or from high computational and communication complexity (BCO, BDCL). Performance can be improved by representing the system with an objective function whose minimum is associated to a good configuration; the various constraints appear as penalty terms in the function. The problem is thus reduced to the search for a minimum, that is often performed via heuristic algorithms like Hopfield neural networks, simulated annealing, reinforcement learning. These strategies usually require a central process to have global information and decide for all cells.

We consider a penalty-function formulation of the channel assignment problem that has been previously solved by search heuristics; we prove that the global minimum of the objective function can be found with $O(n \log n)$ time complexity, and therefore there is no need of heuristics. Finally, we show that the algorithm that arises from this formulation can be modified so that global knowledge and synchronization are no more required, and we give its distributed version. By simulating a cellular network with mobile hosts on a hexagonal cell pattern with uniform call distribution, we show that our technique actually performs better than the best known algorithms.

1   Introduction

Consider a geographic area in which a number of multi-channel transceiver server stations are placed at a suitable mutual distance. Stations are connected to each other by means of a wired network, which we are not concerned about in the present paper. Each station acts as a network entry point for all the mobile radio hosts that can communicate with it at the highest signal-to-noise ratio among all stations. As a result, the geographic area is divided into *cells* whose borders are not well defined, as they depend on the ever-changing radio signal levels. A rough sketch is given in Fig. 1, which still depicts a *good* situation where every cell is a simply connected domain: reflections, interference and terrain configurations, as well as handoff policies, might substantially alter it. However, a more regular setting such as the common hexagonal pattern will be considered in the experiments, because it is still a commonly used test ground. Modeling real-world instances is in fact a hard and delicate work, and the situation is made worse by the policies of many cellular service providers who keep their data reserved, or provide them in a strictly confidential way, so that the scientific community cannot share them and permit reproducibility of results.

In the typical situation, a server station can be received by server stations in other cells. In this case, mutually interfering stations must employ different communication channels (i. e. frequency bands, time slices or codes from an orthogonal set), in order to avoid *co-channel interference* (interference caused by transmissions on the *same* channel). In its simplest (and most unrealistic) form, the channel

---

[*]Università di Trento, dipartimento di Matematica, via Sommarive 14, 38050 Povo (TN) — ITALY.

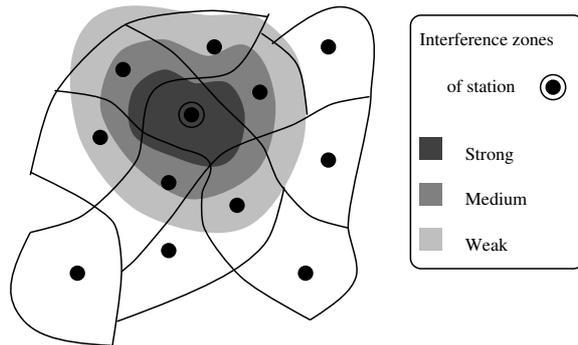[†]E-mail: `battiti/bertossi/brunato@science.unitn.it`

Figure 1: A simple geographic cellular network and the interference area of a server

assignment problem is equivalent to the Euclidean graph coloring problem, hence it is NP-hard. This problem can be treated with simple greedy heuristics [8] [3] [1]. Because a server station must communicate with several mobile hosts at once, however, we must assign more than one channel to each server. When this problem is handled with graph-coloring heuristics, we need to substitute every node with a clique of cardinality equal to the required number of channels, in order to give the appropriate number of channels to each transceiver while respecting the interference constraints. This approach causes a quadratic explosion of the problem.

Moreover, radio interference is *additive*, and simple adjacency restrictions (like those in the graph or list coloring problems) are not sufficient to model the complexity of the real-world situations. If the interference phenomena are strong enough, even stations that use different channels may interfere, provided that they operate on adjacent frequency bands or on subsequent time slices (propagation delays may cause a time slice to partially invade another one). This problem becomes significant when the overall frequency spectrum has to be minimized. Indeed, the strong need of radio bands for several purposes makes the reserved bandwidth for cellular communications rather small ($\Delta f \approx 60\text{MHz}$ in the 900MHz band for the GSM system [2]), and hardware techniques[1] can't do all the job by themselves.

