Performance Evaluation of Mobile Wireless Networks: A New Perspective *

J. Jobin, Michalis Faloutsos, Satish K. Tripathi
Department of Computer Science and Engineering
University of California, Riverside
Riverside, CA, 92521, USA
{jobin,michalis,tripathi}@cs.ucr.edu

Abstract

We propose a methodology to simplify the analysis of wireless network simulations. Wireless simulation models are plagued by the vast parameter space found in literature; they are described by a wide variety of parameters such as user speed, cell size, etc. Consequently, performance metrics are not easy to interpret or compare across different models. We discuss how to reduce this parameter space by proposing a set of metrics that describe the network at a higher level of abstraction: our metrics depend on fewer parameters by aggregating related parameters.

The cornerstone of our approach is the *steady state utilization* metric which quantifies the inherent capacity of a network to support a particular workload. A knowledge of the steady state utilization level of a system can help to maximize resources while minimizing loss. We also discuss the *Steady State Arrival Rate* which keeps the network operating at this steady state utilization level. Furthermore, we propose a new way to evaluate the performance of a network using a metric - *effective utilization* - that captures the utilization and the loss at the same time.

Finally, as a case study, we use our methodology to evaluate the how advance reservations affect network performance. We show how, contrary to intuition, reservations can hurt performance, if we consider dropped calls as lost income.

---

[*] This material is based upon work supported by the National Science Foundation under Grant No. 9985195, DARPA award N660001-00-1-8936 and Dimi matching fund DIM00-10071.

1 Introduction

The motivation for this work is the overwhelming complexity of mobile wireless network models which makes it hard to evaluate their performance. We propose a methodology for evaluating performance in a simple, yet effective way. The methodology introduces a new metric which does not deal with the specifics of a system, thereby making it easy to compare different models without having to worry about the underlying details. We also discuss a conventional performance metric and propose an enhancement to improve its accuracy.

Network models are difficult to compare because of the differences in their implementation. Models have a large set of parameters (e.g., speed of user, size of cell, bandwidth requirement of user, etc.) and there are no standard values for these parameters. Thus, simulation results from different models can be inconsistent and hard to compare. One would like to be able to understand how different model parameters affect the results of simulations. Thus, we need a framework that makes it easy to compare models and provides a better insight into the way that protocols behave in different environments. Ideally, we would like as few parameters and metrics in the model as possible.

Most research that aims for a simple or standardized model attempts to hide as many low level details as possible. The idea of simulation abstraction has been considered in [2]. Fall [4] suggests decreasing the number of objects in a simulation by aggregating some of them. A standard framework for simulation has been suggested in [3]. None of them, however, are specific to the wireless network domain.

We propose a methodology to address the above issues. Our main contributions can be summarized as follows:

- The central idea underpinning our methodology is reducing the large number of system parameters and metrics to a few that capture the essence of the network. In this regard, we introduce the novel concept of *Steady State Utilization*, which aims to capture the inherent capacity of a network for a given user behavior.

- Using steady state utilization, we define *Steady State Arrival Rate* to define the arrival rate that will keep the system utilization at the maximum possible, without losses.

- We suggest *effective utilization* as a more insightful metric that will combine both the utilization and the loss in the system.

- Using our methodology, we show how we can obtain a better insight into the effect of bandwidth reservations on network utilization. Results show that reservations can actually hurt performance, contrary to intuition.

The rest of the paper is as follows: In Section 2 we discuss the background. In Section 3 we discuss the main ideas of our paper, followed by an analysis of some of the ideas in Section 4. We present our simulations and the results in Section 5 followed by a case study in Section 6. We conclude in Section 7.

2 Background and Model

We simulate a cellular mobile wireless network. The geographic region is divided into hexagonal cells with six neighbors each. This model is often found in previous work [19, 15, 18]. A user enters the system in any cell, provided there is enough bandwidth available. If a user attempts to enter the system, but cannot do so, due to lack of available bandwidth, it is said to be *blocked.* Once it is in the system, it spends some time in each cell which is called the cell latency [1]. The user keeps moving from one cell to another (neighboring cell) provided it finds the required bandwidth in the next cell. If it is not able to move to the next cell, it is *dropped.* Each cell has a base station which is responsible for keeping track of the bandwidth usage.

