Ad Hoc Resource Allocation in Cellular Systems

Timothy X Brown[1]
Interdisciplinary Telecommunications Program
Dept. of Electrical and Computer Engineering
University of Colorado, Boulder, CO 80309-0530
Tel: (303) 492–1630 Fax: (303) 492–1112
timxb@colorado.edu

***Abstract*–A fundamental question in a wireless cellular system is how to allocate base stations, radio transceivers, and channels in an efficient manner. This question becomes difficult when operators must react to burgeoning demand. This paper explores a pure ad hoc approach where call requests are given the first available channel that meets signal quality requirements, radio transceivers are added immediately to base stations without sufficient radios, and new base stations are placed immediately at a caller's location when their request can not otherwise be met. This ad hoc resource allocation has the advantage that it requires no prior planning or assessment of traffic demands. We view this as a case that bounds current cellular resource allocation practice. We show that the ad hoc performance in terms of total required resources to meet a demand in a given area is similar to more carefully planned systems that are given prior information of the traffic distribution.**

I. INTRODUCTION

Ideal two-dimensional cellular systems place radio base stations (BS) according to regular hexagonal grids [10, 11, 12]. Current cellular and PCS systems, with their small cell sizes, depart significantly from the ideal hexagonal layout due to terrain variations, difficulties in site acquisition, and space variations in mobile station density [7]. In quickly deployed ad hoc military or emergency communication systems, little or no planning may be possible. Distributed campus wireless LANs may add communication elements in a distributed fashion with little coordination between departments. Such smaller operators may not have sufficient resources or expertise for a formal deployment plan. Further, little theory guides how to evolve cellular networks over time to meet growing demand. The question we address is how to allocate network resources in such environments.

Cellular resource allocation exists on three time scales. On the seconds time scale, a call request arrives from a user and the cellular system must decide the BS and channel combination to assign to the user. At the next time scale (days or weeks) the operator may add more radio receivers to a BS that has excessive blocking. At the longest time scale an operator may choose to add additional BS to cover traffic concentrations that reach the limits of frequency reuse or to cover holes in coverage.

Cellular operators often take a demand based approach whereby they start with a deployment roughly based on hexagonal cellular design principles in order to provide sufficient coverage, and then split cells as dictated by demand [10, 11, 12]. Radio receivers are added to overloaded cells. Channels are allocated usingfixed channel assignment (FCA) or some form of dynamic channel assignment (DCA) [9, 13] that avoids the complex optimization associated with a fixed channel assignment FCA [8].

In this paper, we take this approach to an extreme in what we denote ad hoc cellular resource allocation (ACRA). ACRA will be formally defined later but can be summarized as follows. In ACRA we use the simplest form of DCA that assigns the first channel that meets the signal quality requirements. Radios are added immediately to BS with insufficient radios to carry a call request when call blocking is too high. New BS are placed immediately at a caller's location when no channel can carry the call and call blocking is too high.

We feel this algorithm has merit for several reasons. First, as noted already formal planning is difficult and ACRA can be considered as a bound on the performance of any algorithm that incorporates more formal planning. Second, although we would expect ACRA to produce somewhat random patterns of BS deployment, previous work on random deployments has shown that under log-normal shadow fading they can have performance approaching ideal hexagonal systems [3, 4, 5]. Third, we would expect ACRA to do better than random placement since the area near a BS has sufficient signal to use any available radio channel. Thus, we would expect new BS to be well spaced relative to existing BS. Finally, ACRA assumes nothing about the traffic distribution or how it might evolve over time and so provides a framework for BS deployments over time.

The next section presents a user and radio model. Section III presents the ACRA algorithm. Section IV evaluates the performance of the model under a number of different scenarios and results are presented in Section V. Section VI discusses practical aspects of the algorithm. We will show that the ACRA model yields cellular deployments that are close to the resources in an ideal hexagonal layout, and adapts to a number of different scenarios.

1. This work was completed under NSF CAREER Award NCR-9624791 and NSF Grant NCR-9725778.

II. MODEL:

This section details the model used for this paper. The model covers both sparse range-limited cellular systems and low-power, interference-limited, high-density grids of BS such as described by Cox and references therein [6, 7]. To focus on resource allocation many of the details are abstracted or stripped away.

