

Mobility prediction's influence on QoS in wireless networks: a study on a call admission algorithm

Jean-Marc François and Guy Leduc
Research Unit in Networking
Department of Electrical Engineering and Computer Science
Université de Liège, Belgium
{francois,leduc}@run.montefiore.ulg.ac.be

Abstract

Several mechanisms increase the QoS level of mobile networks thanks to an underlying mobility prediction method (i.e. a means to predict a mobile's next access router). This paper aims at studying how the accuracy of the prediction method can influence the network QoS in the particular context of call admission control. It shows that (a) the mobiles behaviour must be adapted according to the prediction scheme accuracy in order to achieve good performance and (b) the admission algorithm can be modified to increase its fairness and to give mobiles an incentive to do such an adaptation.

1. Introduction

The quick development of wireless networks brings new challenges to the network community. In particular, it is commonly admitted that those networks' next generations should provide new kinds of multimedia services requiring a high QoS level.

Unfortunately, wireless networks introduce large delay jitters, poor bandwidth and high error rates, making decent QoS levels difficult to achieve; this is particularly true when *mobile nodes* (MN) experience frequent handoffs causing abrupt network conditions' modifications. This is where *mobility prediction* takes place.

Mobility prediction's purpose is to obtain information related to the next *access router(s)* (AR) a mobile node will be linked to. The underlying idea is that nothing but a proactive method can allow efficient counter-measures against packet

losses, packet reordering, authentication delays,... caused by a handoff.

Defined in such a general way, one can easily understand that a lot of methods are related to mobility prediction, giving very varied results in terms of precision and quality.


Admission control and prediction are directly related since a frequent admission control requirement is to reserve resources so as to improve the network's future QoS; the amount of resources to reserve at a given time thus depends on the users upcoming behaviour.

This diversity leads to a situation where various (sometimes similar) admission control algorithms are simulated with very different prediction schemes.

Estimating the added value brought by a more precise prediction method is not an easy task; to which extent is it interesting to implement a more reliable (and often more complex) method? This paper aims at answering this question in the particular context of call admission control. We will see that this analysis leads to an improved admission method aimed at improving both network performance and fairness; this admission method can be implemented in a simple or in a more complicated way, yielding results in accordance with the added complexity.

The remainder of this article is organized as follows. The first parts explain what admission control is, how it is related to mobility prediction and how this prediction can be modeled. Two admission methods are then defined. Sections 4, 5 and 6 give the results of simulations involving those methods and show the effect of prediction accuracy. Section 7 shows how an admission control scheme can be modified to give mobiles an incentive to improve their predictions. The last section concludes this paper.

This work has been partially supported by the Belgian Science Policy in the framework of the IAP program (Motion P5/11 project) and by the European E- Next NoE.





Once the mobile has entered the current cell, after a certain amount of time, a prediction trigger is emitted.

The prediction scheme predicts the behaviour of the mobile until its next handoff (the dashed line).

One of the basic purpose of a prediction scheme is to identify the mobile's next cell. The opposite figure shows a wrong next cell prediction.

The abstract model's π **parameter** is the ratio between right predictions and the total number of triggers emitted.



A prediction scheme *can* predict, at each trigger, the amount of time before the mobile's next handoff.

The abstract model's δ **parameter** is the mean trigger-to-handoff duration; the σ **parameter** is the standard deviation of this duration, which is considered a random process.

Figure 1. This figure depicts a simple abstract model that modelizes the main characteristics of a mobility prediction scheme.

2. Call Admission Control and mobility prediction

2.1. Overview

Admission control is tightly linked to mobility prediction, because reservations must be done in advance and at the right moment. This prediction is usually done statistically: mobiles' movement behaviours (*e.g.* mean cell travelling time, most likely next AR(s)) and call habits are analysed and averaged, taking care of their possible variations during the day and throughout the week.

Others use simple ([6]) or more advanced ([4, 11, 12, 13]) techniques to track the individual motion of each mobile, hence getting an accurate idea of when and where each of them will handoff. A natural way to tackle this issue is to estimate each mobile's position (*e.g.* [8]), but other methods can be considered (*e.g.* [5]).

A mobile's next cell can also be inferred from its typical travelled cells ; the prediction then amounts to guess

the next cell given the previous ones. This can be achieved by various methods such as information theory ([2, 9]) or Markov chains ([7]).

2.2. An abstract model

The following aims at showing how prediction's performance influences CAC's effectiveness without restricting ourselves to a specific prediction mechanism. It has been chosen to model any mechanism with an abstract model defined by a few parameters. It is depicted in figure 1.