Most mobile wireless network simulation models are quite difficult to evaluate. Even when they can be evaluated, it may still be difficult to compare them against each other. This is due to the large number of parameters that these models consider. This problem has been recognized in general simulation. The solution proposed is to reduce as many parameters as possible and hide low level details. The idea of simulation abstraction has been considered in [2]. They try to abstract unnecessary details from a simulation. Fall [4] suggests decreasing the number of objects in a simulation by aggregating some of them. A standard framework for simulation has been suggested in [3].

The above mentioned abstraction concept and the object aggregation method have been applied to general simulations. We were not able to find a practical application of these concepts to the specific case of cellular wireless networks. The standard framework as sug-

---

[1] Cell latency is also referred to as *cell dwell time* in literature

gested by Bajaj [3] shares our same objectives. They assume that everyone is going to use the same framework, but we believe that this is not always feasible to assume. Therefore, having parameters that hide the system-specific details is important.

We also consider the problem of measuring network utilization. Previous work considers utilization and loss separately [6, 7, 8, 9]. This is often misleading. For example, consider two networks: One of them claims to have 90% utilization with a 5% loss rate and the other has 85% utilization with a 3% loss rate. Also assume that 10% of the dropped users were in the process of downloading files that accounted for 3% and 7% of the total bandwidth in the respective networks. (This bandwidth can be considered to be wasted since the files will need to be downloaded again.) It is hard to determine which is the better network on the basis of these figures alone. The main reasons for this difficulty are because utilization and loss are being considered separately and also because utilization does not take into account the wasted bandwidth.

3 Main Ideas

As discussed in the previous section, the main problem with existing network simulation models is the large parameter space. This makes it hard to evaluate the system performance or to compare different systems. The need for a framework that reduces the number of parameters while still retaining the essential elements of the model is apparent. In this section, we discuss some of our parameters and metrics that provide a first attempt in this direction.

3.1 Parameters

Steady State Utilization

Steady State Utilization is a parameter that gives us an insight into the load a system can support without losses. We start with a fully loaded system (maximum number of users in each cell) and let the users move around in the system. Users are permanent; they do not leave unless they get dropped. In the beginning, as they move, some of them find themselves in cells that are too crowded and hence get dropped. Initially, the drop rate is significant. Gradually, this rate slows down until we reach a point where it is practically zero. We call the utilization at this point the *steady state utilization* of the system. Thus, we define steady state utilization as the maximum sustainable utilization of the system with permanent users and zero arrivals. The steady state utilization of a system depends on the type of load, i.e., the cell latency of the users, the bandwidth required by the users, etc. The results of our experiments, as presented later, suggest that a network can be optimally utilized only if it is operating at this steady state utilization level.

Steady State Arrival Rate (SSAR)

The Steady State Arrival Rate is the arrival rate needed to keep the system functioning at its steady state utilization level. This is required because Steady State Utilization, as defined above, assumes that all users are permanent, whereas in real life they have call durations. This implies that they would terminate after some time. To keep the system functioning at the steady state utilization level, this would require more users to arrive. This arrival rate is called the *Steady State Arrival Rate.* As long as the users arrive at this rate, the system will stay at its steady state utilization level.

The SSAR is given by:

$$SSAR = Util_{ssu} * \frac{MaxUsers}{T} \tag{1}$$

where $Util_{ssu}$ is the steady state utilization of the system, $MaxUsers$ is the maximum number of users that the system can support (i.e. its capacity), and T is the average call duration. We will discuss the derivation of this formula in the next section.

Relative Arrival Rate (RAR)

The arrival rate of a system with respect to the Steady State Arrival Rate is called the *Relative Arrival Rate.* RAR is equal to 1 at the Steady State Arrival Rate. It is useful when comparing systems (as we will show in later in our experiments).

Cell-User-Bandwidth Ratio

We would like to be able to specify how many users can be accommodated by a cell on an average. This is given by the *Cell-User-Bandwidth Ratio.* This ratio is simply the cell bandwidth divided by the user bandwidth.

3.2 Metrics

Observed Utilization

One of the common metrics used for evaluating system performance is *observed utilization.* This utilization can be considered in terms of bandwidth used in the system or the number of users in the system. We will define observed utilization as the percentage of bandwidth used in the system. Observed Utilization does not take into account the loss; usually, this is specified as a separate metric.