The remaining part of this paper is organized as follows. In Section 2 we give a fast overview of the algorithms that can be found in the literature (Section 2.1), then we describe in greater detail the objective-function approach which is central to this paper (Section 2.2) and we show how the minimization problem can be solved in polynomial time (Section 2.3).

In Section 3 we introduce some modifications of the algorithm in order to distribute it among the cells. For this purpose, we need a communication scheme to broadcast all status changes to a small local cluster (Section 3.1) and a local synchronization method to make sure that a cell allocates a channel only when it has enough knowledge of the status of its neighbors (Section 3.2).

The results of simulations are presented in Section 4, where the new algorithm introduced is compared with those described in Section 2.1. In section 5 we draw some conclusions suggested by the experimental analysis.

2   Channel Assignment Algorithms

In this section we summarize the most successful existing approaches for the problem and present a polynomial-time solution of a penalty-function model studied in [7].

2.1   Combinatorial strategies

Let us briefly summarize the main channel assignment algorithms (for more details see the reviews in [6] and [5]). The same algorithms will be considered for the experimental comparisons of the newly

---

[1]Namely, Minimum Shift Keying and Gaussian Minimum Shift Keying to eliminate spectrum bumps around the channel, and frequency hopping schemes to eliminate constant interferences.

proposed technique, see Section 4.

- The FCA (Fixed Channel Allocation) algorithm. Each cell is assigned a fixed pool of frequencies, so that no near cells can use the same channel. No communication is needed between cells; when all channels are in use, subsequent requests shall be rejected.

- The SBR (Simple Borrow from the Richest) algorithm. Each cell has an assigned pool of frequencies, but a channel can be borrowed from a richer neighbor, provided that its use does not cause interference. When more than one neighbor has a free channel, the cell chooses the richest one. It is more efficient than FCA with low traffic rates, but its performances deteriorate when the traffic rate increases, achieving the same performance of FCA at higher computational and communication costs.

- A DCA (Dynamic Channel Assignment) technique, by which every cell can have access to every channel, as long as it does not cause interference; the cell chooses the channel which is most 'blocked' (due to the interference constraints) in the neighboring cells, so that it gives rise to the least blocking probability. It is better than FCA at low traffic rates, but worse at high traffic rates, because many cells might find no channels at all, due to non-optimal decisions at previous times.

- The BDCL (Borrow with Directional Channel Locking) algorithm. Like SBR, but the choice of the channel to borrow is done by the criterion of the above proposed DCA technique. It is the best combinatorial algorithm we know.

In addition to the above algorithms, we consider in more detail a technique based on objective function minimization, where the constraints of the problem become penalty terms in the function.

2.2   A penalty function heuristic

Many penalty-function heuristics have been applied to the channel assignment problem. Some of them require the rearrangement of the whole cellular system when a new channel is requested [9] [4]. Unless the traffic is very low and service communications among cells are cheap and fast, this approach is of little practical use. Other penalty-function heuristics, such as the one we shall consider next, just rearrange the channel assignment inside the cell where the new communication request is issued, by choosing those channels that minimize the probability of a future channel refusal, trying to keep the system in a suitable status for future requests. This effort of maintaining a good configuration is clearly unnecessary in the former case.

In both cases, two heuristic steps of different nature are usually taken. First, one must determine a good penalty function, where "good" means that its minimum should correspond to a suitable configuration. Then, when this function is determined, one must actually determine its global minimum, or at least find some good approximation thereof. For this purpose, a heuristic technique can be used to search for a local minimum of the penalty function. In this second case the term "heuristic" means that there is no proof of convergence to the global optimum in polynomial time.

Following [7], let n_{CE} be the number of cells and n_{CH} the total number of channels. Every cell i, $i = 1, \ldots, n_{CE}$, has a traffic demand $traf_i$ which changes with time. Let us denote with $d_{ii'}$ the Euclidean distance between the centers of cells i and i', and let $interf_{ii'}$ be a $\{0, 1\}$-valued function which states if the two cells interfere or not; the notation can nonetheless be extended to the case of various degrees of interference.

When a connection or termination request is issued in cell i^*, the frequency allocation in this cell must be optimized. The status of channel allocation is given by a $\{0, 1\}$-valued matrix A_{ij} whose entry (i, j) is 1 if and only if channel j is currently in use in cell i. The new channel allocation for cell i^* is stored in vector V_j, $j = 1, \ldots, n_{CH}$.

