



CHALMERS

Chalmers Publication Library

Predictive Resource Allocation Evaluation with Real Channel Measurements

This document has been downloaded from Chalmers Publication Library (CPL). It is the author's version of a work that was accepted for publication in:

IEEE International Conference on Communications

Citation for the published paper:

Ahmad, S. ; Reinham, R. ; Muppisetty, L. et al. (2017) "Predictive Resource Allocation Evaluation with Real Channel Measurements". IEEE International Conference on Communications

Downloaded from: <http://publications.lib.chalmers.se/publication/249158>

Notice: Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source. Please note that access to the published version might require a subscription.

Chalmers Publication Library (CPL) offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all types of publications: articles, dissertations, licentiate theses, masters theses, conference papers, reports etc. Since 2006 it is the official tool for Chalmers official publication statistics. To ensure that Chalmers research results are disseminated as widely as possible, an Open Access Policy has been adopted. The CPL service is administrated and maintained by Chalmers Library.

(article starts on next page)

Predictive Resource Allocation Evaluation with Real Channel Measurements

Suhail Ahmad, Rikard Reinhausen, L. Srikar Muppirisetty, Henk Wymeersch
Department of Signals and Systems, Chalmers University of Technology, Gothenburg Sweden
email: {suhaila, rrikard}@student.chalmers.se, {srikar.muppirisetty,henkw}@chalmers.se

Abstract—Mobile services, especially video streaming, has seen a rapid usage increase in recent years. Base stations (BSs) need to employ smarter and efficient resource allocation strategies to maintain high quality of service (QoS) to users at all the times. Predictive resource allocation (PRA), is one such novel scheme, in which BSs seek to anticipate the user demands and offer service to users in advance. As a result, the QoS can be improved, network load can be distributed over time, while at the same time offering efficient utilization of BS power. In location-aware PRA, the BS exploits location information of the users to predict the channel expected variations and adapt the BS resources accordingly. We evaluate PRA strategies based on an empirical study of the radio channel variation from measured location-aided channel radio maps using a smart phone. We observed that gains offered by the PRA scheme are highly dependent on user mobility patterns.

I. INTRODUCTION

Technological advancements in smart phone development, combined with ease of use and low cost, have made mobile devices increasingly prevalent. This has resulted in a rapid growth of mobile data traffic. Statistics of mobile data usage indicate that engaging activities like video and photo form the bulk of this traffic. The Cisco Visual Networking index [1] and Ericsson Mobility Report 2016 [2] clearly show the rising popularity of video among mobile users. This rise is mainly due to the extensive video sharing and video content generating platforms. To provide seamless video playback at all times puts hard requirements on bandwidth and coverage. In a scenario when a mobile user is streaming a video on the go, much variation can be seen in the quality of the video. The video may also stall at times during the stream. This is due to two main factors affecting the Quality of Service (QoS): the coverage and the network load.

Network load can be reduced on the network layer and through video compression to give users a high-quality video experience. On the lower layers, prediction of future user demands and future channel qualities can further even out network loads [3]–[5]. The reasoning behind doing this is the predictability of user movement and data usage. According to [6], there is a 93% predictability in user movement, based on a 3-month-long record. Moreover, predictability was not found to be lower than 80% for any given user. While there have been many activities on these topics from different communities, here we only mention a few. In [7], stored radio maps were used for predicting transmission rates. A similar approach is seen in [8], where a user pre-buffers low quality

segments before entering a bad coverage area. The above studies mention the need for best QoS video delivery, the existing works on utilization of channel information is very limited. A look at the studies on power consumption are also presented here to understand the efficiency in power usage that can be derived. In [9], the authors consider the impact of predictive resource allocation on base station (BS) power consumption and the network streaming quality. Finally, [10] presents a predictive resource allocation (PRA) that improves rate predictions for enhanced and efficient delivery of stored videos. However, the performance of PRA on real channel traces was not evaluated, instead the channel conditions were modelled by using the distance from a connected BS and calculating the received power with the help of a path loss model. Different settings were used which rendered different results, however the trends and relative results coincide.