Wasted Utilization

Observed Utilization, as we saw does not include the loss. So, we need a metric that will quantify the utilization wasted due to loss. This is called the *Wasted Utilization.* Often, a user has to be dropped in the middle of a call. The negative consequences of such an action can vary depending upon the type of the user and the activity that it was performing at that time. For instance, if a user is dropped in the middle of sending an email, the bandwidth wasted is not too high. On the other hand, if a user is dropped in the middle of transferring a huge file, then the bandwidth wasted can be substantial. This wastage is termed the *wasted utilization.* Wasted utilization has a direct impact on the utilization of the system because now we have to consider what fraction of the utilized bandwidth is being effectively used.

Effective Utilization

Effective Utilization tries to incorporate utilization and loss. As mentioned above, even though the system seems to be getting utilized at a certain level, depending upon the wasted bandwidth, the actual utilization (which accounts for the successfully used bandwidth only) is less than the observed utilization. We call this the *effective utilization.* We show how using effective utilization instead of the conventional network utilization can significantly affect our performance analysis. Effective utilization is not entirely a new concept; it is analogous to *good throughput* or *goodput* often found in literature [17].

3.3 Simulation Abstraction

We now try to abstract some of the common simulation parameters and metrics using our newly defined ones. Figure 1 shows some of this abstraction.

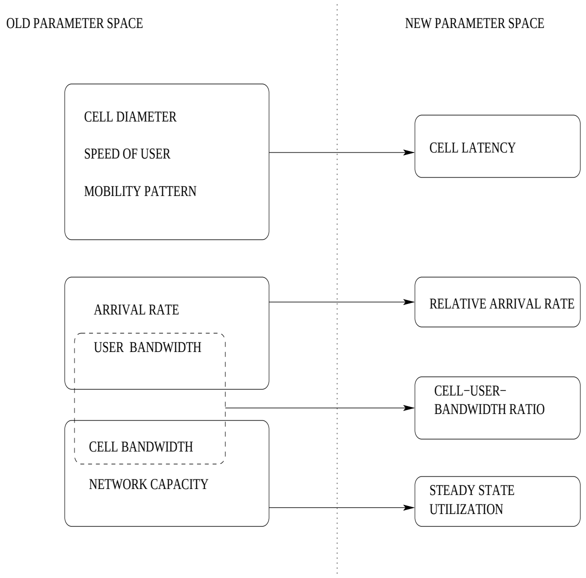

Figure 1: **Reduction of parameter space**

The left column shows some of the parameters and metrics that are usually considered in performance evaluation. The right column shows our equivalent parameters and metrics. For instance, cell latency accounts for the time spent by a user in a cell which is determined by the diameter of the cell, the speed of the user, and the mobility pattern of the user. Our steady state utilization concept subsumes both the cell capacity and the network capacity. It also gives us the Steady State Arrival Rate which in turn accounts for the arrival rate and the bandwidth requirements of the arriving users.

Thus, we can see that it becomes easier to evaluate performance by using the above parameters and metrics which try to hide details.

4 Analysis

In this section we analyze steady state utilization and Steady State Arrival Rate in greater detail using operational laws. Table 1 shows the notation used in our derivations. We derive

N_{users}	Number of users in the system
$N_{arrivals}$	Number of users that arrived in the system
$loss$	Percentage of arrivals that were lost due to blocks or drops
$MaxUsers$	Maximum number of users that the system can support
$Util$	Utilization
$Util_{ssu}$	Steady State Utilization
λ	Arrival Rate
$SSAR$	Steady State Arrival Rate
T	Time

Table 1: **Notations**

expressions for steady state utilization and Steady State Arrival Rate, and establish a relation between the two.