An objective function is built whose minimum is likely to be a good solution of the new allocation for cell i^*. First, a term to privilege those solutions without interference (all terms depend on V, our
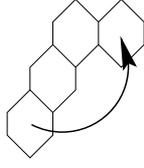
3

Figure 2: Building a reuse scheme: the basic move

unknown solution) is introduced:

$$a(V) = \sum_{j=1}^{n_{\mathrm{CH}}} \sum_{\substack{i=1 \\ i \neq i^*}}^{n_{\mathrm{CE}}} V_j A_{ij} \, \mathrm{interf}_{ii^*} \, .$$

This term adds 1 for each cell interfering with i^* which uses a channel in use in i^*. Second, the requests of the cell i^* should be respected as much as possible:

$$b(V) = \left(\mathrm{traf}_{i^*} - \sum_{j=1}^{n_{\mathrm{CH}}} V_j \right)^2 .$$

The only reason to make this term quadratic is that it must be nonnegative: (an absolute value would also work). Third, a *packing condition* is added: a channel should be reused as near as possible (outside the interference zone), to restrict the blocking probability in other cells.

$$c(V) = - \sum_{j=1}^{n_{\mathrm{CH}}} \sum_{\substack{i=1 \\ i \neq i^*}}^{n_{\mathrm{CE}}} V_j A_{ij} \frac{1 - \mathrm{interf}_{ii^*}}{d_{ii^*}}.$$

This subtracts a positive term for each cell outside the interference zone which reuses a channel employed in cell i^*; the larger the distance, the smaller the subtracted term. Next, changes in the present allocation of the cell i^* should be minimized:

$$d(V) = - \sum_{j=1}^{n_{\mathrm{CH}}} V_j A_{i^* j}.$$

This subtracts 1 every time a channel currently used by cell i^* is chosen for the next configuration (this means that a mobile host needs to change its channel as rarely as possible). If some frequency hopping technique is used, however, this requirement does not make much sense, as the mobile host is equipped for frequent configuration changes. Last, experimental evidence shows that to achieve a good performance the channel reuse should follow a regular scheme (for example, a compact pattern [10]). This is achieved by introducing the $\{0, 1\}$-valued matrix $\mathrm{res}_{ii'}$ whose entry (i, i') is 1 if and only if cells i and i' belong to the same reuse scheme (i.e. should use the same channels if possible). Common reuse schemes follow some sort of "knight" move (for instance, the one shown in Fig. 2).

$$e(V) = \sum_{j=1}^{n_{\mathrm{CH}}} \sum_{\substack{i=1 \\ i \neq i^*}}^{n_{\mathrm{CE}}} V_j A_{ij} (1 - \mathrm{res}_{ii^*}).$$

Note that all terms are arranged to go towards a lower value when the constraints are satisfied. Let us combine them in a single objective function to minimize:

$$J(V) = A \cdot a(V) + B \cdot b(V) + C \cdot c(V) + D \cdot d(V) + E \cdot e(V),$$

where A, B, C, D and E give different importance to the various constraints.

2.3 The polynomial algorithm BBB

So far a possible penalty function $J(V)$ has been heuristically determined. To minimize $J(V)$, [7] employs Hopfield neural networks, but actually the minimization of this function is straightforward and does not require any heuristic search technique. In fact, we can rewrite $J(V)$ as a quadratic function in which the quadratic term *depends only on the number of channels*, and not on the single channels used. Let us rewrite

$$a(V) = \sum_{j=1}^{n_{\mathrm{CH}}} V_j \, a_j, \qquad \text{where} \qquad a_j = \sum_{\substack{i=1 \\ i \neq i^*}}^{n_{\mathrm{CE}}} A_{ij} \, \mathrm{interf}_{ii^*};$$

the term a_j simply counts the number of cells in the interference zone of i^* which use the channel j. The term $b(V)$ can be rewritten as

$$b(V) = \left(\sum_{j=1}^{n_{\mathrm{CH}}} V_j \right)^2 - 2 \, \mathrm{traf}_{i^*} \sum_{j=1}^{n_{\mathrm{CH}}} V_j \, .$$