With a few values, this simple model summarizes the main information brought by a prediction algorithm. Those parameters should have an intuitive meaning and act upon the behaviour of the method using predicted events. At each trigger, the prediction mechanism reports an identifier of the next AR and the estimated delay before the next handoff (δ seconds). This last measurement is optional as some schemes have not been designed to give this information. The trigger-to-handoff delay is only estimated by the prediction algorithm, so the difference between this estimation

and the real delay can be seen as realizations of a random distribution; let σ be the standard deviation of this distribution.

As one can see, this abstraction is only a rough sketch and several real systems' characteristics have been ignored, mainly:

- At most one trigger can be emitted before a given handoff;
- The prediction success ratio is fixed once for all; one cannot model behaviours such as an increasing prediction success ratio when the trigger is delayed.
- Each trigger only gives the most likely next AR (it could guess the n next ARs, or several ARs with their probability of being next);
- The only information given about the trigger-to-handoff duration are its mean value and standard deviation. Nothing more is given about the duration distribution which hence will most probably be chosen gaussian.

Let's now study the effect of those parameters on a concrete CAC algorithm.

3. The PCR algorithm

This CAC algorithm has been introduced in [4] by M.H. Chiu and M. Bassiouni; it is a typical CAC using mobiles' movements guesses extensively. Its aim is to reduce the handoff blocking probability in a cellular network by means of reserved channels. It is briefly described below; the interested reader can refer to [4] for more information.

Several mobiles are moving in a wireless network made of cells having a fixed number of channels. All the cells are assumed to have the same number of channels, and a given mobile can only use one channel at a time. The effect of soft handoffs, frequency borrowing or space diversity techniques is not studied here. This implies that the proposed CAC scheme quite well matches 2G GSM but should be re-evaluated in order to be applied to 3G networks ([10, 3, 1]).

Each mobile position is monitored so as to extrapolate its next cell, C . If a mobile has an ongoing call, it reserves a channel of C ; if there is no free channel, the reservation is pushed on a unbounded FIFO reservation queue (a reservation is popped out every time a channel is freed until the queue is empty). When a mobile wants to start a new call, it has to find a free (unreserved) channel in its current cell; if there is none, the call is rejected. A mobile can of course cancel a previous reservation if it finds out that it will be useless (because of a sudden motion change or because its current call has ended).

[4] proposes three different ways to deal with handoffs, called PCR1, PRC2 and PCR3. We won't study the intermediate algorithm PCR2 here, but focus on extreme cases, explained below. When a handoff occurs:

PCR1	If the call has a prior reservation, it is allocated. Else, if there is a free channel, it is allocated. Else, the call is dropped.
PCR3	If there is a reserved channel, it is allocated. Else, if there is a free channel, it is allocated. Else, the call is dropped.

As one can see, the difference is that channels can be reserved either for a particular mobile (in that case, every reservation holds a mobile identifier) or "anonymously". This algorithm can be studied using the abstract prediction scheme introduced before. Note that as this scheme only allows one prediction per handoff, the only way to notice a wrong prediction is to wait for the actual handoff (and to see if it occurs with the predicted cell). The PCR algorithm and the way it has been mixed with the abstract prediction scheme have been summarized in figure 2.

4. Simulations

We can now study how PCR behaves according to the parameters of the prediction scheme. The aim is to find what characterizes a *good* predictive algorithm (in the CAC context) and to what extent its performance can modify those of PCR.

4.1. Simulations parameters

The simulations presented here closely mimic those of [4]. Their setup is the following.

The terrain is composed of 30 circular cells laid-out as a bee's nest. The cells have a 1000m radius and overlap, thus, introducing a "handoff hysteresis" (see figure 3). The mobility model has been chosen so as to approximate the motion of a mobile in a random direction¹. Both this direction and the mobile speed is updated at each time step according to a gaussian distribution. This model is precisely described in [14] (each distribution has a unitary variance in radian for the direction, in meters per second for the speed, the mean speed is 18 m/s and the randomness factor a has been fixed to 0.4; this yields a mean cell traveling time of 110s). Figure 3 shows how a typical path looks like.

Using several cells (instead of doing the experiments with only one) allows to use reasonable cell traveling time and channel holding time distributions, without relying on approximations such as those presented in [15].

¹The mobility model used in [14] is not described precisely enough to be implemented here; thus, another realistic one has been used.