4.1 Steady State Arrival Rate

Let the users be homogenous in terms of bandwidth requirements. Assume that the call duration of a user on an average is T. If we observe a system with arrival rate λ for time T, the number of users that arrived in the system can be given by:

$$N_{arrivals} = \lambda * T \tag{2}$$

Some of these users that arrive are lost due to blocks and drops. Assuming that $loss$ is the fraction of users lost, the number of users that exist in the system can now be expressed as:

$$N_{users} = (1 - loss) * \lambda * T \tag{3}$$

Dividing both sides of (3) by the maximum number of users, $MaxUsers$, we get:

$$\frac{N_{users}}{MaxUsers} = (1 - loss) * \frac{\lambda * T}{MaxUsers} \tag{4}$$

The left side of (4) is the system utilization:

$$Util = (1 - loss) * \frac{\lambda * T}{MaxUsers} \tag{5}$$

Solving for λ, we get:

$$\lambda = \frac{Util}{1 - loss} * \frac{MaxUsers}{T} \tag{6}$$

When the system is operating at its steady state utilization, loss is zero. Substituting this in (6) we find the arrival rate corresponding to steady state utilization, which we defined as Steady State Arrival Rate (SSAR):

$$SSAR = Util_{ssu} * \frac{MaxUsers}{T} \tag{7}$$

Remember that steady state utilization was defined based on the assumption that there were no new arrivals or departures in the system. Now, if we assume that there are departures, we will need arrivals to offset them and keep the system operating at the steady state utilization. This arrival rate is the $SSAR$.

5 Experimental Results

We assume a wireless cellular network with homogeneous cells. The users are homogeneous with the same bandwidth requirement and same mean value of cell latency. For changing the cell capacity (in terms of number of users), we change the ratio of the cell bandwidth to the user bandwidth. We perform the experiments with a network of 18 cells, with each cell having six neighboring cells. The simulation programs are written in the C version of CSIM. First, we wanted to consider how different parameters affect the steady state utilization. Second, we wanted to see what the Steady State Arrival Rate means and how we could use it.

5.1 Parameters That Affect Steady State Utilization

First, we conduct some experiments to see the the effect of changing the Cell-User-Bandwidth ratio on the steady state utilization. Then, we experimented to study the effect of cell latency on steady state utilization.

The effect of Cell-User-Bandwidh ratio on steady state utilization:
In Figure 2(a), we show three plots. The lowest plot corresponds to the case where we have a single class of users and the Cell-User-Bandwidth ratio is 20. As seen, the steady state utilization level is about 37%. The topmost plot is for the case where we increase this ratio to

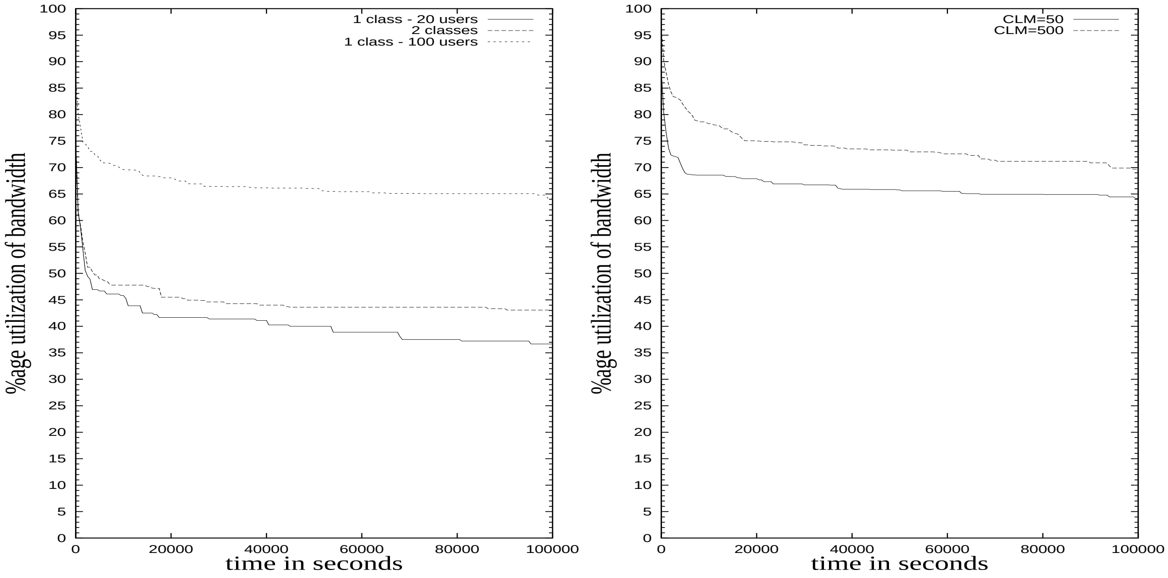

(a) steady state utilization vs. cell capacity and classes of users.

