Simplifying the Analysis of Wireless Cellular Network Simulation

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

Abstract

We propose a systematic approach for simplifying the analysis of wireless cellular network simulations. Wireless network simulation models have a variety of parameters and there is no consensus on what values are reasonable for these parameters. This variety in the parameter space combined with the lack of standard values makes it difficult to analyze simulation results or compare results from different simulation models. Our approach tries to address this problem by evaluating how a network model performs with respect to its own optimal operating point. We introduce the novel concept of *steady state utilization* to define this optimal point. Steady steady utilization aims to capture the inherent capacity of a network in a system-independent manner. Using this concept, we show how we can analyze results and compare across different simulations. Simulations suggest that steady state utilization is a good indicator of the optimal operating point.

Keywords
steady state utilization, network capacity

1.INTRODUCTION

We propose a framework for simplifying the analysis of results of wireless cellular network simulations. Wireless simulation models suffer from an abundance of parameters. Some of the commonly found ones in literature ([1, 2, 3, 4, 5]) are cell latency, network size, shape of the cell, etc. Unfortunately, there are no standard accepted values for them. Consequently, varying assumptions are frequently made for some of the values ([1, 2, 3, 4]). This makes the analysis and comparison of simulation results a complex task.

The problem can be illustrated by the following examples. Consider a network that is operating at 85% utilization with 1% loss. Is it being efficiently used? As another example, assume that we wish to evaluate a protocol for two networks - A and B. A has an arrival rate of 30 users/second, 80% utilization, and 5% loss, whereas B has an arrival rate of 40 users/second, 70% utilization, and 2% loss. Which network is being better utilized? As we can see, there is no simple answer to these questions.

The problem is that there are too many variables involved. Without knowing the size of the network, or the traffic load, or a host of other parameters, it is not possible to decide which network is being better utilized. We need some method to eliminate this dependence on parameters. Our approach tries to achieve this system-independence by evaluating a network model in terms of how it is performing versus what it is capable of. In other words, we try to analyze its performance independent of other models by assessing its performance relative to its own optimal performance. We choose network utilization as the measure of performance to facilitate this assessment and compare it against an optimal utilization - the *steady state utilization.* We define steady state utilization to be the maximum utilization of a network without loss, for a given user behavior.

The rest of the paper is as follows: In Section 1, we discuss the background. Section 2 discusses the main ideas. We show the experimental results in Section 3 and conclude in Section 4. An earlier version of this work which contains preliminary ideas and results appears in [10].

2.BACKGROUND

We use the model for wireless networks commonly found in literature ([3, 4]). A geographic region is divided into cells each of which has a fixed number of neighbors. Each cell has a base station for managing bandwidth usage among the users. Users enter the system in a random cell and move from one cell to the next adjacent cell until they are dropped. We use a wrap-around network: when users reach the edge of the network, they continue on to the other side of the network.

Most wireless simulations have a large number of parameters. There has been research in the area of the general simulation environment. Huang *et al.*[7] propose simulation abstraction to reduce unnecessary details. Fall[8] discusses decreasing the number of simulation objects by aggregating some of them. Bajaj *et al.* [9] propose a standard framework to facilitate easy comparison across models and suggest that everyone use the same framework. However, this is not always convenient; indeed, to the best of our knowledge, there still exists no single framework that is universally used. Thus, it becomes clear that it is important to have parameters that are few in number and hide system-specific details.

Our approach is analogous to the MERIT framework for comparing protocols as described in [6]. To quote, MERIT aims to *"rank any protocol by comparing it to a theoretical optimum rather than to a competing protocol."* The MERIT approach is to provide a benchmark relative to which all protocols can be compared. Analogously, our benchmark metric is the steady state utilization, which defines the maximum potential capacity of a network.

3.OUR FRAMEWORK FOR ANALYZING RESULTS

At the core of our approach lies the concept of *steady state utilization.* We use it to capture the essence of the inherent network capacity. Using it as a comparison benchmark, we will show how it simplifies the analysis of simulation results.

Steady State Utilization:

Steady state utilization is an indicator of the maximum load that a system can support without loss. We start with a fully loaded system where all cells have the maximum number of users. It is a closed system and no more users enter the system after this initial stage. We then let the users move around in the system. We assume that all users are permanent and do not leave the system unless they are dropped. The movement of users in the system leads to a situation where sometimes a cell has too many users and hence some of them get dropped. This drop rate is quite drastic in the beginning; however, it slows down gradually over time and becomes practically zero. We call the utilization at this point the steady state utilization.

