

Large-scale App-based Reporting of Customer Problems in Cellular Networks: Potential and Limitations

Yu Jin*, Nick Duffield*, Alexandre Gerber*, Patrick Haffner*, Wen-Ling Hsu*, Guy Jacobson*, Subhabrata Sen*, Shobha Venkataraman*, Zhi-Li Zhang†

*AT&T Labs - Research †Computer Science Dept., University of Minnesota
* {yjin,duffield,haffner,hsu,guy,sen,shvenk}@research.att.com †zhzhang@cs.umn.edu

ABSTRACT

In this paper, we study the *Location-based Reporting Tool (LRT)*, a smartphone application for collecting large-scale feedback from mobile customers. Using one-year data collected from one of the largest cellular networks in the US, we compare LRT feedback to the traditional customer feedback channel – customer care tickets. Our analysis shows that, due to the light-weight design, LRT encourages customers to report more problems from anywhere and at any time. In addition, we find LRT users access network services more intensively than other mobile users, and hence are more likely to experience and are more sensitive to network problems. All these render LRT feedback a valuable information source for early detection of emerging network problems.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations

General Terms

Measurement, Management

Keywords

Cellular network, troubleshooting, app-based reporting tool

1. INTRODUCTION

With the rapid growth in mobile voice and data services, effective management of large-scale cellular data networks is critical to meet customer demands and expectations. Due to the vast complexity involved, problems may occur in a number of different places, e.g., mobile handsets, software and apps running on the handsets, or within the cellular network infrastructure – the latter itself spans large geographical regions, consisting of thousands of cell towers, radio spectrum access controllers, and a whole gamut of other network elements and servers, supporting millions

of users. Identifying and pinpointing – not to mention troubleshooting – these problems can be an extremely challenging task.

Traditional troubleshooting approaches utilize network measurements, e.g., RTT, loss rate, collected by cellular service providers either at various locations in the network [1, 2] or at mobile handsets [3, 4]. However, such measurements do not necessarily reflect customers’ experience on the network. Due to this reason, direct customer feedback through the traditional customer problem report channel – customer care tickets – as a valuable source of information for troubleshooting problems in cellular networks has recently attracted more attention [5, 6]. A customer ticket is issued when a customer calls the technical support line and reports the problem to a customer agent, and it records the whole conversation between the two parties. Despite their usefulness, customer tickets involve high overhead. A customer needs to call in and wait on the line for a customer agent to speak to and spends time diagnosing the problem with the agent. Hence customer tickets depend heavily on the availability of customer agents. Because of this, a light-weight channel is demanded for real-time customer problem reporting and troubleshooting.

The increasing popularity of smartphone devices and the roll-out of more complex software and apps make possible a new channel of large-scale location-based customer trouble-reporting in cellular networks. Users can launch performance tests on their mobile handsets and inform service providers of any problems through smartphone apps. The *Location-based Reporting Tool (LRT)* is an example of such apps¹. LRT enables customers to report any performance problem by simply pressing a button, and the report sent via LRT, which we refer to as a *LRT message*, contains important information regarding the user’s location in the network (see an overview of LRT in Section 3). Since its debut in one of the largest cellular networks in the US, LRT has received more than 1 million downloads and millions of LRT messages have been collected in the past year. In this paper, we focus on making sense of these LRT messages. *How different are they from the traditional customer care tickets? What are the advantages and limitations of this new channel for detecting emerging network problems? In addition, since LRT users are self-selected – the user needs to actively choose to download and use LRT, are they a good representative sample of the entire mobile application, especially, in terms of troubleshooting network issues?*

To answer these questions, in this paper we study the unique characteristics of LRT feedback compared to customer care tickets (see Section 4). Our study demonstrates that the light-weight and simple design of LRT encourage customers to report more prob-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

W-MUST’11, August 19, 2011, Toronto, Ontario, Canada.
Copyright 2011 ACM 978-1-4503-0800-7/11/08 ...\$10.00.