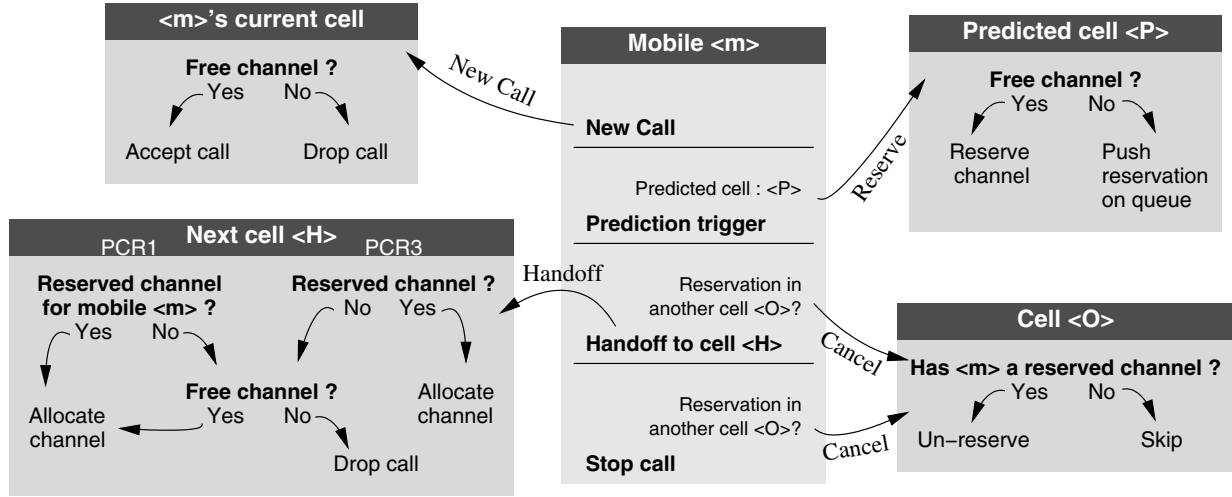


Figure 2. The PCR algorithm. This figure depicts the way PCR1 and PCR3 algorithms operate and how they can interoperate with an abstract mobility prediction scheme.

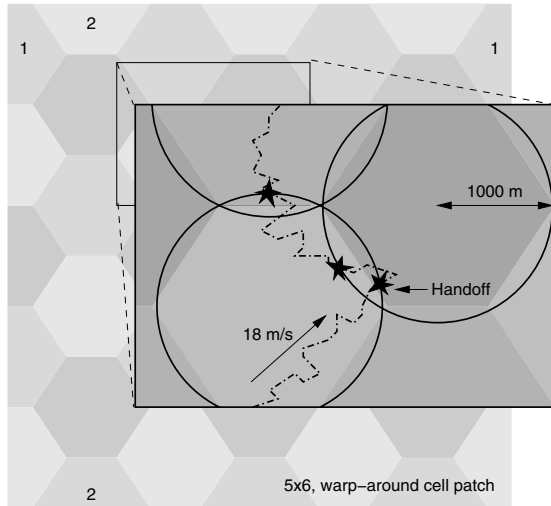


Figure 3. The simulation setup. This picture gives an overview of the cells' organization. The dashed line shows a simulated mobile motion.

Each simulation involves 4000 mobiles and lasts 12 hours.

The calls' duration is exponentially distributed with a mean of 3 minutes. New calls arrive according to a Poisson process whose frequency is adapted so as to reach a target

cell load, defined as

$$\text{Cell load} = \frac{\text{Arrival rate to the cell} \times \text{Average call duration}}{\text{Number of channels per cell}}$$

Each cell has 18 channels. Each simulation is done with a given cell load. Both PCR1 and PCR3 have been tested.

As stated before, the handoff prediction model is abstracted using 3 parameters: δ (the trigger-to-handoff duration), σ (δ 's standard deviation) and π (the prediction success ratio); their effect is studied in the following. When a missed prediction occurs, the predicted cell is chosen randomly in the cells neighbouring the one where the mobile will eventually go. The trigger-to-handoff duration is modelled by a gaussian distribution with mean δ and variance σ^2 ; the trigger is discarded if this duration is greater than the mobile's cell travelling time.

4.2. Simulations results

4.2.1 Handoff blocking ratio

The result of the simulations is given in the graphs composing figure 4. A reservation is done as soon as a trigger is emitted; the reservation duration is thus equal to the trigger-to-handoff delay. The new call blocking rate is an important factor that helps estimating a CAC scheme's efficiency; it is studied in the next section.