(b) steady state utilization vs. cell latency

Figure 2: The change in the steady state utilization with a change in the cell capacity, the number of classes, and cell latency.

100. This leads to an increase in the steady state utilization to about 65%. Thus, the graph seems to indicate that the steady state utilization is more when users have smaller bandwidth requirements. This is because if a higher bandwidth user is dropped, it will affect the effective utilization more. For the extreme case, consider a network with two cells and one user in each cell with each occupying 80% of the cell bandwidth. In this case, as soon as they try to move, one of them will be dropped, and immediately the utilization drops drastically. Thus, the experiments suggest that **finer granularity of user bandwidth makes the system operate smoother with fewer losses.**

Next, we consider two classes of users. The Cell-User-Bandwidth ratio of one class is 100, and that of the other class is 500. The users are equally divided into the two classes. The result is seen in the middle plot in Figure 2(a), where the steady state utilization is now about 43%. This plot is meant to illustrate the idea of simulation abstraction. If we consider users in terms of bandwidth usage only, we do not have to worry about the classes of users.

The effect of cell latency on steady state utilization:

Here, we consider the effect of cell latency on the steady state utilization level. The first task was to find the steady state utilization of a given network. Figure 2(b) show the case for two different mean values of cell latency - 50 and 500 seconds. As seen in Figure 2(b), when the mean cell latency is 50 seconds, the steady state utilization is about 64%, and when the mean cell latency mean is 500 seconds, the steady state utilization is about 70%. Intuitively, ss the cell latency increases, there would be fewer drops and blocks, and hence the steady state utilization should increase. Indeed, consider the extreme case where cell latency is close to infinity. The users are more or less stationary and the steady state utilization will be close to 100%.

5.2 What Steady State Arrival Rate means

Once the steady state utilization is determined, we incorporate departures in the system. This is done by specifying the call duration. Then, we calculate the arrival rate corresponding to the steady state utilization. This is the $SSAR$ as defined earlier. When Relative Arrival Rate(RAR) equals 1, we have the $SSAR$ and utilization is the steady state utilization. We varied the RAR, thereby subjecting the system to different loads corresponding to different values of the arrival rates. We consider loads for 0.5, 0.8, 1.0, 1.2, 1.5, 2.0, 2.5, and 3.0 times the $SSAR$ for both values of mean cell latency - 50 and 500 seconds.

Figures 3(a) and 3(b) show the comparison between the observed utilization and the effective utilization when the arrival rate is varied. We see that if RAR is below 1, there is no wasted bandwidth. However, the utilization is below the steady state utilization. Thus, even though we are not experiencing any loss, we are not taking full advantage of the available resources.

This is true until $RAR = 1$ (i.e., $\lambda = SSAR$) at which point the loss starts setting in. As RAR increases beyond 1, the effective utilization starts dropping and the loss due to the drops also increases. Thus, we see that for $RAR = 1$, we have the maximum effective utilization without loss. In other words, the maximum load rate that the system can support without loss is the $SSAR$. This suggests that we should operate a network only at a level which is near the $SSAR$. If we try to go beyond this $SSAR$, we are bound to incur more loss and wasted bandwidth.