¹For proprietary reasons, we cannot use the actual name of the app.

lems from *anywhere* and at *any time*. Because of this, LRT messages can indeed help detect emerging network issues much earlier than customer tickets, especially during nights and weekends, when the number of customer tickets are constrained by the limited number of customer agents. Moreover, the location information contained in LRT messages makes isolating problematic components in the cellular network easier.

We further conduct a comprehensive analysis of LRT users. Our study shows that the LRT users represent a self-selected non-uniform sample of the entire mobile user population. Compared to other users, LRT users tend to access network services, both voice and data services, more intensively and at more network locations (see Section 5). Furthermore, due to their preference of applications with stringent performance requirements, e.g., voice-over-IP and media streaming, LRT users are more sensitive to network performance variations. Both properties are desirable from the perspective of troubleshooting network problems, which enables LRT users to sense and report problems in the cellular network much earlier. Our analysis also points out several limitations of LRT, such as the small population size and the bias towards certain kinds of applications, and suggests remedies to make LRT a more practical network troubleshooting solution.

2. BACKGROUND AND DATASETS

Cellular Network Overview. The cellular network under study uses primarily UMTS (Universal Mobile Telecommunication System), a popular 3G mobile communication technology supporting both voice and data services. Fig. 1 depicts the key components in a typical UMTS network. The UMTS network has a hierarchical structure: where each *Radio Network Controller* (RNC) controls multiple node-Bs, and one *Serving GPRS Support Node* (SGSN) serves multiple RNCs (see [6] for details of the UMTS network).

Datasets. Our study uses LRT messages collected in the UMTS network for a one-year time period. To assist our analyses, we utilize additional datasets collected at various locations inside the UMTS network over the same time period, such as voice usage, data usage, Short Message Service (SMS) usage and so forth. We emphasize here that no customer private information is used in our analysis and all customer identities are *anonymized* before any analysis is conducted. Similarly, to adhere to the confidentiality under which we had access to the data, at places, we present normalized views of our results while retaining the scientifically relevant magnitudes.

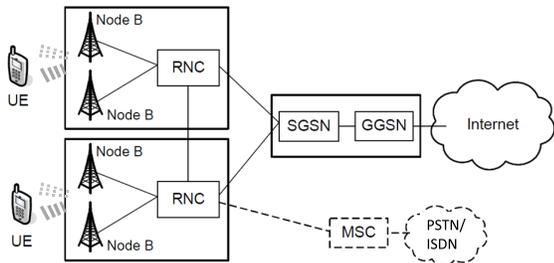


Figure 1: UMTS network architecture.

Customer Tickets. To study the difference of LRT from the traditional customer report channel – customer tickets, we collect all customer tickets during the same time period. Customer ticket is the default way for *all the mobility customers* to inform the service

provider regarding any problem by calling in a customer support line. A customer ticket contains the time of the call and a summary of the entire conversation between the customer and the customer agent during the call. We note that customers may call for a variety of reasons. A large majority of calls are *non-technical* related, e.g., questions about billing, service contracts, etc. Sometimes customers call when experiencing certain *technical* problems, e.g., unable to connect to the network, etc. Similar to [6], in this paper, we refer to customer tickets as these *technical tickets*, which contain none of the following keywords: bill, account, plan and feature.

Mapping Users to Network Locations. One of the key advantages of LRT is that it is a location-based report tool. Each LRT message contains the cell tower name that the customer is connected to. With this, we can easily correlate LRT messages at each level of the cellular network hierarchy (see Section 4.2). However, for other data sources, e.g., customer tickets, this information is not readily available. We infer such information from *GPRS Tunnelling Protocol Control* (GTP-C) messages as follows.

When a customer wants to access the cellular network data service, a *GTP Create* message is sent to the GGSN (recall Fig. 1) to establish a GTP tunnel for the current GTP session, which contains the Location Area Code (LAC) and Cell ID (CID) of the node-B that is currently serving the customer. A *GTP Update* message will be sent to the GGSN to update the latest LAC and CID when the customer travels beyond a certain distance and a SGSN handover happens. When the customer finishes using the data service, the GGSN is informed by a *GTP Delete* message to remove the GTP tunnel and hence terminate the GTP session. By tracking GTP-C messages, we are able to associate customers with network locations with a good accuracy at RNCs or higher level network locations, e.g., SGSNs or cities [7].