Steady State Arrival Rate (SSAR):

Steady state arrival rate is the arrival rate that will keep the system operating at its steady state utilization. As defined earlier, steady state utilization by itself assumes that there are no arrivals and departures and that all users are permanent. This is not realistic. So, in a system where the user lifetimes are finite, there is a corresponding arrival rate which is necessary to keep the system at its steady state utilization. We define this as the SSAR.

Effective utilization

Effective utilization tries to provide a better understanding of the actual utilization by combining utilization and loss. It does not consider wasted utilization and takes into account only the successfully used bandwidth. Therefore, it can never be more than observed utilization. Effective utilization is analogous to *good throughput* or *goodput* which is often found in literature.

Note that effective utilization and steady state utilization are distinct. Effective utilization can have a range of values over the period of a simulation run. Steady state utilization for a given network and user behavior will have only one value.

Steady state utilization as an optimum

We define relative arrival rate (RAR) as the ratio of the actual arrival rate (λ) over $SSAR$. RAR is equal to 1 when the arrival rate is equal to $SSAR$. We vary the RAR by subjecting the system to different arrival rates. Figure 1 shows the utilization when RAR is varied. Observed utilization increases with increased load (i.e., larger value of RAR). However, the effective utilization increases only up to the point where RAR is just slightly more than 1. Beyond that, the effective utilization starts dropping because loss sets in. In other words, the system seems to operate at its maximum potential capacity without loss when the arrival rate is the SSAR. Thus, $SSAR$ is the maximum load that a system can support without experiencing loss, and the utilization at this point is the steady state utilization. Our approach is centered on evaluating how close to this optimum a network can achieve.

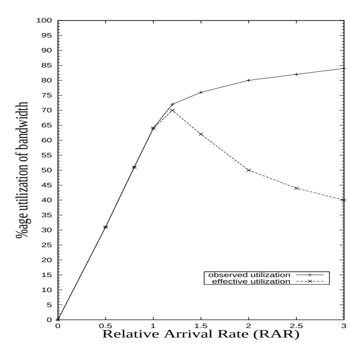

Figure 1: Observed and effective utilization vs. relative arrival rate

Using our approach for simulation analysis

Our approach can be used in the following manner:

(1).Given a network and its utilization, find its steady state utilization. (2).Compare the given utilization to its steady state utilization and see whether the network is being efficiently utilized.

Using our approach, it is quite simple to see how efficiently a network is being utilized. Thus, the steady state utilization is a reference point for understanding the utilization of a network.

This approach can also be applied to the comparison of two networks. For instance, assume that we have to evaluate a given protocol on two different networks - A and B. Assume that A has a steady state utilization of 90% and an observed utilization of 73%. B has a steady state utilization of 82% and observed utilization of 78%. We can say that A is not being as efficiently utilized as B. The important thing here is that we do not have to deal with the lower level details of the model such as size of the cell, shape of the cell, user bandwidth, direction of movement, etc.

In [10], this approach is applied to the study of advance reservations in wireless networks to provide Quality of Service (QoS). It makes use of the steady state utilization to show how reservations are not always useful and can sometimes degrade performance leading to lower QoS.

Having shown that steady state utilization can give us a better understanding of the capacity of a network, we now study how various parameters affect it in an effort to understand its robustness and establish how much of an inherent nature it possesses.

4.EXPERIMENTAL RESULTS

We conduct experiments to study the behavior of steady state utilization with the variation of different parameters. We consider a wrap-around network to avoid edge effects. We use the most common mobility model found in literature, the random walk model [11]. Here, a user moves from one cell to another randomly and is independent of all other users. Later, we study two other mobility models.

A:Cell latency

Cell latency is the amount of time that a user spends in a cell. The users are homogeneous and have the same cell latency for a particular simulation run.