In this paper, we apply the methodology from [10] to real channel data, collected using several physical campaigns, in order to understand how a location and prediction-based resource allocation algorithm would fair, compared to traditional resource allocation schemes, and if gains from synthetic channels are maintainable in real channels. We have evaluated three resource allocation schemes: equal share allocation and rate-proportional allocation were implemented along with the PRA algorithm for comparison purposes. Performance was evaluated in terms of video degradation, power consumption and throughput of the network.

II. SYSTEM MODEL

A network of M users (indexed by m) and K BSs is considered. BSs serve users with video streams, which are available at the BSs. Time is slotted in slots of duration τ , indexed by t . In each slot, the wireless channel is shared between multiple users, while keeping the achievable data rate constant for each user.

A. Signal Model

The BS association for each user is known for all times, with the set of users assigned to BS k at time t denoted by $U_{t,k}$. The received power at time t at user m is denoted by $P_{t,m}$, with corresponding rate

$$r_{t,m} = \tau B \log_2 \left(1 + \frac{P_{t,m}}{N_0 B} \right), \quad (1)$$

where B is the bandwidth and N_0 is the noise power spectral density. We assume that at time 1, the received powers

$P_{1,m}, \dots, P_{T,m}$ are known (see Section III) for T time slots in the future. We denote the corresponding rates by $\mathbf{r}_{1,m}^{\text{LA}} = [r_{t,m}]_{t=1}^T$ and $\mathbf{r}^{\text{LA}} = [\mathbf{r}_{1,m}^{\text{LA}}]_{m=1}^M$. We further introduce \mathbf{x} , of same dimensionality as \mathbf{r} , with $x_{t,m} \in [0, 1]$ denoting the fraction of airtime that user m is assigned from a BS in time slot t . Hence, for an allocation \mathbf{x} , the predicted rate of user m at time t is given by $x_{t,m}r_{t,m}$.

B. QoS Model

All users request video content at time 1, with the same streaming rate A (bps). Hence, the minimum video content available in the buffer of user m , expressed in bits, at each time slot δ , required for smooth streaming is calculated to be $D_{\delta,m} = A\tau\delta$. Hence, for smooth playback, we require $\sum_{t=1}^{\delta} x_{t,m}r_{t,m}^{\text{LA}} \geq D_{\delta,m}$. We define the video degradation (VD) experienced by user m at time t as

$$\text{deg}_{t,m}(\mathbf{x}) = \max \left(A\tau t - \sum_{t'=1}^t x_{t',m}r_{t',m}^{\text{LA}}, 0 \right). \quad (2)$$

C. Goal

Our goal is to find \mathbf{x}_t to provide good QoS to each user. To this end, we use three resource allocation methods:

- 1) *Equal share*: each user gets the same amount of time in each time slot. This can be written as $x_{t,m} = \tau/T$.
- 2) *Rate proportional*: each user gets an amount of time in each time slot, proportional to its rate: $x_{t,m} = \tau/T \times r_{t,m}^{\text{LA}} / (\sum_k r_{t,k}^{\text{LA}})$.
- 3) *PRA*: the allocation accounts for the QoS of each user. For one formulation [10] of this, we have

$$\min. \quad \sum_{m=1}^M \sum_{t=1}^T \left(\frac{1}{D_{\text{tot}}} \text{deg}_{t,m}(\mathbf{x}) + \frac{\lambda}{MT} x_{t,m} \right) \quad (3a)$$

$$\text{s.t.} \quad \sum_{m \in U_{t,k}} x_{t,m} \leq 1, \forall t, k \quad (3b)$$

$$x_{t,m} \in [0, 1], \forall t, m \quad (3c)$$

$$\sum_{t=1}^{\delta} x_{t,m}r_{t,m}^{\text{LA}} \geq A\tau\delta, \delta \in \{1, \dots, T\}, \quad (3d)$$

in which λ is a trade-off parameter, indicating the importance of minimizing air time over video degradation. Both components of the objective are normalized, where $D_{\text{tot}} = MA\tau T(T+1)/2$ is the total cumulative demand. We note that (3) is a linear program, which can be solved efficiently.