All BS and mobiles have identical transmit power, antenna gains, etc., and the path loss is an inverse power law with path loss exponent ε. All antennas are omnidirectional with no variation in the vertical radiation pattern (i.e. they are isotropic). Rayleigh fading (aka small-scale fading [14]) is treated by micro-diversity techniques in the radio channel and considered outside the scope of this paper (so-called local mean statistics [1]). Shadow fading (aka large-scale fading) is modeled as independent log-normally distributed multiplicative noise, Ψ, on the signal strength received from each BS. It is well modeled by a log-normal density [6]:

$$p(\psi) = \frac{1}{\psi\sqrt{2\pi}\sigma}e^{-\frac{1}{2}\left(\frac{\ln\psi}{\sigma}\right)^2}, \tag{1}$$

so that the fading factor has mean 1 and one standard deviation includes from 1/σ to σ.

Users are uniformly distributed, do not move while communicating, and remain on the same channel throughout the call. N orthogonal channels are available and each channel carries one user. Being orthogonal, adjacent channel interference is ignored and only co-channel interference is considered. In an FDMA-based cellular system, each channel is a separate frequency. In a TDMA-based cellular system each channel is a frequency and time-slot pair. In a CDMA-based cellular system each channel is a frequency and spreading code pair. For simplicity, we assume N FDMA channels.

Calls are generated via a Poisson process with exponential call holding time. Call set-up is instantaneous with no distinction between user originating or user terminating calls. We consider uniform and non-uniform traffic distributions across the plain.

Channel quality is measured as the uplink and downlink carrier to interference ratio (C/I):

$$C/I = \psi_S d_S^{-\varepsilon} \Big/ \left(N_B + \sum_i \psi_i d_i^{-\varepsilon}\right), \tag{2}$$

where d_S is the distance from the mobile to the signal BS, $\{d_i\}$ is the set of distances to co-channel interfering BS's (on the downlink) or mobiles (on the uplink), and N_B is the normalized background noise power.

The performance of a deployment strategy is in terms of the total resources to meet a performance target. It is measured by fixing a required C/I threshold, T, a maximum call blocking rate, B_{Max}, and measuring the total number of BS and radio receivers required to meet C/I and blocking requirements.

III. AD HOC CELLULAR RESOURCE ALLOCATION:

ACRA consists of three components: channel assignment, radio transceiver allocation, and new BS placement. We pause to emphasize that the ACRA algorithm is unrealistic in the sense that it allows resources to be instantly allocated or reallocated, but it does provide a basis for understanding how more realistic ad hoc deployments behave (See Section VI.).

A. Dynamic Channel Assignment

Many DCA schemes are known [9] and the goal of this paper is not to find the best of these schemes. The goal instead is to show representative performance with DCA.

We consider a simple form of DCA. At each call arrival, the user attempts to access the strongest BS. The BS can assign any channel it is not using in its cell to the new user subject to two criteria. The first criteria is that the uplink and downlink for the chosen channel must meet the required C/I for the radio. The second criteria is that the addition of this new call cannot cause any existing radio link to drop below the required C/I. Within this criteria, several channels may be available which could be ranked by the carrier or interference signal strength or the C/I. A number of algorithms are available (e.g. [2]). We choose a simple algorithm. The channels are numbered from 1 to N. The smallest numbered channel that meets the criterion is assigned.

Note that the mobile restricts his access to the strongest BS. Even if the mobile queried all BS, only the BS with the strongest signal, call it b, will ever be used by a mobile in DCA. The C/I on a given channel is maximized by using the strongest available BS for the signal. If this channel is not occupied at b then b, by definition, is that BS. If this channel is occupied at b, than using any other (weaker) BS yields a downlink C/I < 1. This contrasts with FCA where if none of the assigned channels are available at b, then the mobile must attempt at another BS. Note also, that under shadow fading the strongest BS is not necessarily the closest.

All N channels are checked with every arrival. This means that BS transceivers can use any of the N channels. A call is blocked if either the strongest BS does not have a sufficient number of transceivers, or if no channel has sufficient C/I.

B. Base Station Transceiver Allocation.

The system tracks the number of call arrivals and the number of call arrivals that are blocked and computes the measured call blocking rate. If a call is blocked by the DCA algorithm, and the measured call blocking rate is less than B_{Max}, the maximum call blocking rate, then the call is blocked. If the measured call blocking rate is greater than B_{Max} and the call could be carried if the strongest BS had an additional radio transceiver, than a radio transceiver is added to the BS and the call is carried.