3. OVERVIEW OF LRT

LRT is a smartphone application that provides customers a means to submit feedback on their network experience to their cellular service provider. LRT has a simple design, allowing users to report problems by simply pressing a button².

Three major problem categories and five subcategories are predefined, see Table 1. We note that the five subcategories may change along with different versions of the application. However, the three major categories – coverage, voice and data – remain the same. Today, LRT can be installed on a selected number of smartphone devices and requires access to the data service. We expect more mobile devices will support running LRT in the near future.

Table 1: Predefined LRT problem categories.

Major	Subcategory
Coverage	No Coverage
Voice related	Dropped Call <i>and</i> Failed Call Attempt
Data related	Data - Can't Connect <i>and</i> Data - Too Slow

In addition to these predefined problem categories, a user can also submit free-text comments about the event. Other additional features, such as viewing nearby free Wi-Fi locations, etc., are also provided. In our dataset, the free-text comments are empty in most cases. Therefore, in this paper, we rely on the predefined categories to classify reported problems. Fig. 2 illustrates the breakdown of different reported problems in four quarters during our observation

²When data service is not accessible, LRT messages will be buffered and then delivered after the connection has been re-established.

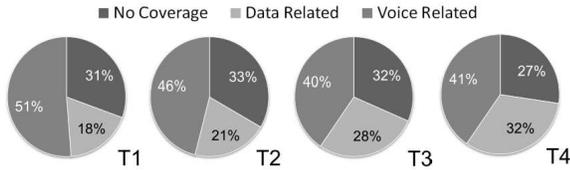


Figure 2: LRT message categories.

period, one calendar year (denoted as T1 to T4, in chronological order). Though the number of LRT messages received are different in the four quarters, voice-related problems always constitute the largest fraction (more than 40%) and the coverage problems account for approximately 30% of the problems reported. We find also that data-related problems become more significant over the calendar year, increasing from 18% to 32%, and this is consistent with our observations with respect to the growth of usage and expectation from mobile customers on data services.

4. COMPARING LRT MESSAGES TO CUSTOMER TICKETS

As LRT is a new approach for mobile customers to report network problems, the LRT messages have characteristics distinct from more traditional ways of reporting problems, i.e., customer care tickets. In this section, we compare these two channels of customer feedback, and our analysis highlights opportunities and limitations in detecting emerging network issues using LRT.

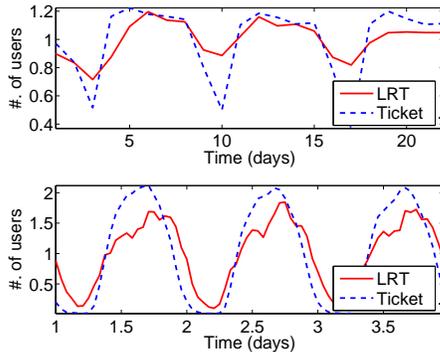


Figure 3: Number of ticket callers or LRT senders across time

4.1 LRT: Opportunities and Limitations

Any time reporting. Customer care tickets require phone interviews between the customers and the agents and thus are limited by the availability of the customer agent resource. In comparison, *the flexibility of LRT enables users to report problems at any time.* This difference is reflected in the statistics of the tickets or messages generated across the day that we now describe.

Fig. 3 displays the number of users who send LRT messages and customer tickets over a two-week time period (For proprietary reasons, we normalize the actual numbers of tickets or messages by their respective means.). For ease of visualization, we show two graphs; the top graph plots a single point for each day, while the bottom graph plots one point per hour. Clearly, there is a strong time-of-day effect and day-of-week effect in both channels. In both graphs, LRT messages exhibit a lower variability across time, especially during nights and weekends, when there are limited customer agent resources and so much fewer customer tickets can be

generated. As we show in Section 4.2, using LRT messages, we can quickly detect network performance issues that happen during times when few customer tickets are available.

