

The Effect of Packet Loss on Redundancy Elimination in Cellular Wireless Networks

Cristian Lumezanu
Georgia Tech
lume@cc.gatech.edu

Katherine Guo
Alcatel-Lucent Bell Laboratories
kguo@bell-labs.com

Neil Spring
University of Maryland
nspring@cs.umd.edu

Bobby Bhattacharjee
University of Maryland
bobby@cs.umd.edu

ABSTRACT

Network-level redundancy elimination (RE) algorithms reduce traffic volume on bandwidth-constrained network paths by avoiding the transmission of repeated byte sequences. Previous work shows that RE can suppress the transmission of 20-50% bytes when deployed at ISP access links or between routers. In this paper, we focus on the challenges of deploying RE in cellular networks. The potential benefit is substantial, since cellular networks have a growing subscriber base and network links, including wired backhaul, are often oversubscribed. Using three large traces captured at two North American and one European wireless network providers, we show that RE can reduce the bandwidth consumption of the majority of mobile users by at least 10%.

However, cellular links have much higher packet loss rates than their wired counterparts, which makes applying RE much more difficult. Our experiments also show that the loss of only a few packets can disrupt RE and eliminate the bandwidth savings. We propose *informed marking*, a lightweight scheme that detects lost packets and prevents RE algorithms from using them for future encodings. We implement RE with informed marking and deploy it in a real-world cellular network. Our results show that with informed marking, more than 60% of the bandwidth savings of RE are preserved, even when packet loss rates are high.

Categories and Subject Descriptors

C.2.1 [Computer-communication networks]: Network architecture and design; C.2.3 [Computer-communication networks]: Network operations; C.4 [Performance of systems]: Performance attributes

General Terms

Performance, design, measurement

Keywords

cellular networks, redundancy elimination, loss

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.

IMC'10, November 1–3, 2010, Melbourne, Australia.

Copyright 2010 ACM 978-1-4503-0057-5/10/11 ...\$10.00.

1. INTRODUCTION

Cellular wireless networks are increasingly becoming a ubiquitous medium for Internet access. Cellular devices are set to overtake PCs as the most common Web access devices worldwide [9]. The increasing subscriber base and increasingly rich applications strain the capacity of both the cellular wireless links and the wired backhaul network that connects cell towers to the Internet.

Network-level redundancy elimination (RE) [19, 20, 5] has emerged as an efficient way to reduce traffic volume on bandwidth-constrained network paths. Redundancy elimination algorithms rely on deploying a cache at each end of the network path. For each packet traversing the path, the ingress node finds common sequences of bytes within packets that were previously sent (and stored in its cache) and replaces them with fixed-size pointers to the cache. At the egress, packets are decoded by replacing encoded content with data from the cache. Previous work found that RE can suppress the transmission of 20-50% bytes at ISP access links [20, 4]. More recent work has shown that RE is feasible even as an end system service, when the only redundancy to be eliminated is that within (not across) pairs of hosts [2].

Cellular network links are clear candidates for the deployment of redundancy elimination: they are bandwidth constrained [18, 6] and are likely to remain so for several years. Because RE is successful as an end-system service [2], RE may provide savings on the constrained wireless link, not just in the infrastructure shared by many users. Yet, applying RE to cellular networks is not straightforward: loss in the radio network can cause the two caches to become unsynchronized and hinder the decoding of repeated content.

In this paper, we study the feasibility of network-level redundancy elimination in high-loss environments such as cellular wireless networks. Using packet-level simulation driven by three large cellular traces and live experimentation on a North American CDMA 1x-EvDO network, we show that even the loss of a few packets can eliminate the bandwidth savings obtained with RE. To alleviate the harm caused by packet loss on RE, we propose *informed marking*, a simple mechanism to detect lost packets and prevent RE from using them in future encodings. Compared with RE in a lossless environment, RE with informed marking is able to preserve more than 60% of the bandwidth savings in a real network with high loss rates (above 10%).