III. RECEIVED POWER PREDICTABILITY FROM MEASUREMENTS

The above PRA method relies critically on the assumption of predictability of the received power. To test this assumption, we have performed a measurement campaign with off-the-shelf cell phones.

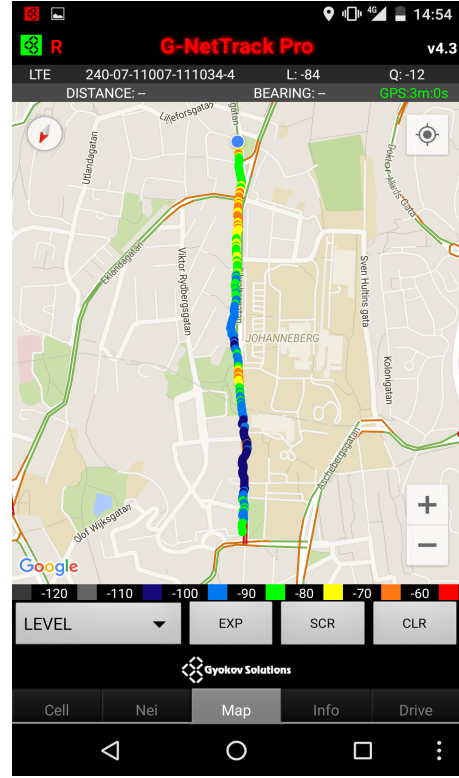


Fig. 1. Screenshot of the RSRP values along the walking route using the GNet Track Pro application.




Fig. 2. RSRP variation for the walking scenario data, with mean function and variations around the mean.




Fig. 3. Autocorrelation of the RSRP values after removal of the mean.

A. Measurement Methodology

The measurements were gathered from campaigns performed in Gothenburg, Sweden, during the spring of 2016. They were collected by using a Google Nexus 5X smartphone. Data was logged using GNet Track Pro by Gyokov solutions. The application allowed logging of signal strength, GPS position and many other parameters. It is important to note that an emphasis was made to maintain the same conditions, as much as possible, throughout all the measurement iterations. An empirical collection of data includes data collection at different times of the day. Therefore, for all the cases the measurement iterations were carried out during various times of the day to get a more ideal averaging of the readings. This was done until the characteristics, the mean and variance, for each position could be assumed to be approximated from the measurements. The phone was kept heads up on the right-hand front facing pocket of the jacket for all the measurements taken. These measurements were then mapped to a one-dimensional space with the distance from a fixed point as one of the parameters. The starting and ending points are located at known fixed geographical positions, allowing the same track to be recorded several times. Three LTE parameters (received signal strength indicator (RSSI), reference signal received power (RSRP), and reference signal received quality (RSRQ)) were collected [11]. RSSI provides information regarding the total received power which includes thermal noise and all interferences, RSRP on the other hand is the linear mean of the power contributions of the resource elements (REs), in a specific bandwidth, that carry cell-specific reference signals. A RE is defined as one 15 kHz subcarrier by one symbol. Finally, RSRQ is the ratio between RSRP and the RSSI and as such depends on the power of the serving cell. For our purpose, RSRP was the parameter which was used for modeling the available throughput.

Three measurements campaigns were conducted: one walking in a street, one in a tram going through a tunnel, and one




Fig. 4. Histogram of the RSRP values after removal of the mean.

on a bus touring the city. Here we only report results of the first campaign, as GPS measurements were of good quality throughout the path.

B. Measurement Results

Fig. 1 shows a snapshot of the RSRP values along the track. We observe large variations of the RSRP, depending on the position. The aggregate of all measurements is shown in Fig. 2, along the path. The figure shows that, over the span of multiple days, the RSRP at a given position is relatively stable, with variations due to environmental factors as well as GPS errors. Peaks in the RSRP are due to direct line of sight connections with BSs, while valleys are due to shadowing by large buildings and other structures. To complement these results, we have also analyzed the channel statistics after removing the mean from Fig. 2. Fig. 3 shows the autocorrelation function of the RSRP. We see that the RSRP values decorrelate after around 160 meters. Fig. 4 shows a histogram of all mean-removed RSRP values. We observe a relatively Gaussian behavior, with a standard deviation of around 3.3 dB.

