



HAL
open science

Throughput Prediction in Cellular Networks: Experiments and Preliminary Results

Alassane Samba, Yann Busnel, Alberto Blanc, Philippe Dooze, Gwendal
Simon

► **To cite this version:**

Alassane Samba, Yann Busnel, Alberto Blanc, Philippe Dooze, Gwendal Simon. Throughput Prediction in Cellular Networks: Experiments and Preliminary Results. CoRes 2016, May 2016, Bayonne, France. hal-01311158v1

HAL Id: hal-01311158

<https://hal.science/hal-01311158v1>

Submitted on 10 May 2016 (v1), last revised 19 May 2016 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Throughput Prediction in Cellular Networks: Experiments and Preliminary Results

Alassane Samba¹, Yann Busnel², Alberto Blanc³, Philippe Dooze¹,
Gwendal Simon³

¹Orange Labs

²Crest (ENSAI) / Inria Rennes

³Telecom Bretagne - IRISA

Throughput has a strong impact on user experience in cellular networks. The ability to predict the throughput of a connection, before it starts, will bring new possibilities, particularly to the Internet service providers. They could adapt contents to the quality of service really reachable by users, in order to enhance their experience. First this study highlights the prediction capabilities thanks to different algorithms and data gathered at different network levels. Then we propose a simple approach based on machine learning to predict the throughput using a few data related to the context of use.

Keywords: cellular network, throughput prediction, machine learning

1 Introduction

In cellular networks, Quality of Service (QoS), throughput in particular, depends on user context (radio channel quality, speed, distance from base station, etc.). To enhance the Quality of Experience (QoE), content providers implement *adaptive* delivery strategies, where the quality and the characteristics of the delivered content are adjusted to match the QoS of each user. These adaptive strategies are reactive : at a given time of the delivery, the throughput for the next x seconds is predicted based on the real throughput observed during the past y seconds.

Yet, the content providers take some key decisions at the beginning of the delivery. For instance, most web services have several style sheets for their web page, with a variable number of elements and information. The decision of which style sheet to deliver should be taken in the early moment of the connection although no past throughput observation is available. This calls for new mechanisms to provide a rough prediction of the throughput for the next x seconds, using only contextual information.

Existing solutions present several shortcomings preventing their widespread usage. For instance, a well-studied approach is to estimate the bandwidth by a series of short path measurements, including Round-Trip Time (RTT), and packet loss rate, but this approach requires exchanging data before making a decision. Some proposals rely on instant Channel Quality Indicator (CQI) to estimate the instantaneous bandwidth, but they do not target throughput prediction on a larger time frame. We discuss these approaches in Section 2.

We aim at identifying which contextual parameters are the most relevant to predict the throughput during a session in a cellular network. We have conducted a large-scale trial where users have performed a file download test a thousand times in several locations, mobility and radio connection configurations. Data related to context of use, radio access network performance and traffic quality have been collected for each test. In this paper, we describe this trial and we provide some early results of the statistical correlations between the main contextual information and the actual throughput.

2 Background and motivation

It is well known [MSMO97] that the throughput of a Transmission Control Protocol (TCP) connection can be accurately predicted from a set of measures related to RTT, window size and packet loss rate.