A second observation from Fig. 3 is that customer tickets peak in the late morning, while LRT messages peak mostly in the late afternoon. This is most likely due to the lag between the time when problems occur and when the user reports them. Indeed, detailed notes in customer tickets indicate that many customers report problems that happened a day before (or sometimes, even earlier). This reporting lag causes the the highest customer care call volume to occur around the beginning of the day. In contrast, LRT messages, due to their low cost, are sent more promptly. The peak of LRT messages in the afternoon coincides with daily periods of heavy traffic load; as network load increases (and especially when subscribers use demanding applications such as video streaming), the overall performance may degrade, and this causes users to send more LRT messages. We describe in Section 4.2 how LRT messages enables us to detect emerging network issues earlier than with customer tickets.

Anywhere reporting. Generating a ticket typically involves a substantial overhead on the part of the customer, e.g., he needs to wait on the phone until an agent is available to discuss his problem, and then spend time diagnosing with the agent. This overhead tends to discourage users from reporting every problem that they encounter; instead, most customer tickets are reported when users persistently encounter the problem. Thus from a location perspective, most tickets concern a user’s *primary usage locations* (i.e. where he stays most of the time), such as home or work place [6]. In contrast, sending an LRT message involves little overhead beyond pressing a button. Because of this, *LRT messages encourage reporting problems that occur anywhere the customer goes*, and we now show that this does indeed make a difference in terms of locations where customers provide feedback.

Recall that we use GTP-C messages to obtain the trajectory of a user’s physical path. We can estimate the *primary usage location* for the user at the RNC level. Similarly, we define the *primary LRT location* as the RNC that a user is mapped to when most of the LRT messages are sent³. Comparing these two metrics lets us see which network locations a user complains about most in LRT messages.

Fig. 4 shows the percentage of LRT users whose primary usage location differs from their primary LRT location, as a fraction of the LRT users who sent at least x LRT messages in September 2010. We present this analysis for both data-related LRT messages and all LRT messages separately in Fig. 4. We observe that, unlike customer tickets, in most of the cases (more than 72% for all LRT messages and more than 60% for data related LRT messages), LRT users complain about places different from their primary usage locations. One explanation for this behavior may be that customers prefer using customer tickets to report problems at their primary usage locations, in order to ensure that they can interact with a live customer agent and hence the problems can be resolved appropriately. In comparison, at the other locations (e.g., places that they pass through), customers tend to report problems via LRT, since these problems are less disruptive to users’ normal activities. In Fig. 5, we show the number of users observed (averaged on an hourly basis) at each RNC vs. the number of LRT messages received complaining about that RNC during one week in August 2010. Note that many RNCs associated with more LRT complaints only have a small population (note both the x -axis and the y -axis

³Our designation of a user’s primary usage location can be affected if they use Wi-Fi for data usage. We shall address such measurement bias in our future work.

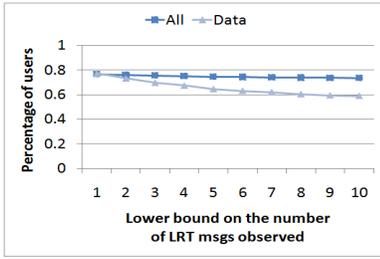


Figure 4: Proportion of LRT users whose primary LRT locations do not match their primary usage location.

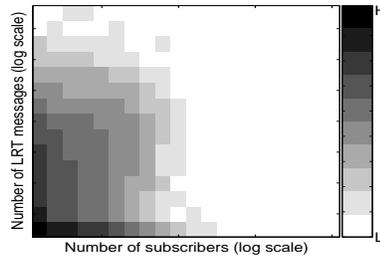


Figure 5: # of users vs. # of LRT msgs associated with each RNC

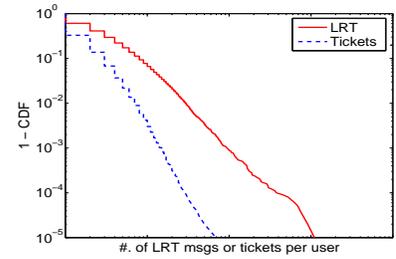


Figure 6: CCDF of the Number of tickets or LRT msgs per user

are in log scale). This suggests that at locations with a small stable user population, LRT can be a good complement to customer tickets for detecting network issues.