Our contributions are as follows: First, we present a measurement study on the effectiveness of redundancy elimination in cellular wireless networks using three real-world

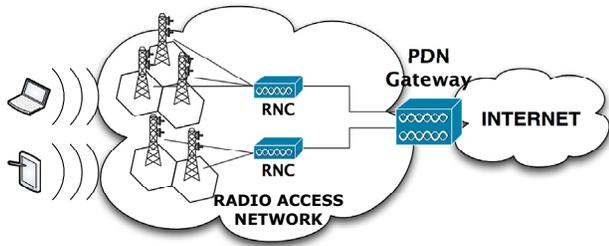


Figure 1: Simplified 3G/4G wide-area wireless network architecture.

packet traces gathered at wireless service providers in North America and Europe. Our analysis shows that most mobile users can save bandwidth with RE; some reduce their traffic volume by as much as 50%. Second, we observe that the high packet loss rates in cellular networks can pose problems for RE. We use simulations and live experimentation to understand how severe these problems are and show that even a few losses eliminate the benefit of RE. This is the first study on the effect of loss on redundancy elimination. Finally, we propose and evaluate informed marking, a lightweight mechanism that helps RE recover most of the bandwidth savings lost because of packet losses.

2. BACKGROUND

In this section, we describe the general architecture of cellular wireless networks, discuss existing RE techniques, and present the challenges that RE techniques face on cellular wireless links.

2.1 Cellular Wireless Networks

A cellular wireless network consists of a radio access network (RAN) that attaches to the Internet using a packet data network (PDN) gateway. We show a simplified example in Figure 1. We omit details and focus only on the path packets take from mobile station (MS) to the Internet. PDN gateways aggregate data from multiple radio network controllers (RNC) that receive mobile traffic over the air using base stations and wireless antennas. Data packets sent or received over the RAN are tunneled between the mobile devices and the PDN gateway.

The primary bandwidth bottleneck in wireless networks is in the airlink between the mobile users and the RNC [11]. WSPs have two solutions to keep the bandwidth demand below the available capacity: increase capacity or reduce traffic volume. Although 4G technologies, such as LTE, promise to increase the peak rates in the radio network to 100 Mbps, their deployment is expensive and will likely take several years, at which point they may remain bottlenecks.

2.2 Redundancy Elimination Overview

Network-level redundancy elimination (RE) is an efficient way of reducing traffic volume on bandwidth-constrained network paths. RE algorithms rely on deploying a packet cache at each end of a network path. The main idea is to compress outgoing packets by replacing sequences of bytes that have appeared in earlier packets with fixed-size pointers to the cached version of those packets. At the other end, packets are decoded by following the pointers and replacing encoded content with data from the cache. To function

correctly and efficiently network-level RE requires the two caches to be synchronized.

RE works because network traffic contains a significant amount of repeated content. Spring *et al.* [20] found between 20% and 50% redundancy in traffic collected at an enterprise access link. More recently, Anand *et al.* [4] confirm these results by analyzing traffic from 11 enterprise networks. They also discover that most bandwidth savings come from repeated sequences of bytes within each user’s traffic: a feature we exploit to deploy RE endpoints on cellular wireless clients.

Many commercial systems [16, 13, 14, 12] use RE in WAN optimization middleboxes placed on ISP access links. Recently, the EndRE system [2] proposed to push RE capabilities to end hosts. Neither middleboxes nor EndRE are suitable for cellular wireless networks. Although efficient for core Internet traffic, middleboxes cannot reduce traffic volume on last mile access links such as the radio access network for mobile users. On the other hand, EndRE can compress last mile traffic and, because it runs above the transport layer, is immune to losses in TCP traffic. However, potential packet loss will make it less effective for data streams that do not require reliable delivery (*e.g.*, UDP). Further, to effectively reduce traffic, EndRE must be installed on all servers that mobile users connect to. With the rapid increase in mobile users and applications, the destinations of mobile traffic become more diverse and application popularity becomes more dynamic. This makes the complete deployment of EndRE a difficult task. Even a partial deployment on popular servers is not trivial as it requires dynamic monitoring of server popularity.