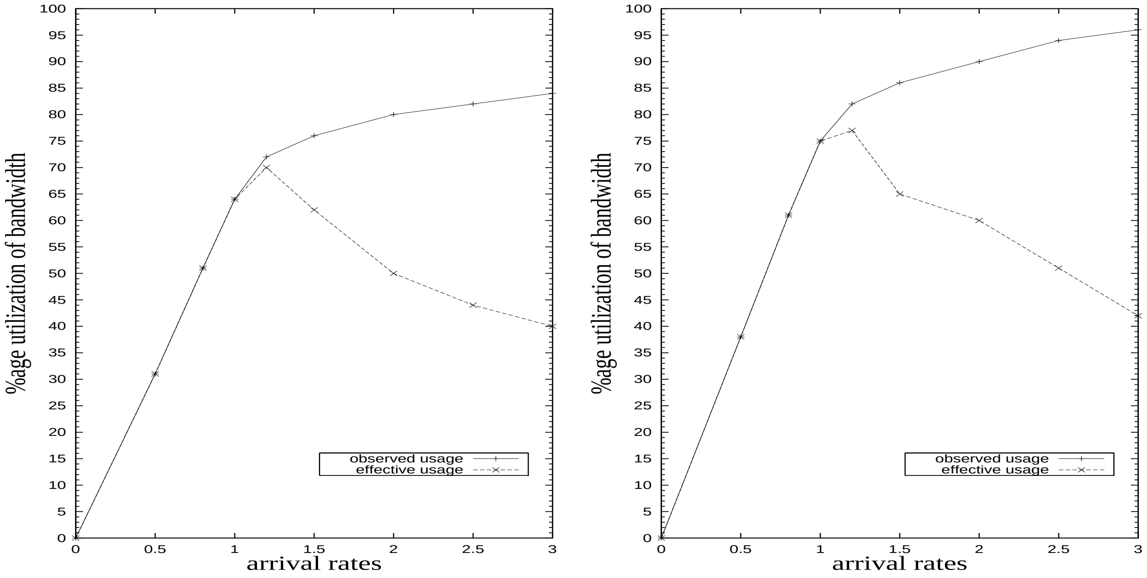

(a) with MCL = 50 seconds

(b) with MCL = 500 seconds

Figure 3: Observed and effective utilization vs. relative arrival rate

5.3 Using Our Methodology

It is straightforward to find the steady state utilization of any network. Once we have this, we can find the Steady State Arrival Rate (SSAR). Next, we can vary the Relative Arrival Rate (RAR) and see how the effective utilization is affected. So, given a network and its load characteristics, one can say whether it is being underutilized or overutilized. Moreover, one can also determine what the arrival rate should be in order to achieve achieve steady state utilization.

The methodology can also be applied to the comparison of two different networks. For instance, consider two networks with different characterstics and a protocol that needs to be evaluated on both. Assume that network A has a steady state utilization of 80% and an effective utilization of 73%. Network B has a steady state utilization of 78% and an effective utilization of 76%. We can say that network A is not being as efficiently utilized as network B.

We should mention that we have tried to keep things simple. In this initial study, we wanted to highlight the use of our approach. Indeed, there are other parameters that affect the performance, such as different classes of users, variable cell latencies, variable call durations, etc. We assume that all users are homogeneous in terms of bandwidth. However, as we showed in our results, even if we have multiple classes of users, it is possible to find the steady state utilization and then use it for comparison purposes. Another assumption is that the mean value of the cell latency is the same for all users. We plan to experiment with changing this in the future. For this study, we have kept our simulation as simple as possible so that we could illustrate our main points.

6 Evaluating the Performance of Advance Reservations

In this section, we apply the steady state utilization metric to a specific problem - to study the benefits of advance reservations in wireless networks. Reservations have been studied as a way of improving the quality of service ([14], [18], [15]). In brief, the idea is to make reservations in some of the neighboring cells that a user is likely to move into ahead of time. Statistical data like a mobility profile ([14], [11]) have been proposed to make the advance reservations as accurate as possible. We study how reservations impact the effective utilization.

We assume a cellular mobile wireless network for the simulation. For simplicity we assume that all users have the same bandwidth requirements and the same cell latency. The mean cell latency is 50 seconds. Each cell can support up to a maximum of 100 users.

For reservations, the next cell has to be predicted in advance, before the user moves into it. We consider two cases. First, we assume we are always able to predict correctly the next cell that the user will move into; this is called *perfect reservation*. Here, we need to reserve only one cell for each user. Of course, perfect reservation is an ideal scenario but it provides us with a good benchmark. Second, we make reservations in two cells. We call this a 2-cell reservation scheme. This is less idealistic and more realistic. Note that the number of cells reserved is a factor of the accuracy of the statistical data and the user behaviour [20].

6.1 Calculating the Steady State Arrival Rate

For optimal performance, the system load should be less than the Steady State Arrival Rate. First, we find the steady state utilization of the system with perfect reservation and the corresponding arrival rate, SSAR. At this load, we have no drops. Then we