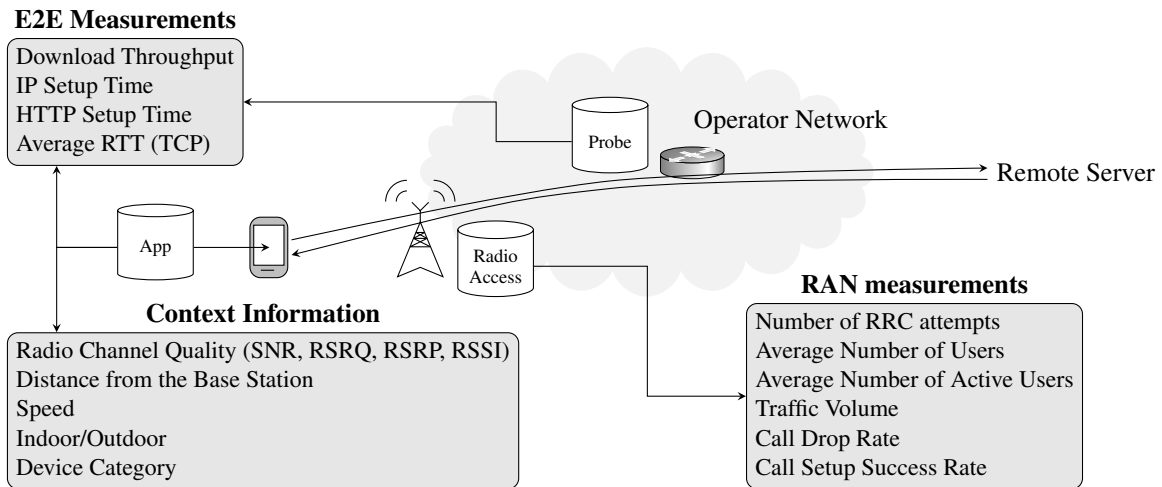


FIGURE 1: Testbed architecture

More recent outcomes [BMW14] adapt analytical throughput models to cellular networks by integrating the effects of the radio channel quality. These models, while extremely accurate, are not capable of making long term predictions.

More practical approaches such as [MSBZ07] and [LDJ⁺15] provide throughput prediction models. The former paper addresses throughput prediction for fixed line connections with machine learning approaches on several parameters, which are in part collected on the server side. This line of research does not address cellular environments. The latter paper addresses the prediction of instantaneous throughput based on some parameters that are collected at the mobile device side, including CQI and Discontinuous Transmission (DTX). This study is however restricted to the prediction of instantaneous throughput. The prediction is not accurate on a longer time frame since the CQI can change quickly, due to the changing radio conditions.

3 Experimental environment and methodologies

As shown in Figure 1, we have used an Android application [One16] to periodically download a file from a remote server, using an HSPA+ connection. For each download the application logs several parameters at the beginning and at the end of the connection, including location, speed, radio quality parameters (such as Signal-to-Noise Ratio (SNR), Reference Signal Received Quality (RSRQ), Reference Signal Received Power (RSRP) and Received Signal Strength Indicator (RSSI)). We were also able to collect data directly from the cellular network from a Radio Access Network (RAN) management system, which includes the number of active users, so-called **RAN measurements**. Finally, using network probes inside the operator network, we collected **end-to-end (E2E) measurements** for each test, including the download throughput we are interested in, the average RTT during the test, the IP setup time meaning the time between socket connection and first Hypertext Transfer Protocol (HTTP) request and the HTTP setup time meaning the time between first HTTP request and start of the download.

We present in this paper the results from a first set of measurements made by three testers. This campaign contains about 2600 observations from 40 different cells (without intra-cell handover). On average, the file was downloaded in eight seconds, which corresponds to the duration of a video segment in adaptive streaming. Unfortunately we miss some values due to the device firmware and Global Positioning System (GPS) availability. For example, RAN state data are missing in 13% of cases and speed in 20%. We used the optimal discretization algorithm defined by [Bou06] to deal with missing values and enhance the models.

Our goal is to predict the average throughput during the download test based only on context information. Several supervised learning techniques can be used to make the prediction. We compare the results of three algorithms, namely Generalized Linear Model (GLM), Neural Networks (NNET) and Random Forests (RF). We gradually add each group of predictors to assess its contribution to the throughput prediction. A

Throughput Prediction in Cellular Networks: Experiments and Preliminary Results




FIGURE 2: Throughput link with some context indicators.

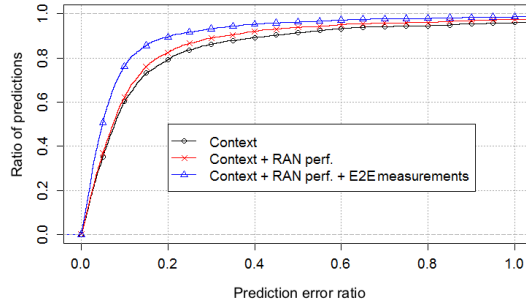


FIGURE 3: ECDF of error rates for the RF algorithms based on the three different types of inputs.

K-fold cross-validation [Koh95] permits to appraise the *generalizability* of models.

4 Evaluation and discussion

We first show in Figure 2 some typical correlations between the actual throughput and the measured context information. Due to space constraints, we cannot detail these correlations in this paper.