All the plots draw handoff blocking percentages as a function of the successful prediction ratio and the trigger-to-handoff delay (here always exactly equal to δ , that is $\sigma = 0$; section 5 studies the effect of σ). Prediction ratios as low as 0.2 have been studied; it is only a little bit higher than the

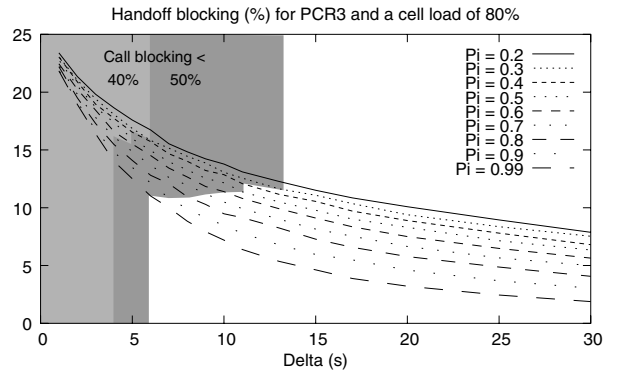







Figure 4. Handoff blocking rate versus prediction parameters. The left column is related to the PCR1 algorithm, the right-hand one to PCR3; side by side plots share the same cell load. The meaning of the grey zones is explained in section 4.2.2.

worst prediction scheme, which would take the next cell at random among the neighbouring cells.

The zero δ borderline case shows what happens when no reservation takes place (since every reservation lasts a very short amount of time); this explains why all the lines merge when δ is small and gives a reference blocking ratio that should not be exceeded (since in this case introducing CAC is worse than doing nothing).

All the graphs related to the same algorithm share a common look. In all cases, a better prediction clearly gives a better blocking ratio in return. The results obtained with a 50% cell load are not very significant: the blocking rates are so low that they make reservations useless.

PCR1 performance are maximized when the reservations last about 5 seconds; long reservations can even give worse results than no reservation at all. As expected, when the prediction is more accurate, the reservations can last longer (up to 10s) since there are less likely to block a channel in a wrong cell. With this particular algorithm, the reservation duration is an important parameter that is discussed in section 5.

The plots related to PCR3 have a very different shape. For a given prediction ratio, the handoff blocking curve decreases with the reservation duration; one can thus see the advantage brought by the fact that reservations are not associated with a particular mobile anymore.

4.2.2 New call blocking ratio

Even if it is preferable not to allow a new call rather than interrupting an ongoing one, one should not forget that there is a tradeoff between those two blocking ratios.

Figure 5 shows a plot for a PCR1 cell loaded at 70% in the same conditions as those explained in section 4. All the points of the decreasing section of the curves are optimal: there is no strict rule to choose one of them.

One can see that as the next cell prediction becomes more accurate, reaching the curve's minimum requires higher a call blocking. A way to approach this minimum while keeping the call ratio between reasonable bounds for high values of π is to restrict the allowed value of the derivative. The figure's grey zone holds the points where reducing the handoff blocking by one percent costs less than three percents in call blocking. This zone has also been drawn on the first column of figure 4.2.2.

In the curves depicted in figure 4.2.2 for PCR3, another strategy has been used. It has been chosen to color the plots according to the call blocking value. The curves' intersection points with those colored zones show how the handoff blocking varies with respect to the next cell prediction accuracy for a constant call blocking value.

5. Timing

5.1. The ideal timing

One could try to take advantage of the local minimum shown in the plots related to PCR1 (figure 4).

In the context presented here, reserving T seconds on average amounts to wait $\delta - T$ seconds after each trigger, yet this suppose that the error on δ is symmetric and has a mean of zero. In what follows, we will suppose that the distribution of this error is gaussian and has a standard deviation given by σ .

This method can't of course be applied to prediction schemes that don't estimate a δ value associated with each trigger.

To verify the effectiveness of this idea, we can use simulations similar to those presented in section 4. Figure 6 shows how the handoff blocking ratio changes with respect to the prediction ratio π and the prediction accuracy σ . In this figure, the mobiles reserve a channel during $N(T, \sigma)$ seconds² before the handoff; negative values are discarded (no prediction trigger is emitted). The other parameters are the same as those used in section 4, the cell load has been fixed to 70% and the chosen optimum is the minimum of the corresponding plot in figure 4 (and is thus a function of π).

² $N(T, \sigma)$ is a random variable from a gaussian distribution with a mean equal to T and a standard deviation equal to σ .