In conclusion, the assumption of predictable received power appears to be reasonable, provided the path of the user along the prediction window is known.

IV. EVALUATION OF RESOURCE ALLOCATION FROM MEASURED DATA

In this section, we will evaluate the performance of the three resource allocation methods from Section II-C using the data described in Section III.

A. Performance Metrics

Three performance metrics are considered, based on the solution \mathbf{x}^* of (3), throughput, degradation, and power consumption. These are defined as follows:

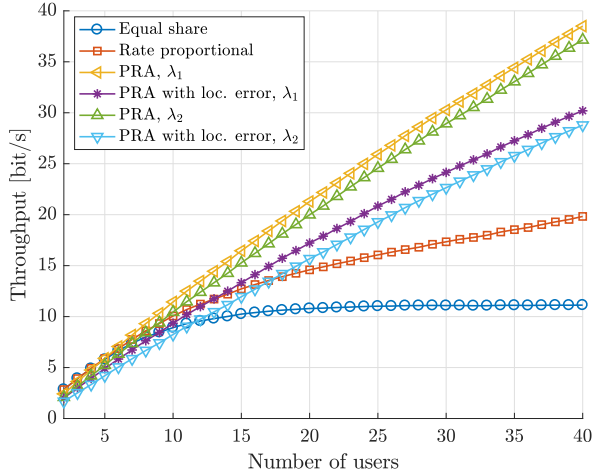


Fig. 5. Throughput comparison as a function of the number of users for equal share, rate proportional and predictive resource allocation. For PRA, $\lambda_1 < \lambda_2$.

- *Throughput*, defined as

$$\text{throughput} = \sum_{m=1}^M \sum_{t=1}^T x_{t,m}^* r_{t,m}, \quad (4)$$

in which $r_{t,m}$ is the actual rate supported by the channel (which may be different from the predicted rate, due to location errors).

- *Average video degradation*, defined as [10]

$$\text{degradation} = \frac{1}{TM} \sum_{m=1}^M \sum_{t=1}^T \text{deg}_{t,m}(\mathbf{x}^*). \quad (5)$$

- *Time-averaged BS power consumption*, defined as

$$P_{\text{avg}} = \frac{1}{T} \sum_{t=1}^T \sum_{k=1}^K (P_{\min} + (P_{\max} - P_{\min}) \sum_{m \in U_{t,k}} x_{t,m}^*), \quad (6)$$

where we recall that $U_{t,k}$ is the set of users assigned to BS k at time t . Here, P_{\min} is the power consumption under no load, and P_{\max} under maximum load. Here P_{\min} is set to zero, so that the degradation simplifies to

$$P_{\text{avg}} = \frac{1}{T} \sum_{t=1}^T \sum_{m=1}^M P_{\max} x_{t,m}^*. \quad (7)$$

We set P_{\max} to 1000 W. This yields a power consumption of $K \times 1000$ W if all K BSs are transmitting during an entire time slot.

B. Simulation Setup

We considered a look-ahead window of 35 seconds with a varying number of users attached to a single base station.¹

¹Due to limitations in the measurement setup, we were not able to obtain RSRP traces along the path for each individual base station.

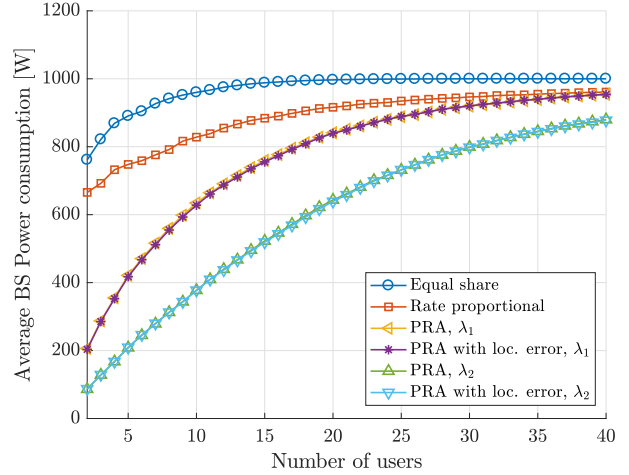


Fig. 6. Base station power consumption for different number of users for equal share, rate proportional and predictive resource allocation. For PRA, $\lambda_1 < \lambda_2$.