Increased Reporting. We find that *the low cost of LRT also encourages users to send more LRT messages*. Fig. 6 shows the number of LRT messages from each LRT user (who have sent at least one LRT message) compared to the number of customer tickets from each mobile user (who have initiated at least one ticket) in a whole month. We observe that more LRT messages are observed from LRT users than the number of customer tickets submitted by the mobile users. In particular, more than 20% of the LRT users send more than 5 LRT messages in a month, where only less than 3% of the users generate more than 5 customer tickets.

Limitations: Despite these advantages of LRT, LRT has a number of limitations. First, in most cases, an LRT message provides limited information, as it only contains a label from one of the five predefined problem categories. In comparison, customer tickets typically record detailed descriptions of the problems that the user has encountered. Second, the uni-directional nature of LRT reporting (i.e., lack of interaction between users and the service provider) is likely add some noise to customer feedback. For example, a user may report a connectivity problem with an LRT message, that in fact is due to a software issue or a problem with the mobile device. Such issues can often be eliminated by interaction with a live agent, as the agent steps through standard trouble-shooting when generating customer tickets.

Because of the limited information and added noise in LRT, diagnosing each individual LRT message may be difficult. Instead, a better way to use LRT messages may be to pinpoint emerging network issues, by analyzing the temporal and spatial correlations in them – intuitively, when a network problem occurs at a particular location, we expect a corresponding burst in LRT messages at that location. A similar approach has been used in detecting network problems using customer tickets [5, 8] with promising results. In the following, using this method, we compare the detection results using customer tickets and using LRT messages.

4.2 Detecting Network Problems using LRT Messages

In our study, we focus on one large city in the US. Fig. 7 compares the time-series of the number of LRT senders (top plot) vs. the number of ticket callers (bottom plot) per hour for a four month time period. We observe two bursts in the LRT time-series. The first one appeared at 2pm-6pm on the 75th day of observation. The ticket time-series shows a burst at 3pm-7pm on the same day. “No coverage” is the dominant complaint associated with both LRT bursts. We also validated that this incident reflected a problem in

the network – the network operators confirmed that in fact, the associated MSC was not processing incoming or outgoing calls, resulting in no or very degraded service to a very substantial number of users in the city.

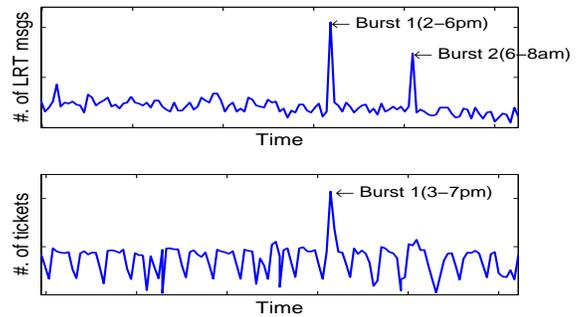


Figure 7: Correlation of LRT bursts with ticket bursts.

Interestingly, the second burst on the LRT time-series has no counterpart burst of customer tickets. Such a LRT burst has been confirmed to be associated with a network outage that happened at a particular RNC. Investigation shows that this LRT burst happened between 6am-8am, and it is possible that during this time period, people may not notice the problem or do not bother to report problems (perhaps if they are busy commuting). Another reason may be the limited customer care resources that handle complaints in early mornings. LRT, on the other hand, has no such constraints, and hence can detect such network problems missed by using customer tickets alone.

Further, LRT messages also report the node-B to which the user is currently connected, and this can help easily isolate problematic components in the cellular network. For instance, the first LRT burst in Fig. 7 can be isolated to a failed node-B, since most of the LRT messages are associated with one particular node-B. Likewise, we can attribute the second burst to an RNC failure. In contrast, customer tickets do not contain this information, and inferring this information with GTP-C messages may not be sufficiently accurate at node-Bs and cell towers (see [7]).