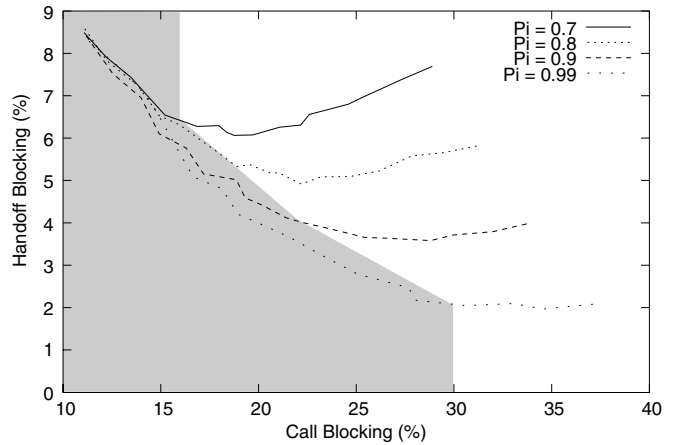


Figure 5. Handoff versus call blocking rate in a cell loaded at 70%, using PCR1. The grey zone shows points where the absolute value of the derivative is higher than $1/3$: gaining 1% in handoff blocking costs less than 3% in call blocking. The points at the border of this zone could thus be considered as optimal.

One can verify that the values obtained when σ equals 0 are the same as those appearing in figure 4.

5.2. Tradeoff

The resulting plot shows the effect of a rough handoff time estimation. As one could have guessed, the next cell prediction is the main goal in order to achieve good blocking ratios; however, figure 6 shows that this assumption is not always true for low values of π .

For example, to improve the blocking ratio by 1%, a prediction scheme characterized by (π, σ) parameters equal to $(.6, 5)$ can reach $(.6, 1)$ or $(.7, 5)$. For most schemes, this makes a serious difference and one of those points will be easier to achieve than the other.

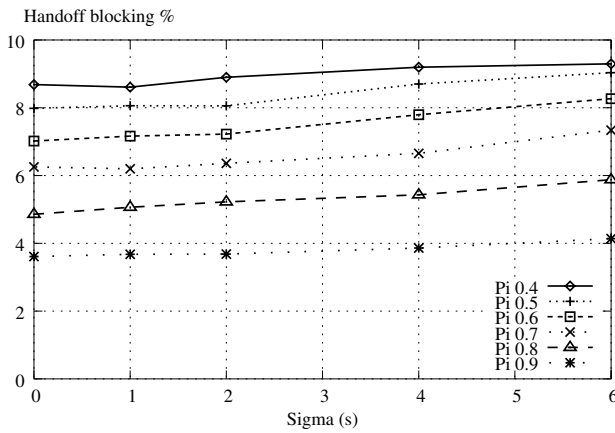


Figure 6. The effect of π and σ on PCR1 performance. This plot is similar to those given in figure 4, where the reservation delay is always 5 seconds plus a random, gaussian value characterized by a standard deviation of σ seconds.

6. Optimal duration as a function of network parameters

We have seen that PCR1 has an optimal reservation duration once the mobiles movement behaviour has been fixed. One could wonder how to choose a reservation duration given a certain wireless network's mobiles characteristics.

If one considers the mobiles' calling behaviour (*i.e.* the calls frequency and duration, which determine the cell load) and prediction parameters as fixed, the only way a mobile can modify the handoff blocking probability is by changing its movement pattern. In this context, the mean duration between the handoffs is certainly the most important factor.

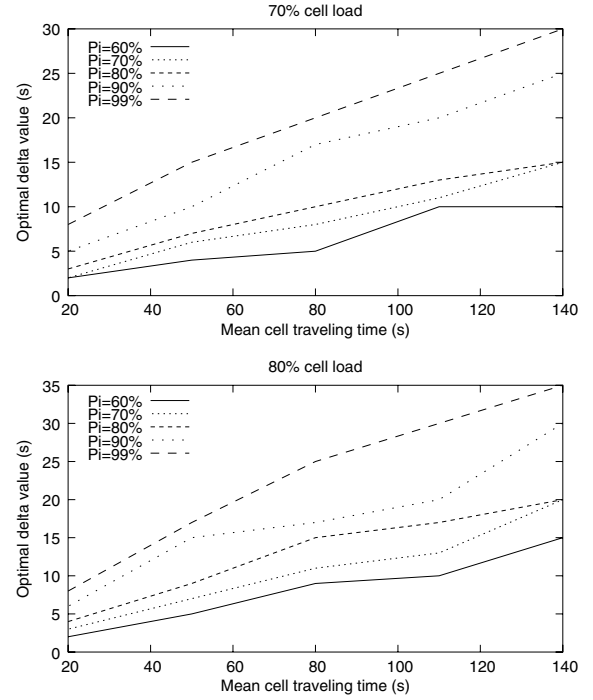


Figure 7. Optimal reservation duration as a function of the mean cell traveling time with different cell loads and next cell prediction ratios. The prediction parameter σ is equal to zero.