C. New Base Station Placement

If the measured call blocking rate is greater than B_{Max} and the call can not be carried through DCA or another radio transceiver allocation, then a new BS with one radio transceiver is placed at the location of the call requester.

When a new BS is deployed, it now becomes the strongest BS for some currently carried calls. For such calls, the call and its radio transceiver are immediately transferred to the new BS. Though in the long run this rearrangement mechanism may not be necessary, it was found in simulations that it smoothed transient behaviors.

IV. SIMULATION:

To evaluate the dynamic channel assignment for different channel and BS layouts, we use a simulated cellular system. This section describes the details of this simulation. The simulation proceeds as follows:

1. Choose scenario parameters: the total number of channels, N; total traffic in Erlangs, E, background noise, N_B; path loss exponent ε; shadow fading standard deviation, σ; required C/I threshold, T; grade of service, B_{Max}; and traffic spatial distribution.
2. Call arrivals and departures are generated according to the total Erlangs and their spatial distribution.
3. For each call arrival: the call is placed according to the ACRA algorithm.
4. For each call departure, the call is simply removed.

Each of these steps is described in detail.

The Poisson call process has E Erlangs total. In this simulation time is not important, and we only are concerned with the sequence of call and departure events. We note:

$$E = \lambda\mu\,, \tag{3}$$

where λ is the arrival rate and μ is the call holding time. The call arrival rate is fixed at λ and the call departure rate depends on the number of calls in progress, n_u. Given n_u, the rate of call departures is n_u/μ. The probability that the next event is a call arrival is given by:

$$p_{\text{arrival}}^{n_u} = \frac{\lambda}{\lambda + n_u/\mu} = \frac{E}{E + n_u}\,. \tag{4}$$

Otherwise, the next event is a departure. We divide the simulation into two phases. A construction phase where resources are added via ACRA for a period of A_A arrivals. After this initial period blocking statistics on the next A_E arrivals is recorded in an evaluation phase.

For call arrivals, the user is first placed in a disc of unit radius centered on the origin. We consider two distributions. The first is a uniform distribution across the disc. The second chooses a radius and angle from the origin. The angle is chosen uniformly from $[0, 2\Pi]$. The radius is given by $|x|$ mod 1, where x is a normal Gaussian. The latter user distribution, denoted Gaussian, simulates a variable user concentration across an area that occurs in practical systems.

The signal power from the user to each BS is defined via:

$$p_{ij} = \Psi_{ij} K d_{ij}^{-\varepsilon} \tag{5}$$

where p_{ij} is the power of the signal from user i to BS j, K is a constant containing the transmit power, antenna gains, etc., d_{ij} is the distance from user i to BS j, ε is the path-loss exponent, and Ψ_{ij} is a the log normal shadow random variable distributed as in (1). Since we are only concerned with power ratios and all BS and mobiles are identical, we let $K = 1$ and use p_{ij} equally for both up and down link powers. If c_i is the channel of user i, and b_i is the BS with the strongest signal to user i, then the carrier to interference ratio for channel c at user ν follows from (2):

$$\left(\frac{C}{I}\right)_{\text{downlink}}^{c,\,\nu} = \frac{p_{\nu b_\nu}}{N_B + \sum_{\{i \mid c_i = c,\, i \neq \nu\}} p_{\nu b_i}} \tag{6}$$

$$\left(\frac{C}{I}\right)_{\text{uplink}}^{c,\,\nu} = \frac{p_{\nu b_\nu}}{N_B + \sum_{\{i \mid c_i = c,\, i \neq \nu\}} p_{i b_\nu}} \tag{7}$$

Channel c meets all C/I criteria if (6) and (7) are above a required C/I threshold, T, for every ν (i.e. for the new call and every existing call). If no channels meet the criterion or b_ν has no available radios, the call is blocked, otherwise the smallest numbered c that meets the criterion is assigned.

The blocking probability, p_B, is simply the ratio of the number of blocked arrivals over the total number of arrivals.