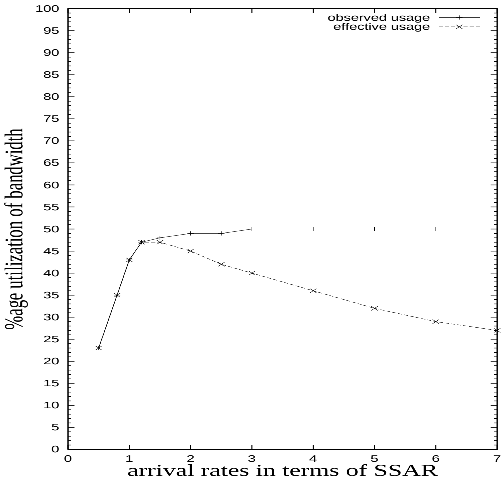

Figure 4: **Variation of utilization with the arrival rate**

subject the system to different values of the arrival rate, λ. Figure 4 shows the case when the arrival rate λ is varied from $0.5 * SSAR$ to $7.0 * SSAR$. The x-axis is the arrival rate and the y-axis is the percentage utilization of the total bandwidth of the system. The graph shows that the observed utilization increases with the load λ until it reaches the 50% utilization level. Intuitively, this can be explained by the fact that in perfect reservation, each user occupies a slot in two cells each - one in which it is currently residing, and the other in which it has made the reservation.

On the other hand, when we go beyond the SSAR, the effective utilization decreases as the load increases. This is because as we start having more users in the system, it becomes more likely that too many users might find themselves in a cell with insufficient bandwidth, and some of them will get dropped.

6.2 Using effective utilization and Steady State Arrival Rate to Study Reservations

Reservations hurt the performance of the network. Reservations reduce the utilization of the system without offering any substantial advantage. Figure 5 shows a comparison between reservation and no reservation. The x-axis is the arrival rate and the y-axis is the effective utilization as a percentage of the total bandwidth of the system. Under low and medium loads, no reservation is better than having reservations, even if they are perfect. Not having perfect reservations only degrades the performance further. The 2-cell reservation scheme exhibits much lower effective utilization as compared to no reservations.

A point to be noted is that eventually, perfect reservation does start outperforming no reservation, but this is only when the load gets high (beyond $4 * SSAR$.) However, at high loads, the effective utilization starts decreasing due to more drops (Figure 4). This implies that for good performance, reservations do not seem to help.

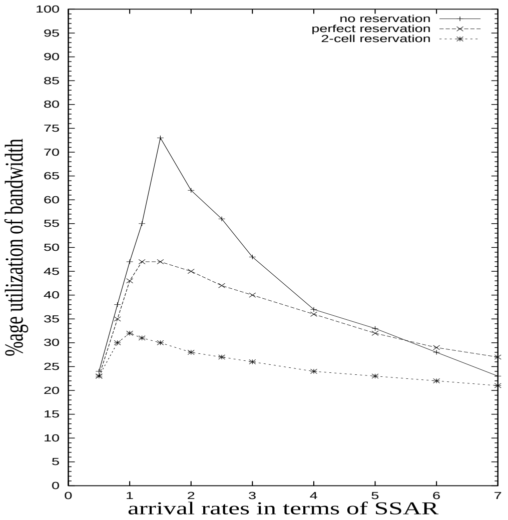

Figure 5: **Effective utilization with and without reservation**

Thus, our results indicate that a system might actually perform better without reservation. Nevertheless, we should mention here that these are preliminary results. Our results are based on our assumption that dropped calls do waste bandwidth. We plan to do more exhaustive experiments before we can state this conclusively.

7 Conclusions

We proposed a performance evaluation methodology for wireless network simulation models. Our methodology reduces the number of parameters and introduces new metrics that facilitate the easier evaluation and comparison of system performance across different models.

We introduced the novel concept of steady state utilization which captures the inherent capacity of a network for a given workload. Using steady state utilization, we defined *Steady State Arrival Rate* as the arrival rate that will keep the system utilization at the maximum possible level, without losses. Moreover, we proposed *effective utilization* as a more insightful metric that combines both the utilization and the loss in the system.

Using our methodology, we show how we can obtain a better insight about the effect of reservations on network utilization. Results show that despite common belief, reservations can actually hurt performance.

We plan to do more experiments that help us understand the concept of steady state utilization thoroughly. We intend to experiment varying the other parameters that we mentioned, such as cell latencies, call durations, multiple classes of users, etc. The idea of having no reservation is quite intriguing; however, there is a lot of work to be done in this direction before we have solid conclusions.