Figure 7 thus plots the optimal reservation duration as a function of the mobiles' mean cell traveling time. The simulations that yielded those results are similar to those presented in section 4, but with various mobile speeds.

As expected, the optimal duration decreases with the traveling time. This is a direct consequence of the higher handoff frequency: without shorter reservations, the reservation queue would exceed its optimal length.

It is not surprising to see that the various curves' slope grows together with the next cell prediction accuracy because (a) the higher the value of π , the longer the optimal duration (for a fixed mobile speed) and (b) when the speed is high, durations are shorter and their absolute difference tends to get thinner.

7. Merit-based Admission Control

The results presented in the previous section lead to a natural idea: it would be desirable to modify the CAC scheme in order to favour mobiles which are trying to reserve channels at a proper time (*i.e.* minimizing the handoff blocking of the whole system).

Method	Cell load		
	60%	70%	80%
PCR1	2.1%	8%	17.1%
Merit (simple)	2.0% (7.8%)	7.1% (12.8%)	16.7% (2.3%)
Merit (probabilistic)	1.9% (9.2%)	6.7% (19.6%)	15.2% (12.6%)

Table 1. Handoff blocking obtained with and without the different merit based admission control.

PCR1 can be modified as follows to achieve this goal; when a mobile m performs a handoff:

1. If it has a reserved channel, it is allocated;
2. If there is a free channel, it is allocated;
3. If the mobile has a pending reservation (*i.e.* in the reservation queue), it applies a *fairness test* (discussed below) which tries to find a less meriting reservation r' . If it succeeds, the channel associated with r' is allocated to m and r' is pushed back at the beginning of the queue (r' will be the next reservation to be popped out of the queue since it was already associated with a channel).
4. Else, the handoff is blocked.

Thus, the idea is that if a mobile m has a pending reservation when it performs its handoff, it should be allowed to borrow the effective (*i.e.* associated with a channel) reservation of another mobile m' if the reservation behaviour of m is more likely to reduce the overall blocking ratio than m' 's. A first, simple *fairness test* is thus one that favours mobiles near the optimal reservation duration T pointed out in section 5.1. Let t be the duration³ of the pending reservation r of the mobile performing a handoff and t' the duration of an effective reservation r' that could be borrowed.

Simple test

This test allows to borrow r' if $|T - t| < |T - t'|$; this simply means that a reservation can be borrowed if it is further apart the optimal time (*i.e.* reservations close to T will be left untouched).

The modified algorithm has three strong points.

First, it adds some fairness to the system; very long reservations are not profitable anymore, since the longer, the more likely they are to be replaced by a shorter, more sensible one thanks to the *fairness test*. One could say that this scheme favours mobiles meriting it, *i.e.* those that try to get a fair amount of the network's resources.

Second, since a reservation that lasts for about T time units is rewarded, the overall handoff blocking ratio decreases.

³The *duration* of a reservation is defined as the amount of time between the reservation creation and the mobile handoff.

Third, the modification has a minimal impact on the new call blocking ratio. Indeed, the above procedure's third item does not add any new reservation, but borrows a channel to allow a handoff which would otherwise have been blocked. It could prevent a new call to take place, but allows a hand-off to occur instead: this is exactly the spirit of the admission control we want to define.

Figure 8 plots the results of a simulation of this algorithm. Mobiles have been divided into four classes with different mean reservation durations: 5, 10, 20 and 30 seconds (gaussian distributed with a standard deviation of 1 second, the distribution is thus a *gaussian mixture*); they are experiencing cell loads of 60, 70 and 80%. Each class has the same number of mobiles and all of them use a π factor equal to 0.7. Since the simulation parameters have been chosen equal to those taken in section 4, the ideal reservation duration T has been fixed to 8, 10 or 15 seconds (depending on the cell load and according to the results plotted in figure 4).

This figure shows the results achieved by another *fairness test*:

Probabilistic test

Given the distribution of reservation durations d , one can compute the expected distance to T of r' knowing that it is at least t' . The test is thus $|T - t| < E[|x - T|]$, where x is the random variable associated with reservation durations (thus greater than t'):

$$E_{x \geq t'}[|x - T|] = \frac{\int_{t'}^{\infty} d(x)|x - T| dx}{\int_{t'}^{\infty} d(x) dx} \quad (1)$$

Notice that this requires the knowledge of the distribution d , which imposes monitoring mobiles' reservations durations.