Figure 2 shows how utilization varies with time. Time is plotted on the x-axis in terms of the mean cell latency. The graph shows seven plots corresponding to different cell latency values. As seen, cell latency does not seem to affect the value of steady state utilization for a network; the range of values for the utilization is less than 1%. Indeed,

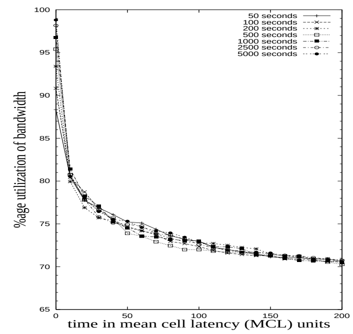

Figure 2: Variation of utilization with cell latency : 48 hexagonal cells

steady state utilization is achieved after a certain number of time units of cell latency. Thus, it is dependent on the number of hops a user makes on an average, as opposed to the actual cell latency. Intuitively this means that after a certain number of hops, the number of users settles down to a point where there are no more drops. This indicates an inherent nature of the network to support a particular number of users. An implication is that if we use the steady state utilization in analyzing simulation results, we do not need to be concerned about the cell latency.

We also studied the effect of cell latency in a hexagon cell network of 18 cells and 80 cells and a square cell network with three different sizes. The range of values for utilization for all cases was less than 2% except in the case of a square cell network with 16 cells where it was less than 3%. This indicates that the independence of utilization from cell latency holds even if the shape of the cell changes. This strengthens our statement that cell latency does not affect steady state utilization.

B: Network size

First, we consider a hexagon cell network. Figure 3 shows that there is negligible effect of network size on the steady state utilization. The utilization in the 18 cell network is only marginally better - about 0.5%. This implies that if all other factors are constant, there is no dependence on the size.

Similar results were obtained for a square cell network. We experiment with three sizes (16,49, and 81); the utilization seems to be unaffected by size, except that the 16 cell network is only marginally better - about 0.5%.

These results indicate that size is inconsequential except when it is too small. Intuitively, we believe this is because the remainder of the parameters

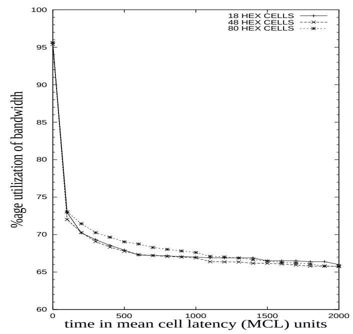

Figure 3: Variation of utilization with network size: 18, 48, and 80 hexagonal cells

are the same, the users are randomly distributed uniformly throughout the cells, and the cells themselves are all the same. In a smaller network, what happens in one cell could have a more pronounced effect on what happens in other cells, i.e., the correlation between cells is higher and hence, there is a larger margin of error. Indeed, the 18 and 16 cell networks seem to differ a bit from the others.

C: Shape of the cell

We experiment with two kinds of cells - hexagon and square. Hexagon cells are the most commonly found cells in literature ([3, 4, 1]). We consider three different network sizes for both types of networks. As we saw earlier, the network size does not seem to affect the steady state utilization. So, a hexagon cell network can be safely compared to a square cell network of equivalent size without being concerned about any side-effects.

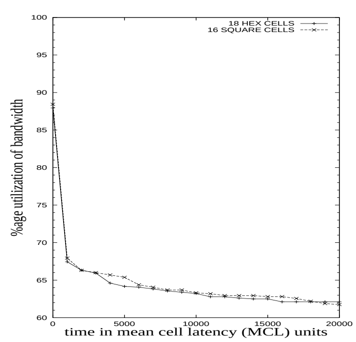

Figure 4: Effect of shape on utilization: 18 hexagonal vs 16 square cells

As seen in Figure 4, a hexagon cell network of 18 cells has a similar utilization as a square cell network of 16 cells. This suggests that the shape of the individual cells in a network does not play an important role in the final results. We obtained similar results when we compared a 48 hexagon cell network to a 49 square cell network and an 80 hexagon cell network to an 81 square cell network.

These figures indicate that the shape of the cell does not affect the steady state utilization. In other words, a hexagon cell network is equivalent to a square cell network.

D: Edge effects

To study edge effects, we experiment with a *bounce-back* network. Here, users moving into cells that lie on the edge of the network do not wrap around to the other side; instead, they bounce back from the edge and keep continuing. Our results show that utilization increases with the size of the network. This is because a bounce-back network introduces edge effects which result in users being dropped. However, the number of cells at the edge of the network as a percentage of the total number of cells is higher in a smaller network as compared to a larger network. Hence more users get dropped in a smaller network because of edge effects and therefore, utilization is lower.