The $\mathrm{traf}_{i^*}^2$ term is constant and can be omitted, while the quadratic term is the square of the number of channels reserved for the cell i^* (the number of 1's in vector V). Let us rewrite it in a way similar to the other ones:

$$b(V) = \left(\sum_{j=1}^{n_{\mathrm{CH}}} V_j \right)^2 + \sum_{j=1}^{n_{\mathrm{CH}}} V_j b, \qquad \text{where} \qquad b = -2 \, \mathrm{traf}_{i^*} \, .$$

Clearly, b does not depend on j. The other terms can be rewritten as follows:

$$c(V) = \sum_{j=1}^{n_{\mathrm{CH}}} V_j c_j, \quad d(V) = \sum_{j=1}^{n_{\mathrm{CH}}} V_j d_j, \quad e(V) = \sum_{j=1}^{n_{\mathrm{CH}}} V_j e_j$$

where

$$c_j = - \sum_{\substack{i=1 \\ i \neq i^*}}^{n_{\mathrm{CE}}} A_{ij} \frac{1 - \mathrm{interf}_{ii^*}}{d_{ii^*}},$$

$$d_j = -A_{i^*j}, \qquad e_j = \sum_{\substack{i=1 \\ i \neq i^*}}^{n_{\mathrm{CE}}} A_{ij} (1 - \mathrm{res}_{ii^*}).$$

The term c_j evaluates the packing condition for channel j; the term d_j rewards the choice of channel j if it was already in use; the term e_j penalizes the use of a channel outside the reuse scheme.

We can collect the single coefficients into global ones:

$$w_j = A \cdot a_j + B \cdot b + C \cdot c_j + D \cdot d_j + E \cdot e_j;$$

the global objective function is then

$$J(V) = \left(\sum_{j=1}^{n_{\mathrm{CH}}} V_j \right)^2 + \sum_{j=1}^{n_{\mathrm{CH}}} w_j V_j,$$

where, as we have already pointed out, the square term is just the square of the number of assigned channels.

To minimize $J(V)$ we calculate the weights w_j for each channel; each calculation requires at most n_{CE} steps to test interferences, reuses and packing. Globally, the calculation of the weights w_j requires

time $O(n_{CE}n_{CH})$. If we had fixed the number n of channels that we want to assign, the minimization would be achieved by taking the channels j whose w_j are the least (the quadratic term is constant among the solutions with the same number of channels). To take advantage of this, we calculate a permutation σ_j, $j = 1, \ldots, n_{CH}$, such that the vector $(w_{\sigma_j})_{j=1,\ldots,n_{CH}}$ is sorted in increasing order. The sort requires time $O(n_{CH} \log n_{CH})$. At last, let us call J_n the minimum of the objective function restricted to n-channel solutions. Its value is

$$J_n = n^2 + \sum_{j=1}^{n} w_{\sigma_j}, \qquad n = 0, \ldots, n_{CH},$$

and the difference between the minima for successive values of n is

$$J_n - J_{n-1} = 2n - 1 + w_{\sigma_n}, \qquad n = 1, \ldots, n_{CH}.$$

So, a simple scan of the vector w_{σ_j} is enough to find the minimum for all n, that is the global minimum, in time $O(n_{CH})$.

Hence, the global minimum of the objective function,

$$\min_{V \in \{0,1\}^{n_{CH}}} J(V) = \min_{n=0,\ldots,n_{CH}} J_n,$$

can be found in total time $O\big(n_{CH}(n_{CE} + \log n_{CH})\big)$. The procedure returns also the number of channels in the optimal solution, say n^*, therefore the channels to be assigned to cell i^* are

$$n_{\sigma_1}, n_{\sigma_2}, \ldots, n_{\sigma_{n^*}}.$$

Let us denote this penalty-function minimization algorithm as BBB (where the initials of the authors are used in the lack of a better term). The immediate advantage of BBB is that one finds the true global optimum of $J(V)$ and not only an approximation without performance guarantees. In addition, let us consider that Hopfield networks, like many other techniques, require the function coefficients to vary only in a certain range in order to ensure stability and convergence of the search process. In other words, the precise values of the coefficients are critical not only for weighting the various constraints (which is precisely what they are introduced for), but also for making the heuristic search procedure succeed. Algorithm BBB does not use any heuristic search algorithm, therefore it is not restricted to those coefficient values that ensure convergence, and one can let them vary over all the nonnegative real range, thus having more freedom in tuning them.