Both fairness tests have a common drawback: mobiles doing reservations at the right moment still could have their reservations borrowed; they only are favoured statistically.

Choosing a different anchor timing could still add fairness to the system; here, fairness and performance are orthogonal issues.

The simulation's results emphasize that PCR1 is not fair: mobiles reserving channels for a longer period of time experience a better service at the expense of the global hand-off blocking ratio. On the other hand, the results of merit-based methods are divided up homogeneously. In particular,




Figure 8. This simulation shows the handoff blocking ratios experienced by four classes of mobiles under a 60, 70 and 80% cell load. They are characterized by different mean reservation durations (from 5 to 30 seconds), otherwise the simulation parameters are the same as those described in section 4. The chosen value of T are consistent with the results shown in figure 4.

classes characterized by low and high δ parameters systematically see their blocking decrease or increase, respectively.

The advantage of the more complicated, probabilistic test is revealed by table 1. It clearly shows that it performs better in terms of handoff blocking.

The figures in parentheses give the relative difference with PCR1. The simple test doesn't take care of the fact that reservation lasting for about T time units are likely to last during a much longer amount of time. This behaviour is particularly harmful to mobile classes with small δ values.

As expected, the call blocking is not that much affected by the merit-based schemes; relatively to PCR1, the simple test increases it by about 2% while the probabilistic one increases it by about 4%.

The difference between the two merit tests is brought to the fore by figure 9. The first plot shows the reservation time estimated by the probabilistic (equation 1) and the simple test (which approximates $E_{x \geq t'}[|x - T|] \simeq |t' - T|$). Several reservation duration densities $d(\cdot)$ have been considered:

- the gaussian mixture used in the simulation above (*i.e.* 4 gaussian distributions centered on 5, 10, 20 and 30 seconds with an extended deviation of 1 second);
- uniform densities over $[0, M]$ with $M = 10, 20$ or 30 seconds.

The optimal time is $T=10s$.

This plot helps to see which effective reservations can be borrowed: a mobile that has a pending reservation that lasted t seconds can borrow an effective reservation that has already lasted t' seconds if the curve ordinate at abscissa t' is greater than $|t - T|$.

The second plot draws the difference between the curves relative to the probabilistic test and the curve relative to the simple test. The positive portion of the curves shows when

borrowing an effective reservation is easier with the probabilistic test (*i.e.* borrowing is allowed for a longer interval centered on T).

The simple test is clearly over-pessimistic when most reservations last less than the optimal time, and over-pessimistic when dealing with long reservations. This fact explains the different behaviour of the merit tests. Considering the distribution related to the simulation above, one can observe that a mobile performing a handoff with a pending reservation that lasted t seconds—with $\pm 1 < t < \pm 20$, *i.e.* a difference to the optimal time less than about 9s—can borrow any other effective reservation using the probabilistic estimation, but cannot borrow those that lasted between t and $20 - t$ seconds with the simple test. Being centered on the optimal time, the 10s class does not experience this difference (for those mobiles, the $[t, 2T - t]$ range is small). The 5s class is, on the contrary, badly off: its mobiles cannot borrow reservations in the 5 to 15 seconds range (this is important since most reservations last more than 5 seconds, and explains the high blocking ratio experienced by this class in the simulations). Longer reservations experience less problems, since the probabilistic curve come closer to the simple test's and less mobiles are likely to be involved in long reservations.

8. Conclusion

It is usually admitted that motion prediction can enhance the quality of service provided by wireless networks. This paper discusses how the accuracy of the prediction acts on the level of service in the particular case of cell access control.

It has been shown that, with certain access control protocols, reservations must be done at the proper time; this