Each user traverses the path independently from different starting points with different velocities (drawn uniformly between 1m/s and 10 m/s). So for each simulation, each user starts at a random position along the path, with a random traveling direction, which the user follows for 35 seconds. We arbitrarily set B to 1 Hz, so rates should be seen with respect to this value. We set $A = 0.055$, so the video rate corresponds to 0.055 bits/s/Hz (e.g., corresponding to 1.1 Mbit/s over 20 MHz bandwidth). The background noise power is assumed to be -80 dBm, and the predicted signal power is assumed to be equal to the true power experienced by the user. The received power is taken as the mean from Fig. 2. Results were collected for 500 Monte Carlo runs. For the PRA method, we also investigate the impact of location errors, modeled as a shift of the initial position along the traversing path, for which we generated a uniformly distributed error between 1 m and 10 m. We set two different values of λ to determine the trade-off between power consumption and video degradation, $\lambda_1 = 1.3$ and $\lambda_2 = 2.5$. This latter value gives the priority of reducing airtime, while being tolerant to some video degradation.

C. Results and Discussion

Fig. 5 shows the throughput for the considered resource allocation strategies. It can be observed that PRA offers higher throughput compared to the equal share and rate proportional schemes. In fact, the PRA throughput grows approximately linearly with the number of users, which is the best that can be achieved. When the number of users are few, all methods perform equally well, since the system is under-utilized and the throughput grows linearly with the demand. However, equal share and rate proportional schemes see a relative drop of throughput after more than 10 users are in the system. As expected, the location uncertainty degrades the performance for the PRA scheme, especially when there are more users.

V. CONCLUSION

Predictive resource allocation relies on assumption of the predictability of the future received power. Theoretical studies have shown significant gains in terms of various performance metrics, both at the user side and the infrastructure side, provided this assumption is true. In this paper, we have reported the outcome of a measurement campaign with off-the-shelf hardware, in order to verify this assumption, and to perform predictive resource allocation based on real data. Our main finding is that indeed future received power can be predicted over relatively long time horizons, provided the future user locations are known. Our simulations indicate that under low mobility and low number of users, performance gains in terms of user quality are reduced compared to synthetic data, as fewer degrees of freedom and fewer chances to exploit high powers are available to the resource allocation algorithms. For all investigated scenarios, predictive resource allocation led to significant reduction in base station power consumption, while improving aggregate throughput.

ACKNOWLEDGMENT

The authors would like to thank Martin Dahlgren for his help in the experiments and Dr. Mouhamed Abdulla for the discussions. This research was funded, in part, under the ERC Proof of Concept Grant No. 680904 (MapCOM).