3   The distributed polynomial algorithm dBBB

We first note that algorithm BBB can be improved by storing at each cell a permanently sorted array of weights to be updated at each change of state in the nearby cells. The sorting time can thus be cut down to a simple update of the sorted array at each call. Moreover, by considering only the local interference, the calculation of the weights does not depend on the total number n_{CE} of cells, but it can be performed in $O(1)$ time.

We need, however, to simplify our objective function by eliminating non-locality. There are only two global terms in the function:

- The "reuse scheme" given by the array $\text{res}_{ii'}$. The experimental tests (Fig. 8, Section 4) show that it is not influential on the overall system performance.

- The "packing condition", whose weight decreases at increasing distances. The same experimental tests cited above prove that we can restrict the "packing condition" to the $2r$-th ring of neighboring cells (where r is the interference radius).
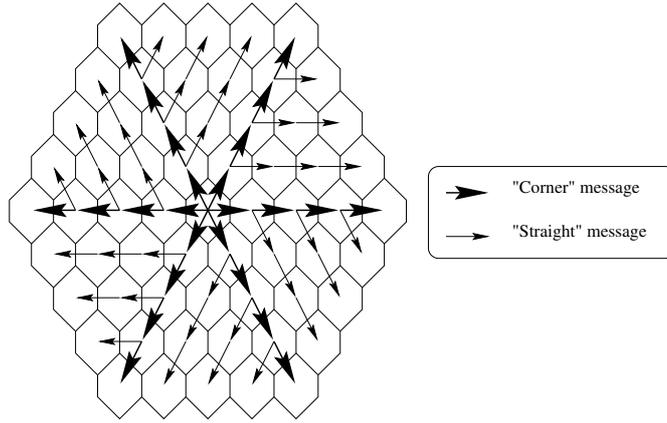
Figure 3: A local broadcasting scheme

```
1.  source:
2.    for d in 0 ... 5 do
3.        send (1, 4, 1, m) along direction d;
4.  relay:
5.    upon receipt of (c, r, h, m) from direction d do
6.        if h < r then
7.            send (c, r, h + 1, m) along direction d + 3   (mod 6);
8.            if c then
9.                send (0, r, h + 1, m) along direction d + 4   (mod 6);
10.       Act according to the received message (c, r, h, m);
```

Figure 4: The local broadcasting algorithm

3.1 A communication scheme

Once the objective function is localized, we just need a good communication strategy to replace the central authority which took all the decisions in the previous algorithm. When a cell initiates or terminates a call, it must broadcast its new status to its neighbors (up to the $2r$-th ring). To do so, we need a simple broadcasting scheme, like the one presented in Fig. 3.

Let the message be structured as a tuple (c, r, h, m), where c is a flag indicating if the message must be duplicated: if c is set, the message shall be called a "corner" message; r is the maximum distance from the source the message must reach, h is the number of steps the message has taken up to now and m is the message itself. Suppose that in each cell the directions of incoming and outgoing messages are numbered clockwise modulo 6. Then if a message arrives on direction d, the opposite direction will be $d + 3$ (mod 6).

The algorithm can be implemented as in Fig. 4. The source sends six "corner messages" to its neighbors (lines 2–3); when a cell receives a message, if it is not far enough (line 6), it must relay it to the cell on the opposite side (line 7); if it is a corner message, the cell must also propagate a non-corner ($c = 0$) copy of the message to the clockwise-next direction (lines 8–9). This copy shall be subsequently propagated only in straight line. After the propagation of the message, of course, the relay cell must modify its record about the broadcasting cell according to the message content m (line 10).