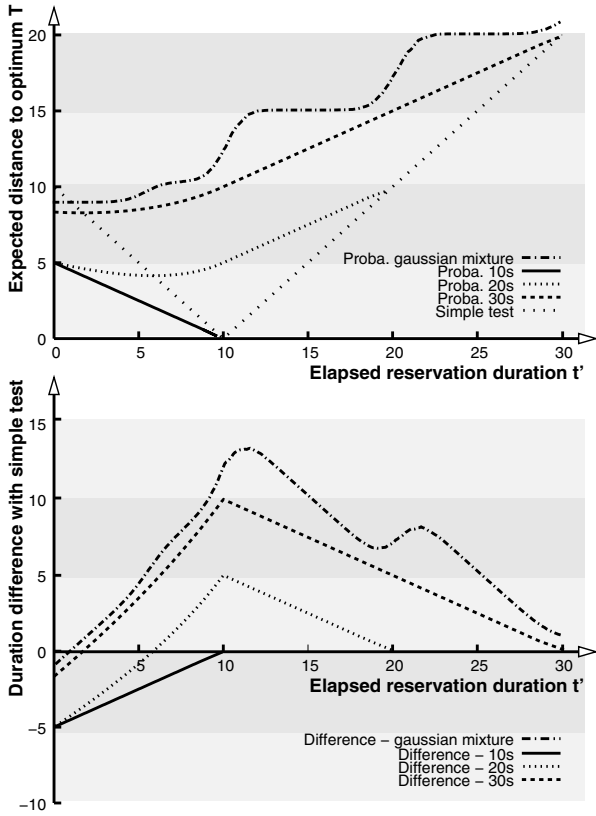


Figure 9. Comparison between the probabilistic and simple tests. The curves of the first plot draw the reservation duration estimations for two reservation duration distribution types; the first is the gaussian mixture used in the simulation above, the second a uniform distribution over an interval of 10, 20 or 30 seconds. The optimal time is $T=10s$.

ensure that the handoff blocking ratio is minimized and that channels are not reserved for too long. Being able to do a reservation at the proper time implies some conditions on the prediction mechanism, which have been quantified using simulations based on an abstract prediction model.

This naturally leads to a modified access control scheme that penalizes early and late reservations. This has two consequences: the network fairness is improved (selfish mobiles are discouraged) and the overall handoff blocking is minimized. This has been confirmed by simulations.

In the future, the impact of movement prediction's accuracy on other protocols should be studied. The analysis of existing prediction schemes could allow the setup of an abstract model more accurate than the one presented here.

References

- [1] N. Baghaei and R. Hunt. Review of quality of service performance in wireless LANs and 3G multimedia application services. *Elsevier; Computer Communications*, 27(17):1684–1692, November 2004.
- [2] A. Bhattacharya and S. Das. LeZi-update: an information-theoretic approach to track mobile users in PCS networks. In *Proc. of ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'99)*, Seattle, August 1999.
- [3] Y. Chen. *Soft Handover Issues in Radio Resource Management for 3G WCDMA Networks*. PhD thesis, Queen Mary, University of London, September 2003.
- [4] M.-H. Chiu and M. Bassiouni. Predictive schemes for hand-off prioritization in cellular networks based on mobile positioning. *Journal on selected areas in communications*, 18(3):510–522, March 2000.
- [5] J.-M. François, G. Leduc, and S. Martin. Learning movement patterns in mobile networks: a generic approach. In *Proc. of European Wireless 2004*, pages 128–134, Barcelona, Spain, February 2004.
- [6] S.-J. Lee, W. Su, and M. Gerla. Wireless Ad Hoc multicast routing with mobility prediction. *Mobile Networks and Applications*, 6(4):351–360, August 2001.
- [7] G. Liu and G. Maguire. A predictive mobility management algorithm for wireless mobile computing and communications. In *Proc. of international conference on Universal Personal Communications (ICUPC'95)*, Tokyo, Japan, November 1995.
- [8] T. Liu, P. Bahl, and I. Chlamtac. Mobility modeling, location tracking, and trajectory prediction in wireless ATM networks. *IEEE JSAC*, 16(6):922–936, August 1998.
- [9] A. Roy, S. Das, and A. Mistrá. Exploiting information theory for adaptive mobility and resource management in future cellular networks. *IEEE Wireless Communications*, 3(5):59–65, August 2004.
- [10] M. Shafi, S. Ogoose, and T. Hattori. *Wireless Communications in the 21st Century*. Wiley-IEEE, 2002.
- [11] W.-S. Soh and H. S. Kim. Dynamic guard bandwidth scheme for wireless broadband networks. In *Proc. of IEEE Infocom'01*, pages 572–58, Anchorage, AK, April 2001.
- [12] W.-S. Soh and H. S. Kim. QoS provisioning in cellular networks based on mobility prediction techniques. *IEEE Communications Magazine*, 41(1):86–92, January 2003.
- [13] W.-S. Soh and H. S. Kim. Dynamic bandwidth reservation in cellular networks using road topology based mobility predictions. In *Proc. of IEEE Infocom'04*, Hong Kong, Mars 2004.
- [14] V. Tolety. Load reduction in Ad Hoc networks using mobile servers. Master's thesis, Colorado School of Mines, 1999.
- [15] M. Zonoozi and P. Dassanayake. User mobility modeling and characterization of mobility patterns. *IEEE Journal on selected areas in communications*, 15(7):1239–1252, 1997.