References

[1] Piyush Gupta, P.R.Kumar. *The Capacity of Wireless Networks.* 38th IEEE Transactions on Information Theory, pp.388-404, vol.IT-46. no.2. March 2000

[2] Polly Huang, Deborah Estrin, John Heidemann *Enabling Large-Scale Simulations: Selective Abstraction Approach to the Study of Multicast Protocols* Proceedings of the Sixth International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems (MASCOTS '98), Montreal, Canada. July 1998.

[3] Sandeep Bajaj, et al. *Improving Simulation for Network Research* USC Computer Science Department Techical Report, March 1999.

[4] Thomas C. Fall *A Framework for the Simulation Experimentation Process* Proceedings of the 1997 Winter Simulation Conference

[5] Joel Short, Rajive Bagrodia, Leonard Kleinrock. *Mobile Wireless Network System Simulation* ACM Mobicom '95. pp.195-209

[6] Anup K. Talukdar, B.R. Badrinath, Arup Acharya. *Intergrated Services Packet Networks with Mobile Hosts: Architecture and Performance.* Wireless Networks. vol.5, No.2, March 1999, pp.111-124.

[7] Jay R. Moorman, John W. Lockwood, Sung-Mo Kang. *Real-time prioritized Call Admission Control in a Base Station Scheduler.* Proceedings of the third ACM international workshop on wireless mobile multimedia 2000. pp.28-37.

[8] Bo Li, Chuang Lin, Samuel T. Chanson. *Analysis of a hybrid cutoff priority scheme for multiple classes of traffic in multimedia wireless networks.* Wireless Networks Vol.4, 1998, pp279-290

[9] Youssef Iraqi, Raouf Boutaba. *A Novel Distributed Call Admission Control for Wireless Mobile Multimedia Networks.* Proceedings of the third ACM international workshop on wireless mobile multimedia 2000. pp.21-27.

[10] Samir R. Das, Robert Castaneda, Jiangtao Yan, Rimli Sengupta. *Comparative Performance Evaluation of Routing Protocols for Mobile, Ad Hoc Networks.* Proceedings of the 7th International Conference on Computer Communications and Networks (IC3N), Lafayette, LA, October, 1998, pp.153-161.

[11] Fei Yu, Victor C. M. Leung. *Mobility-Based Predictive Call Admission Control and Bandwidth Reservation in Wireless Cellular Networks.* Proceedings of the IEEE Infocom 2001 Conference, Anchorage, Alaska. April, 2001.

[12] S.Bandhopadhyay, Krishna Paul. *Evaluating the Performance of Mobile Agent-Based Message Communication among Mobile Hosts in Large Ad Hoc Wireless Network.* Proceedings of the 2nd ACM International Workshop on modeling, analysis, and simulation of wireless and mobile systems. Aug.1999

[13] Xiaoyan Hong, Mario Gerla, Guangyu Pei, Ching-Chuan Chiang. *A Group Mobility Model for Ad Hoc Wireless Networks.* Proceedings of the 2nd ACM International Workshop on modeling, analysis, and simulation of wireless and mobile systems. Aug.1999

[14] Vaduvur Bharghavan, Jayanth P. Mysore. *Profile-Based Next-Cell Prediction in Indoor Wireless LANs.* IEEE Singapore International Conference on Networking. 1997.

[15] Suresh Singh. *Quality of Service Guarantees in Mobile Computing.* Journal of Computer Communications. Vol.19, 1996. pp.359-371.

[16] Rohit Ghai, Suresh Singh. *An Architecture and Communication Protocol for Picocellular Networks.* IEEE Personal Communications Magazine. Vol.1(3), 1994. pp.36-46.

[17] Hari Balakrishnan, Venkata N. Padmanabhan, Srinivasan Seshan, Randy Katz *A Comparison of Mechanisms for Improving TCP Performance over Wireless Links* IEEE/ACM Transactions on Networking. Vol.5, No.6, Dec 1997. pp.756-769

[18] A.R.Aljadhai and T.F.Znati. *A Framework for Call Admission Control and QoS Support in Wireless Environments.* Proceedings of IEEE INFOCOM '99. pp.1019-1026

[19] Kevin Brown, Suresh Singh. *A Network Architecture for Mobile Computing.* IEEE INFOCOMM '96

[20] J.Jobin, Satish K. Tripathi, Swapna Gokhale, Michalis Faloutsos. *Using Statistical Data for Reliable Mobile Communications* Accepted for publication