3.2 A Mutual Exclusion technique

Last, to avoid conflicts in channel choice we must ensure that, when a cell is changing its configuration, none of its neighbors up to the reuse distance does the same thing simultaneously. For this we can
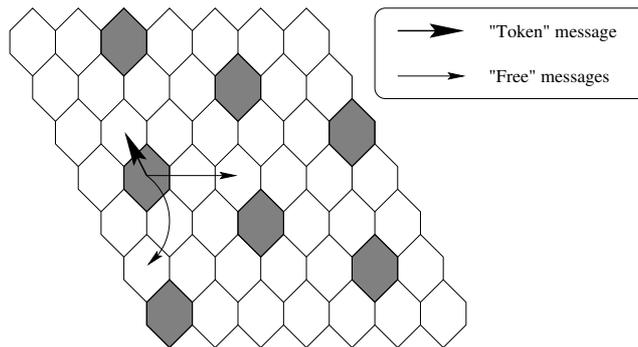
Figure 5: A multiple token-passing scheme

implement a multiple token-passing protocol such that no two tokens are nearer than the reuse distance. If the reuse distance is two and the network is a regular hexagonal grid, let us refer to Fig. 5. The grey cells are possessing the token; when a cell is done with it, it sends a "token" message "upwards" (following the thick arrow) and two "free" messages along the thin arrows (this requires two other cells to act as relays). Before entering the critical state, a cell must wait for one "token" and two "free" messages. The "token" message ensures that all preceding cells in the token-passing chain are safe, while the two "free" messages declare the safety of the two potentially blocking cells which are possibly using the token at the same time in its neighborhood. To take the border effects into account, however, the two leftmost columns and the two uppermost rows should wait for just one "free" message, the cells in the upper left corner don't have to wait for any "free" message, while the lower row cells should generate a token (without waiting for one) whenever they get enough "free" messages.

The actual algorithm for a 7×7 grid with a reuse distance of 2 is presented in Fig. 6 (the directions are numbered clockwise from 0 to 5 starting from west). Three variables are used to store the status of the cell. The boolean *Token_flag* is **true** if and only if the cell is holding the token; the counter *Free_count* contains the number of "free" messages that still have to be received before the cell is able to use the token; the flag *Critical_section_flag* is **true** if and only if a channel request or release has been issued in the cell. The entry to the critical section is only allowed when *Token_flag* is **true** and *Free_count* is 0.

All cells must initially call the *Init* procedure to initialize their status, except those initially possessing the token (the grey ones in Fig. 5). The cells that initially possess the token must start calling the procedure *Have_token_at_the_beginning*. Cells that need to process a channel request (or release) must call the procedure *Enter_critical_section* that sets the *Critical_section_flag*, so that the algorithm runs the channel-assignment procedure when possible, and waits until that same bit is reset.

Of course the main loop is executed concurrently, and the *Check* procedure must end before any other message is received. For routing purposes, since they have to travel two cells, the "Free" messages have an integer part. When a cell receives a "free(0)" message, it just has to retransmit it as "free(1)" to a direction that depends on the incoming path (lines 30–33), while a "free(1)" message must be operated on place (lines 34–36) by procedure *Check*, that verifies if the cell has to enter the critical section, skip it or wait because it isn't ready.

4   Experimental analysis

We consider a 7×7 hexagonal grid, as the one shown in Fig. 5, like in most of the literature. A fixed server station is placed at the center of each cell, while a number of mobile hosts is free to move across the whole land. The total number of available channels is 70 and the co-channel interference is extended to the second ring of neighbors. The grid does not wrap around like a torus.

A C++ program has been written to run a comparative simulation of five algorithms (FCA, SBR, DCA with local optimization, BDCL and BBB).

8

```
 1.   Procedure Have_token_at_the_beginning:
 2.   ┌  Token_flag ← true;
 3.   │  Free_count ← 0;
 4.   └  do Check

 5.   Procedure Init:
 6.   ┌  if (node is in the first row) then
 7.   │      Token_flag ← true
 8.   │  else
 9.   │      Token_flag ← false;
10.   │  Free_count ← 2;
11.   │  if node is in the two upper rows then
12.   │      decrease Free_count;
13.   │  if node is in the two leftmost columns then
14.   │      decrease Free_count;
15.   └  Critical_section_flag ← false;

16.   Procedure Enter_critical_section:
17.   ┌  Critical_section_flag ← true;
18.   └  wait until Critical_section_flag = false;

19.   procedure Check:
20.   ┌  if Token_flag and Free_count = 0 then
21.   │  ┌  if Critical_section_flag then
22.   │  │      Run the channel assignment procedure;
23.   │  │  Send Token along direction 1;
24.   │  │  Send Free(0) along directions 3 and 4;
25.   └  └  do Init

26.   Main polling loop:
27.   ┌  Upon receipt of Token do
28.   │  ┌  Token_flag ← true;
29.   │  └  do Check
30.   │  Upon receipt of Free(0) from direction 1 do
31.   │      send Free(1) along direction 5;
32.   │  Upon receipt of Free(0) from direction 0 do
33.   │      send Free(1) along direction 3;
34.   │  Upon receipt of Free(1) do
35.   │  ┌  decrease Free_count;
36.   └  └  do Check
```