Despite many advantages of LRT messages over customer tickets in detecting emerging network issues, LRT-based detection may not be applicable at network locations where the LRT user population is small. Although LRT has received millions of downloads, the *LRT users still only account for a very small percentage of the mobile users in the network* and hence at many locations there are not enough LRT users to generate statistically significant bursts when problems occur. Moreover, LRT users are self-selected – the user needs to actively choose to download and use LRT. This nat-

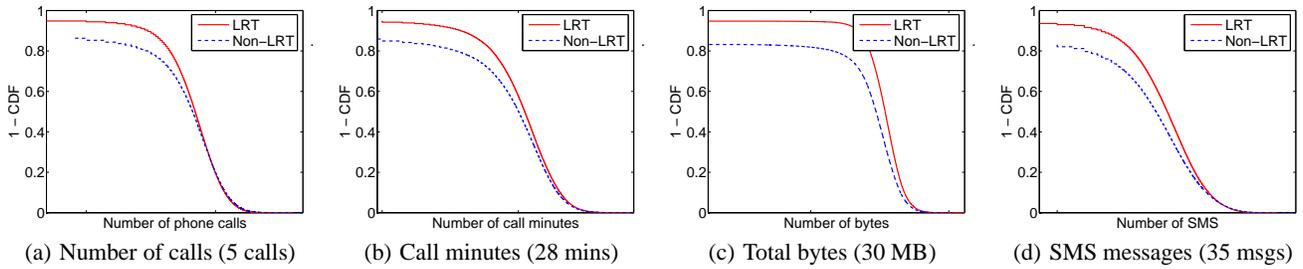


Figure 8: Comparing LRT users to Non-LRT users

urally leads to questions about *whether LRT users are good representatives of the whole user population. Do the problems reported by LRT users indeed affect other users?* To answer these questions, we compare from various aspects between LRT users and non-LRT users in the next Section.

5. ANALYSIS OF LRT USERS

We now compare the network usage patterns of LRT users with other mobile users, in order to analyze whether they are a representative sample of the entire set of users. For this analysis, we select two popular smartphone devices which support LRT. Let U denote the mobile customers who use one of the above two smartphone devices exclusively during the whole calendar year of observation. We denote $U_p \in U$ as the set of users who have sent at least one LRT message, which we refer to as the LRT user group; the rest of the users (denoted as $U_n := U - U_p$) comprise the non-LRT user group⁴. We find that U_p only accounts for a small percentage of U .

Data/Voice/SMS Usage. We first compare how these two groups of users access mobile services over a one-week time period in August 2010, including their voice usage (the number of calls made in Fig. 8[a] and the total call minutes in Fig. 8[b]), data service (the total bytes – both uploading and downloading – in Fig. 8[c]) and SMS usage (the number of SMS messages sent in Fig. 8[d]). We note that, in each plot, the x -axis is in log scale, therefore the difference between two CCDF curves is much larger than it appears to be. We also include inside the parentheses the difference between the median values of the two CCDF curves.

We observe that *LRT users typically use the network services much more intensively than non-LRT users*. This intensive usage of many different network services over a long period of time makes them more likely to experience network performance problems. We also observe that *LRT users use the network in more locations*. In particular, we find that the activities of LRT users span 2.2 miles (difference between the medians) more than that of non-LRT users during a week-long observation period. This also makes them more likely to experience a performance degradation.

We can also use customer tickets as another measure of whether LRT users indeed experience more problems. We compared customer (technical) ticket rates of LRT users and non-LRT users from August to October 2010, and we found that LRT users consistently report more tickets over time (persistently around 30% higher than non-LRT users). This also suggests their increased exposure and higher sensitivity to different network problems, all of which leads them to generate more customer tickets. These LRT users are also

⁴Of course, only considering these two smartphone devices introduces bias to our analysis. However, due to their predominance in the network, we believe such bias is negligible.

likely to seek additional tools (e.g., LRT) for reporting problems when the customer ticket channel is unreachable or inconvenient to use.