In evaluation mode, blocked calls are cleared, and only the DCA component of ACRA is used. In ACRA mode new resources can be added to carry the blocked call if $p_B > B_{Max}$, otherwise, the blocked call is cleared. If the call is blocked because b_ν has no available radios, then a new radio is added to b_ν and the smallest numbered c that meets the criterion is assigned. If the call is blocked because no channel meets the C/I requirements, then a new BS is constructed at the location of the blocked caller. The downlink and uplink signal strengths

Table 2: Experimental Results

Scenario Number	User Distribution	Min Num BS, M	Shadow Fading, σ	Hex		ACRA	
				Total BS	Total radio	Total BS	Total radio
1	Uniform	1	0dB	21	10600	29	10700
2	"	1	10dB	22	10700	28	10700
3	"	100	0dB	115	11600	183	12100
4	"	100	10dB	135	12100	267	12500
5	Gaussian	1	0dB	n/a	n/a	52	11000
6	"	1	10dB	n/a	n/a	52	11000

Table 1: Simulation Parameters

Parameter	Symbol	Values
Number of Channels	N	1000
Total Erlangs	E	10000
Min number BS to Cover	M	1 or 100
Pathloss exponent	ε	4
Shadow fading std. dev.	σ	0 or 10dB
Required C/I	T	10dB
Grade of Service	B_{Max}	2%
Arrivals during ACRA	A_A	100,000
Arrivals during evaluation	A_E	100,000

(which are very large) are guaranteed to meet the C/I criteria for this caller, and the call is only blocked if every channel would cause existing users to fall below the C/I threshold. In the simulations that followed, this was never the case. At this point the new BS may be the strongest BS for some already existing calls. These calls and their radios are moved to the new BS.

Table 1 gives a parameter summary. The number of user channels is comparable to the number an operator would have with existing standards such as AMPS (395), IS-54 (1185), or GSM (500). The total Erlangs is typical of a single carrier's traffic in a medium sized metropolitan area. The channel parameters and GOS are typical values from [12]. The background noise is specified indirectly in terms of the minimum number of BS to cover the unit radius circle. Assuming M BS have equal coverage area and they are hexagons, then these M BS would have the same total area as the coverage area if:

$$r_M = \sqrt{\frac{2\pi}{M 3\sqrt{3}}} = \frac{1}{\sqrt{0.8270 M}}. \tag{8}$$

Next, we choose a background noise so that with no interferers or shadow fading a user at distance r_M from a BS would just meet the C/I criteria in (6) and (7). Combining with (5):

$$N_B = (0.8270 M)^{\varepsilon/2}. \tag{9}$$

Defining N_B in this way allows us to specify whether we want to consider a range limited system—where a large number of BS are required just to get signal coverage (M is large), and interference limited system—where multiple BS are required solely to generate enough capacity (M is less than or equal to 1).

Each scenario is repeated 10 times and averages reported.[2] We compare the results with a hexagonal system which uses DCA. The hexagonal systems are generated by choosing a number of BS, scaling the BS separation so that these BS lie in the unit radius coverage area, and assigning enough channels to each cell to carry the Erlangs captured by each BS with at most 1% blocking. The correct number of BS for comparable performance is found by trial and error. This comparison using a grid of equally spaced BS only makes sense for the uniform distribution.

V. Results:

This section presents results for the ACRA algorithm for the four combinations of interference limited (M = 1) vs. range limited (M = 100) and shadow fading (σ = 10dB) vs. no shadow fading (σ = 0dB) in the uniform user distribution. Two more cases include M = 1, shadow fading vs. no shadow fading and the Gaussian user distribution. Since the traffic distribution is not critical in the range limited case, we do not consider the range limited case with the Gaussian distribution.

Results are shown in Table 2. We note four observations. First, for the interference limited experiments, the total resources do not depend significantly as a function of shadow fading, while in the range limited experiments, the resources increase with increasing shadow fading variability.

Second, the total channel resources are at least 10,000 channels and depend mainly on the number of

2. The final paper will average a larger number of repitions.

BS. This follows from the 10,000 Erlangs of total traffic and the number of trunk groups it is divided into (one per BS).

Third, in the range limited scenario, the ACRA algorithm requires nearly twice as many BS as the hexagonal layout. In the interference limited case, the amount of BS with shadow fading (which would be present in practice) is about 25% more with the ACRA algorithm compared to the hexagonal layout. This may be somewhat surprising considering the completely ad hoc manor in which the resources are employed. This suggest careful planning and optimization is indeed useful in the early stages of cellular design that are mainly concerned with getting coverage, while in later stages where capacity is the main issue, more ad hoc deployments are justified.