Figure 6: The token-passing algorithm

9

Table 1: Coefficients for function $J(V)$

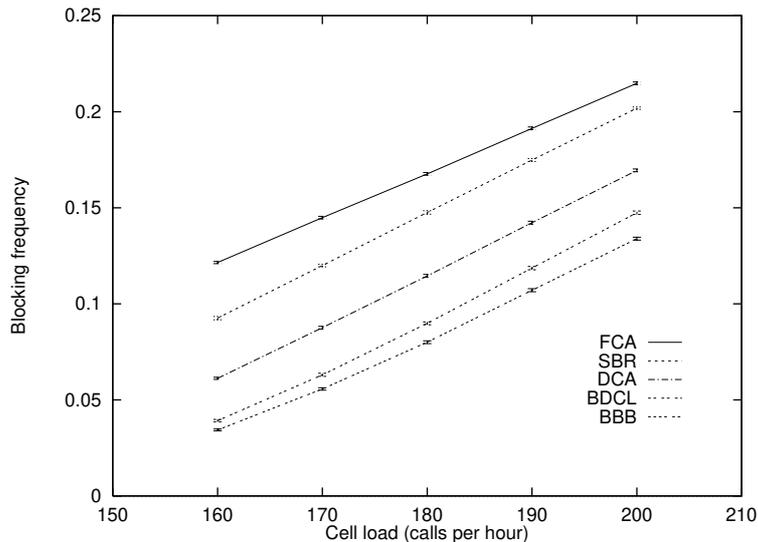| Coefficient	Value
A	7000
B	45
C	1.2625
D	0.01
E	4.17625

Figure 7: Comparison among algorithms. BBB is the algorithm proposed in this paper

For FCA, the reuse scheme (a reuse distance of two cells has been considered) consists of seven partitions of ten channels each. The Euclidean distance between the centers of two neighbors is 1, and the reuse scheme given by the function $\mathrm{res}_{ii'}$ of Section 2.2 is built by iterating the basic "knight" move of Fig. 2, which gives the same pattern as the token placement in Fig. 5. The same scheme has been used to distribute the seven channel groups among the cells for the FCA algorithm.

The coefficients for the function $J(V)$ are shown in table 1.

Each algorithm (FCA SBR, DCA, BDCL, BBB) has been simulated for connection rates of 160, 170, 180, 190 and 200 calls per hour with Poissonian distribution (and hence exponential inter-arrival time), corresponding to a traffic of 8, 8.5, 9, 9.5 and 10 erlangs. The mean duration of a connection is exponential with an average of 180sec. Each simulation consisted of 50 runs of 100000 seconds of simulated time each.

Fig. 7 shows the results of this simulation. Every error bar represents the 95% confidence interval calculated over 50 independent runs. The error bars are so small that they are hardly visible in the figures. As an example, at the cell load of 200 calls per hour, the performance (with 95% confidence interval) if of 0.2147 (0.0006) blocking frequency for FCA, of 0.2018 (0.0006) for SBR, of 0.1694 (0.0007) for DCA, of 0.1474 (0.0009) and of 0.1339 (0.0008) for BBB. The values of the error bars are similar for all other cell loads. From these experiments it is apparent that the new algorithm BBB outperforms all the competing algorithms.