REFERENCES

- [1] "Cisco visual networking index: Forecast and methodology 2015–2020," <http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.pdf>, June 2016.
- [2] "Ericsson mobility report: On the pulse of the networked society," <https://www.ericsson.com/res/docs/2016/ericsson-mobility-report-2016.pdf>, June 2016.
- [3] J. Tadrous, A. Eryilmaz, and H. El Gamal, "Proactive resource allocation: Harnessing the diversity and multicast gains," *IEEE Transactions on Information Theory*, vol. 59, no. 8, pp. 4833–4854, 2013.
- [4] L. S. Muppirisetty, J. Tadrous, A. Eryilmaz, and H. Wymeersch, "On proactive caching with demand and channel uncertainties," in *Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2015, pp. 1174–1181.
- [5] J. Tadrous and A. Eryilmaz, "On optimal proactive caching for mobile networks with demand uncertainties," *IEEE/ACM Transactions on Networking*, vol. 24, no. 5, pp. 2715–2727, Oct 2016.
- [6] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, "Limits of predictability in human mobility," *Science*, vol. 327, no. 5968, pp. 1018–1021, 2010.
- [7] J. Yao, S. S. Kanhere, and M. Hassan, "Improving QoS in high-speed mobility using bandwidth maps," *IEEE Transactions on Mobile Computing*, vol. 11, no. 4, pp. 603–617, 2012.
- [8] H. Riiser, T. Endestad, P. Vigmostad, C. Griwodz, and P. Halvorsen, "Video streaming using a location-based bandwidth-lookup service for bitrate planning," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 8, no. 3, p. 24, 2012.
- [9] H. Abou-zeid, H. S. Hassanein, and N. Zorba, "Enhancing mobile video streaming by lookahead rate allocation in wireless networks," in *IEEE Consumer Communications and Networking Conference (CCNC)*, 2014, pp. 471–476.
- [10] H. Abouzeid and H. S. Hassanein, "Efficient lookahead resource allocation for stored video delivery in multi-cell networks," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2014, pp. 1909–1914.
- [11] "RSRP and RSRQ measurement in LTE," 2016. [Online]. Available: <http://laroccasolutions.com/training/78-rsrp-and-rsrq-measurement-in-lte>

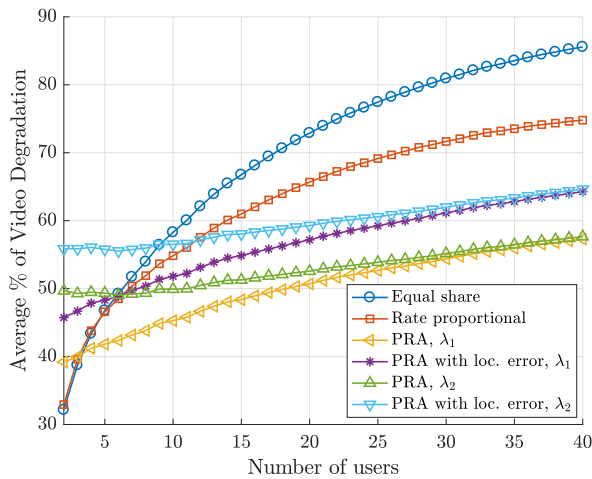


Fig. 7. Video degradation comparison for varying number of users for equal share, rate proportional and predictive resource allocation. For PRA, $\lambda_1 < \lambda_2$.

Fig. 6 shows the average power consumption of the BS for the three resource allocation schemes. PRA lead to lower power consumption than the non-predictive schemes. This is because to satisfy the demands, PRA is able to exploit good channel conditions in the future, requiring less power. However, this saturates to the level of the rate proportional schemes as the number of users increases. The equal share approach always has the largest power consumption, and reaches the maximum power after only 15 users. The impact of location uncertainty to PRA seems to be negligible from the power consumption point of view.

In Fig. 7, we show the video degradation performance of all the schemes. In contrast to throughput and power consumption, now PRA does not always offer the best performance. The conventional schemes have similar video degradation, offering the best performance for a low number of users, but quickly degrading as the number of users increases. PRA has slightly higher degradation for lower number of users, but degrades more gracefully with increasing number of users. When the system has a light load, PRA can obtain a lower objective function by reducing the airtime, translating to lower power consumption. We recall that by changing the trade-off parameter λ , we can reduce the video degradation, though at a cost in power consumption. For the case with maximum number of users (i.e., 40 users), we can see that PRA has a significant gain compared to rate proportional allocation schemes and while PRA with location error offers a somewhat reduced gain. These observed gains in our study are different to what has been reported in [10], which had much larger gain on synthetic channel data. Possible reasons include a difference in look ahead window size (4 minutes in [10], 35 seconds in this paper) and a difference in channel characteristics (synthetic data in [10], measured data in this paper) and different mobility (more mobility leads to better performance, due to increased channel diversity).