The machine learning algorithms predict a throughput, which we then compare to the actual throughput observed during the download test. Figure 3 shows the ECDF of the prediction error ratio, which is measured by the absolute value of the difference between the predicted and the actual throughput, divided by the actual throughput. We highlight two observations. First, the context information can suffice to make accurate predictions. The accuracy of the prediction is close to the prediction using both context and RAN information. Second, the machine learning algorithms are effective. Typically, half of the throughput rates are within 7% of the actual throughput while, for a fixed line, the previous proposals were within 10% of errors for the best half of predictions [MSBZ07].

	GLM	GLM	NNET	NNET	RF	RF
		10-fold		10-fold		10-fold
Context	0.82	0.80	0.84	0.80	0.84	0.84
Context + RAN data	0.88	0.85	0.89	0.86	0.89	0.88
Context + RAN data + other E2E measurements	0.95	0.93	0.95	0.93	0.95	0.94

TABLE 1: Coefficient of determination (R^2) of models.

In Table 1 we compare the results of the three aforementioned algorithms. We present the coefficient of determination (R^2) obtained on each model and on its 10-fold cross-validation. R^2 measures the part of variance explained by a model from global target variable variance. It is calculated as follows : $R^2 = 1 - \frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{\sum_{i=1}^n (\bar{y} - y_i)^2}$, where n is the number of observations, y_i is the target value of the i -th observation ($i \in \{1, 2, \dots, n\}$), \bar{y} is the average value of target and \hat{y}_i is the prediction for the i -th observation. Three algorithms produce roughly equivalent results for throughput estimation according to R^2 . Nevertheless RF models show better generalization capability than others, according to the small differences between R^2 of models from entire dataset and R^2 of 10-fold cross-validation. In fact, RF algorithms are known to be less prone to overfitting. These results also confirm that, unsurprisingly, the more data can be used by the machine learning algorithm, the better the prediction. But the prediction accuracy does not increase significantly with the RAN measurements (around 5% compared to the context information only). The accuracy increases with the use of other E2E measurements (especially average RTT) to estimate throughput, which confirms the results of previous papers. However, these data are not available at the beginning of the connection, while the context information is.

5 Conclusion

Predicting a transmission throughput through cellular network using a small set of information available before the connection is a challenge. Our results confirm the correlation between throughput and the context information, which opens important perspectives regarding the development of adaptive delivery techniques. Our approach uses information about the context and the coverage quality, which are available before the connection. Adaptive behavior can thus be implemented according to predicted throughput by a service provider to enhance QoE. A remaining open problem consists in taking into account sudden changes in coverage conditions that should occur during a connection, like a handover. Supplementary methods that tolerate evolution of coverage conditions and mobility prediction represents good candidate to handle this challenge. In the future works, we will also study the correlations between each parameter of the context information, as shown in Figure 2. Finally, we will make a deeper analysis by using the whole set of measurements, which contains more than 50 testers in a wider range of configurations.

Références

- [BMW14] Nicola Bui, Foivos Michelinakis, and Joerg Widmer. A model for throughput prediction for mobile users. In *European Wireless Conference*, 2014.
- [Bou06] Marc Boullé. Modl : a bayes optimal discretization method for continuous attributes. *Machine learning*, 65(1) :131–165, 2006.
- [Koh95] Ron Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 1995.
- [LDJ⁺15] Feng Lu, Hao Du, Ankur Jain, Geoffrey M. Voelker, Alex C. Snoeren, and Andreas Terzis. CQIC : Revisiting cross-layer congestion control for cellular networks. In *ACM HotMobile Workshop*, 2015.
- [MSBZ07] Mariyam Mirza, Joel Sommers, Paul Barford, and Xiaojin Zhu. A machine learning approach to TCP throughput prediction. In *ACM Sigmetrics conference*, 2007.
- [MSMO97] Matthew Mathis, Jeffrey Semke, Jamshid Mahdavi, and Teunis Ott. The macroscopic behavior of the TCP congestion avoidance algorithm. *SIGCOMM Comput. Commun. Rev.*, 27(3), 1997.
- [One16] EQual One. Get a 360° vision of mobile service experience ! – <http://www.v3d.fr/solution/equal-one/>. Vision 360 Degres (V3D), February 2016.