2.3 RE over Cellular Wireless Links

To efficiently reduce the volume of traffic over the radio network, we must perform RE between each mobile user and the PDN gateway. This requires installing one cache on the mobile device and the other on the gateway. However, high packet loss rates in the radio network can create inconsistencies between caches and prevent RE from functioning correctly. A packet that enters the sender’s cache, but is subsequently lost, will appear (to the sender) to be useful packet for encoding; if it does, the receiver will not be able to decode any subsequent packet that refers to the lost one.

How high can loss rates in wireless networks be? To answer this question, we connected a stationary laptop to the Internet through a CDMA 1xEV-DO USB wireless card and ran an experiment to measure the loss rates that TCP may experience. We used *wget* to download the front page of each of the 500 most popular web servers [3] and *tcpdump* to capture all network traffic (both uplink and downlink) during the downloads. We have chosen popular, well connected destinations to minimize the probability of losses on the wired section of the path. We compute the retransmission rate during the experiment. Because TCP may retransmit packets that are not necessarily lost, the retransmission rate can be viewed as an upper bound for the packet loss rate.

We present the results in Figure 2. We show, in log scale, the number of packets sent and retransmitted during every 1 minute interval for 100 minutes. The retransmission rate tends to oscillate between 0.001 and 0.10, with peaks of over 0.30. In Section 3.4, we analyze the effect of loss rates on the bandwidth savings obtained with RE and show that even small loss rates can affect the correctness of RE.

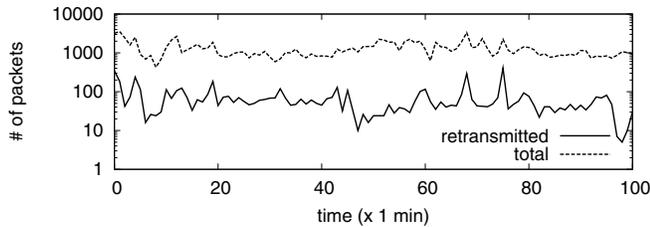


Figure 2: Number of total and retransmitted packets measured in a CDMA wireless network while downloading the front page of 500 most popular web sites (Total packets do not include retransmitted packets).

WSPs have made efforts to reduce the packet loss rate in their networks. This is mainly because high loss rates over the air interface also affect the performance of TCP, which interprets lost packets as a sign of congestion. Link layer retransmission protocols (such as RLP in CDMA [17] and RLC in UMTS [1]) ensure a link layer loss rate well below 1% [7]. However, when an IP packet is segmented into multiple link layer frames smaller than the maximum IP packet size, the IP loss rate is higher than link layer loss rate [10].

Existing RE schemes address loss by retransmitting the lost packets uncompressed. This may be sufficient in wired networks, where packet losses are few, and simply relying on retransmissions is enough. In wireless networks, however, this is not efficient: the large number of retransmissions may consume the bandwidth saved with RE, as we show in Section 3.

3. NORMAL REDUNDANCY ELIMINATION IN A LOSSY ENVIRONMENT

In this section, we analyze three cellular wireless traffic data sets and show that many mobile users can reduce their traffic volume with RE. We also show, using simulations and live experiments, that lost packets can significantly degrade the performance of RE.

3.1 Data Sets

We analyze three network traces, collected using *tcpdump*, at PDN gateways in two North American and one European 3G wireless service provider (WSP) networks. Table 1 shows the properties of each trace. Because we study the effects of RE for individual mobile users, we select the top k users that have sent or received the largest amount of traffic in each trace. We set k to 2,000 for WSP-1 and to 100 for WSP-2 and WSP-3. This ensures that the top k users are responsible for more than 80% of the traffic we capture in each network.