Application preference. In addition to differences in network usage, LRT and non-LRT users also favor different applications⁵. We show the ratio $P(app|u_p)/P(app|u_n)$ in Fig. 9, where $u_p \in U_p$ and $u_n \in U_n$. The dotted horizontal line represents $y = 1$ ⁶. A higher value of the ratio (greater than 1) indicates a higher chance that a LRT user will participate in that particular class of applications. We see that LRT users use more kinds of applications, especially smartphone *app* applications, than non-LRT users. This makes them more likely to be aware of the LRT application and try it out. We note also LRT users are also much more likely to use voice-over-IP and streaming, which are sensitive to variations in network performance.

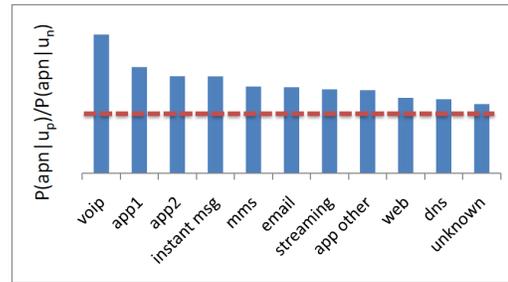


Figure 9: Application preference.

We further break down web traffic according to the content providers. Table 2 displays the top 100 content providers ranked by the ratio $P(apn|u_p)/P(apn|u_n)$ (the ratios associated with these 100 content providers are all above 1). Even though the difference between these two groups of users in terms of web usage is not as significant, they access very different sets of content providers. Interestingly, LRT users visit a lot of e-commerce, popular media, social networking and blog sites. Extensive activity on these sites will again make them more sensitive to network performance changes.

In summary, LRT users are a self-selected group that are quite different from most users in a number of dimensions – how intensive their network usage is, how many network locations they access, how diverse their applications are, how demanding the requirements of their applications are. Thus, LRT users are a little like canaries in a coal mine, since they are very sensitive to

⁵We match level 4 and level 7 headers in packets with predefined manual rules, we classify traffic into 12 application classes. The details are in [9].

⁶*app1* and *app2* are the most dominant smartphone apps in the network.

Table 2: Popular content providers for the LRT users

Category	Count	Examples
E-commerce	17	ebay, amazon,groupon, slickdeals, etc.
Ads	17	adbrite, tapjoyads, admarvel, adsonar, etc.
Media	16	tv.com, shazamid, transpera, turner.com
Tool	14	bit.ly, sitemeter, flurry, recaptcha, etc.
News	10	localwireless, cnn.com, nytimes, go.com
Social Network	7	digg, linkedin, twitter, plusplus, etc.
Blogs	6	wordpress, sharethis, blogspot, blogger, etc.
Weather	4	accuweather, weather.com, etc.
Photo	4	Picasa, flickr, imageshack.us, gravatar
Other	5	secureserver.net, gmail, etc.

the problems and potentially exposed to more of them than typical users⁷. On one hand, this is advantageous as these LRT users may help alert us the emerging network issues much earlier before most users notice a performance degradation. On the other hand, since LRT users use applications (e.g., VoIP, streaming) that are much more sensitive to performance variations, the issues detected by LRT users may not necessarily affect other mobile users. For example, an abnormally high latency may affect an LRT user watching video but is tolerable to a mobile user only sending e-mails. In addition, as we have seen in Section 4, LRT users report many problems happening at the network locations with a small number of users. Troubleshooting based on LRT messages alone may therefore not be cost-effective from a service provider’s perspective. One way to address this may be to prioritize those problems detected from LRT messages according to the potential number of customers affected by each problem, i.e., the stable customers at that location. Furthermore, expanding the LRT population (e.g., through advertisements on popular content providers, or pre-installing LRT on user devices) can help create a more representative sample and potentially detect more network problems. We leave these as our future work.

6. RELATED WORK

There is a rich literature in detecting and troubleshooting network problems in large networks. A majority of work focus on detecting, locating or trouble-shooting wired/wireless IP data network problems using passive or active network measurement data, e.g., via expert rule-based inference [10] or machine-learning techniques [11, 12], or via inference of dependency among network elements, entities and events [13, 14], or correlating bursts of customer tickets with other network events [5]. Our work differs in that we focus on studying a new channel of large-scale customer feedback to cellular service providers. We demonstrate unique characteristics that distinguish this new channel to traditional customer care tickets. Our work sheds light on how to make sense of this new channel and how to apply it for detecting emerging network related issues.