Finally, we note that the algorithm works as readily on the uniform user distribution as the non-uniform Gaussian distribution.

VI. ACRA in Practice

The ACRA algorithm is unrealistic for a number of reasons. First, there are time delays between when a need for resources is identified and when the resources are deployed. In practice, operators monitor call block rates at BS and deploy sufficient radios as the blocking rates become too large. Similarly, holes in coverage, and cells that require splitting to increase capacity are identified, and future sites are mapped out well in advance of the actual need. None-theless, the ACRA results suggest that no matter how little fore-sight is put into this process, as long as resources are deployed where there is a demonstrated need, the total resource usage should be close to that of what an omnipotent operator with full information of traffic demands would build.

Second, the ACRA algorithm deploys radio resources one channel at a time. TDMA and CDMA, and in fact most FDMA based equipment is deployed in minimal blocks of user channels. Further, operators proactively deploy a minimum number of channels. Since the total number of radio channels was similar across different scenarios, it is unlikely these factors would make a significant difference in the end.

Third, BS in ACRA are deployed at the blocked callers location whereas in practice many factors constrain the location of BS. Again, these considerations are likely to only improve on ACRA's good results.

VII. Conclusion

This paper presented ad hoc cellular resource allocation (ACRA), a very simple algorithm for deploying BS and radio transceiver resources in a wireless cellular network. The algorithm was able to find good deployments with approximately the same number of total radio transceivers and 25% more BS compared to an ideal hexagonal layout with uniform distribution of users in an interference limited environment. In a range limited environment, it required approximately twice the resources and justifies the planning used by operators when they first deploy their system. But, as the system matures, when providing sufficient capacity to spatially varying traffic is the goal, ACRA-like algorithms that deploy BS driven more by traffic demands than maintaining any sort of grid are justified.

Further work is building on these initial results including user mobility, and handoff analysis.

References

[1] Bernhardt, R.C., “Performance aspects of two-way transmission in portable radio systems,” *Proc. IEEE VTC’89*, San Francisco, CA, May 1–3, 1989, pp. 279–84.

[2] Bernhardt, R.C., “Time-Slot Management in Digital Portable Radio Systems,” *IEEE T. on Veh. Tech.*, v. 40, n. 1, Feb. 1991. pp. 261–72.

[3] Brown, T.X, “Analysis and Coloring of a Shotgun Cellular System,” in *Proc. of the IEEE Radio and Wireless Conference (RAWCON) 98*, IEEE pub, 1998, pp. 51–4.

[4] Brown, T.X, “Dynamic Channel Assignment in Shotgun Cellular Systems,” n *Proc. of the IEEE Radio and Wireless Conference (RAWCON) 99*, IEEE pub, 1999, pp. 147–150.

[5] Brown, T.X, “Practical Cellular Performance Bounds via Shotgun Cellular Systems,” *IEEE JSAC*, v. 18, n. 11, Nov. 2000, pp. 2443–2455.

[6] Cox, D.C., “Universal Digital Portable Radio Communications,” *Proc. of the IEEE*, v. 75, n. 4, pp. 436–77, Apr. 1987.

[7] Cox, D.C, “A radio system proposal for widespread low-power tetherless communications,” *IEEE T. on Comm*. v. 39, n. 2, pp. 324–35, Feb. 1991.

[8] Duque-Anton, M., et al., “Channel assignment for cellular radio using simulated annealing,” *IEEE. T. on Vehicular Technology*, v. 42, n. 1, p. 14–21.

[9] Prasad, R., *Universal wireless personal communications*, Artech House, 1998, ch. 14.

[10] Lee, W.C.Y., *Mobile Communications Design Fundamentals, 2nd Ed.*, John Wiley & Sons, 1993. p. 210.

[11] Macdonald, V.H. “The Cellular Concept,” *The Bell System Technical Journal*, v. 58, n. 1, Jan 1979, pp. 15–41.

[12] Rappaport, T.S., *Wireless Communications: Principles and Practice*, Prentice Hall, 1996.

[13] Singh, S., Bertsekas, D., “Reinforcement learning for dynamic channel allocation in cellular telephone systems,” In *Advances in Neural Information Processing Systems 9*. MITPress, 1997. pp. 974–80.

[14] Sklar, B. “Rayleigh fading channels in mobile digital communication systems Part I: Characterization” *IEEE Comm Mag.*, pp. 90–109, Jul. 1997