4.2. Mobility models

We study three different mobility models:
Random walk: The users move from one cell to another in a completely random manner. All users are independent of each other.
Directed movement: The users move in a specific direction. The next cell of a user is in a direction that is within a specific angle [$-\alpha$,$+\alpha$] of the previous cell.
Directed movement with stoppage: This is similar to the directed movement model. In addition, at certain points along its journey, the user stops for a while before continuing.

The experiments described so far are for the random walk model. Now, we will consider the other two models. For lack of space, we do not show the graphical results for these two models.

Cell latency:
For both the directed movement models (with and without stoppage) results show that the utilization increases with an increase in the cell latency. This is in contrast to the random walk model where the cell latency did not seem to have any effect. We are in the process of investigating this.
Network size:
For both mobility models, our results show that the network size does not affect utilization. Also, we noticed that smaller networks (16 and 18 cells)

seem to differ by a tiny fraction. This could be because of the stronger correlation between cells in smaller networks.

Shape of the cell:
Our results show that the shape of the cells does not affect utilization in either model. We experimented with hexagon and square cells. This is similar to the result obtained in the random walk model.

Discussion of Results:
Our experiments show that as far as steady state utilization is concerned: (a)Square and hexagon cell networks are equivalent. (b)Network size also does not seem to affect simulations as long as it is not too small. This is applicable in the case of a wrap-around network. In a bounce-back network this is no longer true because of edge effects. (c)Mobility models affect the way in which cell latency impacts steady state utilization. In a random walk model, cell latency appears to have no effect on steady state utilization. However, in the directed movement models (with and without stoppage) steady state utilization changes with cell latency.

Thus, steady state utilization is a robust metric that seems unaffected by a host of parameters. Especially for a random walk model, it seems to be independent of cell latency, shape of the cell, and network size. This can simplify the analysis of simulation results. In the directed movement models (with and without stoppage), it is affected by cell latency but still remains independent of shape of the cell and network size. We believe this makes a strong case for steady state utilization to be used for gauging the inherent capacity of a network and shows that it is a good indicator of the maximum potential of a network.

5.CONCLUSIONS

We discussed how the large variety of parameters and the lack of standard values for them makes the task of evaluation of wireless simulation models difficult. As a solution, we proposed a systematic approach centered on our steady state utilization concept that could be used for this purpose. We showed how steady state utilization is a powerful metric that is apparently unaffected by most parameters, and therefore, can be used as a relative benchmark against which simulation models can be compared. This lends strength to our claim that it does indeed capture the inherent capacity of a network and that it can be used in the analysis and comparison of results from different simulation models.

We plan to experiment with steady state utilization in further detail. We still do not have an explanation as to why it is not affected by the shape of the cell. We are also interested in investigating how different mobility models impact steady state utilization.

References

[1] Youssef Iraqi and Raouf Boutaba, "A novel distributed call admission control for wireless mobile multimedia networks," *Proceedings of the third ACM international workshop on wireless mobile multimedia*, pp. 21–27, 2000.

[2] Vaduvur Bharghavan and Jayanth P. Mysore, "Profile-based next-cell prediction in indoor wireless lans," *IEEE Singapore International Conference on Networking*, 1997.

[3] 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, 1999.

[4] Suresh Singh, "Quality of service guarantees in mobile computing," *Journal of Computer Communications*, vol. 19, pp. 359–371, 1996.

[5] Piyush Gupta and P.R. Kumar, "The capacity of wireless networks," *IEEE Trans. Information Theory*, vol. IT-46, no. 2, pp. 388–404, March 2000.

[6] Andras Farago and Violet R. Syrotiuk, "Merit: A unified framework for routing protocol assessment in mobile ad hoc networks," *ACM MOBICOM*, July 2001.

[7] Polly Huang, Deborah Estrin, and John Heidemann, "Enabling large-scale simulations: Selective abstraction approach to the study of multicast protocols," *Proceedings of MASCOTS 1998. Montreal, Canada*, July 1998.

[8] Thomas C. Fall, "A framework for the simulation experimentation process," *Proceedings of the 1997 Winter Simulation Conference*, 1997.

[9] Sandeep Bajaj, et al. "Improving simulation for network research," *USC Computer Science Department Technical Report*, March 1999.

[10] J.Jobin, Michalis Faloutsos, and Satish K.Tripathi, "Performance evaluation of mobile wireless networks: A new perspective," *ACM MSWIM*, July 2001.

[11] Christian Bettstetter, "Smooth is better than hard: A random mobility model for simulation of wireless networks," *ACM MSWIM*, July 2001.