7. CONCLUSION AND FUTURE WORK

In this paper, we presented comprehensive analyses of a smartphone application LRT, a new channel for collecting large-scale customer feedback. We showed that LRT is a valuable light-weight channel that enables customers to report problems without temporal/spatial constraints. In addition, we found that LRT users access network services more intensively, making them good candidates to

⁷We can see LRT users as precursors of how users will be in a couple of years. Therefore, understanding their issues now can help improve the network for a near future when most users will use more sensitive apps.

sense emerging network problems. Our future work will focus on conducting detailed analysis of different problems reported from various channels, designing more advanced apps that are able to collect real-time performance metrics while a problem is reported. All these will lead to the development of a model for automatic detection and isolation of network problems by combing different customer report channels.

Acknowledgement

Zhi-Li Zhang is supported in part by the NSF grants CNS-0905037 and CNS-1017647, and an AT&T VURI gift grant.

8. REFERENCES

- [1] M. Chan and R. Ramjee. Tcp/ip performance over 3g wireless links with rate and delay variation. In *MobiCom '02*, pages 71–82, 2002.
- [2] X. Liu, A. Sridharan, S. Machiraju, M. Seshadri, and H. Zang. Experiences in a 3g network: interplay between the wireless channel and applications. In *MobiCom '08*, pages 211–222, 2008.
- [3] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, and D. Estrin. Diversity in smartphone usage. In *MobiSys '10*, pages 179–194, 2010.
- [4] J. Huang, Q. Xu, B. Tiwana, M. Mao, M. Zhang, and P. Bahl. Anatomizing application performance differences on smartphones. In *MobiSys '10*, 2010.
- [5] A. Mahimkar, Z. Ge, A. Shaikh, J. Wang, J. Yates, Y. Zhang, and Q. Zhao. Towards automated performance diagnosis in a large iptv network. In *SIGCOMM'09*, pages 231–242, 2009.
- [6] Y. Jin, N. Duffield, A. Gerber, P. Haffner, W.-L. Hsu, G. Jacobson, S. Sen, S. Venkataraman, and Z.-L. Zhang. Making sense of customer tickets in cellular networks. In *INFOCOM'11 Mini-Conference*, 2011.
- [7] Q. Xu, A. Gerber, M. Mao, and J. Pang. Acculoc: practical localization of performance measurement in 3g networks. In *MobiSys '11*, 2011.
- [8] T. Qiu, J. Feng, Z. Ge, J. Wang, J. Xu, and J. Yates. Listen to me if you can: tracking user experience of mobile network on social media. In *IMC'10*, 2010.
- [9] J. Erman, A. Gerber, M. Hajiaghayi, D. Pei, and O. Spatscheck. Network-aware forward caching. In *WWW '09*, 2009.
- [10] G. Khanna, M. Yu Cheng, P. Varadharajan, S. Bagchi, M. P. Correia, and P. J. Verissimo. Automated rule-based diagnosis through a distributed monitor system. *IEEE Trans. Dependable Secur. Comput.*, 2007.
- [11] I. Cohen, M. Goldszmidt, T. Kelly, J. Symons, and J. S. Chase. Correlating instrumentation data to system states: a building block for automated diagnosis and control. In *OSDI'04*, 2004.
- [12] B. Aggarwal, R. Bhagwan, T. Das, S. Eswaran, V. N. Padmanabhan, and G. M. Voelker. Netprints: diagnosing home network misconfigurations using shared knowledge. In *NSDI'09*, 2009.
- [13] P. Bahl, R. Chandra, A. Greenberg, S. Kandula, D. A. Maltz, and M. Zhang. Towards highly reliable enterprise network services via inference of multi-level dependencies. In *SIGCOMM'07*, 2007.
- [14] S. Kandula, R. Mahajan, P. Verkaik, S. Agarwal, J. Padhye, and P. Bahl. Detailed diagnosis in enterprise networks. In *SIGCOMM'09*, 2009.