The elimination of the reuse scheme and the localization of the channel packing condition as in Section 3 finally lead us to the data in Fig. 8, where we compare the original BBB algorithm with its distributed version dBBB. As it is apparent, all differences lie within the confidence intervals. In particular, at the cell load of 200 calls per hour, the performance (with 95% confidence interval)
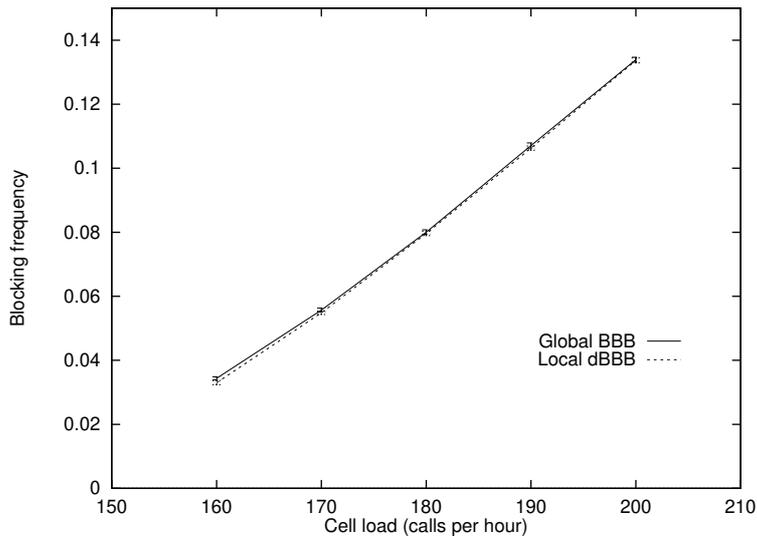
10

Figure 8: Global (BBB) vs. localized and distributed version (dBBB) of the proposed algorithm

becomes of 0.1336 (0.0008), while it was of 0.1339 (0.0008) for BBB. From these experiments one derives the conclusion that the transformation of the algorithm into a localized one does not introduce any significant performance penalty. The localized version is therefore to be preferred, also because it permits the distributed realization previously described.

5 Conclusions

After a brief description of some well known algorithms, a new algorithm BBB has been introduced, based on the exact solution of a penalty-function formulation of the problem. Furthermore, its distributed version dBBB has been presented and the algorithms have been tested through experiments over a statistically significant number of runs with varying load conditions. The conclusion to be derived from the experiments is that the new algorithms achieve lower blocking frequencies than the competitors, at least when compared on the benchmark given by regular hexagonal patterns. In addition, their low computational complexity makes them ideal candidates for real–time realizations.

Due to their nature, most algorithms have only been studied with such restrictive assumptions as "boolean" (non-additive) interference, and, in order to permit a fair comparison with the competition, our tests were restricted to that case. However, our method can easily be applied to a large variety of channel assignment problems with additive interference constraints, adjacent-channel interference and so on. In fact, all that is required for the minimization to be polynomial is that every channel provides a *linear* contribution to the total penalty function.

References

[1] R. Battiti, A.A. Bertossi, and M.A. Bonuccelli. Assigning codes in wireless networks: Bounds and scaling properties. *Wireless Networks*, 1998. To appear.

[2] O. Bertazioli and L. Favalli. *GSM — Il Sistema Europeo di Comunicazione Mobile: Tecniche, Architettura e Procedure*. Advanced Technology Education Services. Hoepli Editore, Milano, 1996.

[3] A.A. Bertossi and M.A. Bonuccelli. Code assignment for hidden terminal interference avoidance in multihop packet radio networks. *IEEE/ACM Transactions on Networking*, 3:441–449, 1995.

[4] M. Duque-Antón, D. Kunz, and B. Rüber. Channel assignment for cellular radio using simulated annealing. *IEEE Transactions on Vehicular Technology*, 42:14–21, 1993.

[5] S. Jordan and E.J. Schwabe. Worst-case performance of cellular channel assignment policies. *Wireless Networks*, 2:265–275, 1996.

[6] I. Katzela and M. Naghshineh. Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey. *IEEE Personal Communications*, pages 10–31, june 1996.

[7] E. Del Re, R. Fantacci, and L. Ronga. A dynamic channel allocation technique based on Hopfield neural networks. *IEEE Transactions on Vehicular Technology*, 45:26–32, 1996.

[8] E. Shamir and E. Upfal. Sequential and distributed graph coloring algorithms with performance analysis in random graph spaces. *Journal of Algorithms*, 5:488–501, 1984.

[9] S. Singh and D. Bertsekas. Reinforcement learning for dynamic channel allocation in cellular telephone systems. *submitted to NIPS96*, 1996.

[10] K.L. Yung and T.P. Yum. Compact pattern based dynamic channel assignment for cellular mobile systems. *IEEE Transactions on Vehicular Technology*, 43:892–896, november 1994.