3.2 RE Algorithm Basics

We use the RE algorithm developed by Spring *et al.* [20] (commonly known as **modp**). For each packet, **modp** identifies content present in earlier packets by generating a list of fingerprints [15] associated with the packet. A fingerprint is a hash over a continuous sequence of bytes of length w . It is impractical to store all fingerprints that can be generated by a packet, since each byte in the packet generates a new fingerprint. Instead, **modp** chooses a fraction $1/p$, and stores only the fingerprints with value $0 \pmod p$. This allows

	WSP-1	WSP-2	WSP-3
when	10/2007	08/2008	05/2009
duration	30 min	1 hr	5 min
technology	CDMA	UMTS	UMTS
top users (k)	2,000	100	100
volume	0.5 GB	1.6 GB	0.8 GB
distribution:			
- TCP	59%	80%	97%
- UDP	36%	18%	2%
- ESP	1%	1%	0%
- others	4%	1%	1%
direction:			
- downlink	62%	89%	82%
- uplink	38%	11%	18%
avg packet size:			
- downlink	328B	952B	1001B
- uplink	199B	183B	322B
avg loss rate:			
- downlink	0.11	0.03	0.03
- uplink	0.12	0.09	0.14

Table 1: Properties of wireless data sets.

for the fingerprints to be chosen independently of their offset in the packet.

Next we describe the process of eliminating redundant traffic on the link between two nodes. For ease of exposition, we consider one direction of traffic and call the nodes sender and receiver. Both the sender and the receiver deploy a packet cache and a fingerprint store. The packet cache stores all packets that were sent or received by the node, while the fingerprint store keeps the fingerprints of all packets stored in the packet cache. For each outgoing packet, the sender computes its fingerprints and verifies whether they are present in the fingerprint store. Matches identify repeated sequences of size w (the size of window over which the fingerprint was computed). We then retrieve from cache the packet corresponding to the fingerprint and expand the matched region byte-by-byte in both directions to find the longest possible match. We replace each matched region with an 8-byte key, that indicates the location of the matched region in cache (packet id and offset) and its length. To obtain the original packet, the receiver identifies the keys in the encoded packet and replaces them with the regions in cache to which they point. Figure 3 is an illustration of the encoded packet and encoding keys.

The bytes of the encoding key represent the bandwidth cost that we have to pay to eliminate redundancy. Ideally, the size of the key is as small as possible. Our design of the 8-byte key improves over the existing **modp** algorithm which uses a 12-byte key. The first four bytes identify the packet in cache. Unlike Spring *et al.* [20], we choose to encode a hash of the packet in cache rather than the fingerprint associated with the repeated sequence, which reduces the size of the encoding key by four bytes. Similar to the original **modp** algorithm, we use the last four bytes to encode three 11-bit integers representing the offsets of the matched region in the current packet and in the cached packet as well as the length of the match. To accommodate the additional bit, we take advantage of the fact that the sum of either offset and the match length cannot be larger than 1,500.

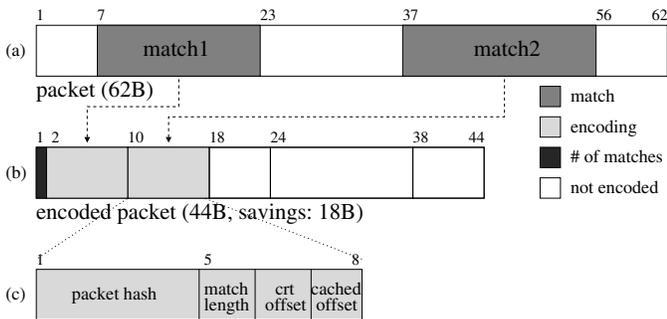


Figure 3: Encoding example: (a) incoming packet has two sequences of bytes seen in previous packets; (b) each match is replaced by an eight-byte encoding key; (c) each encoding key contains the hash of the packet in cache where the match was found (4 bytes), followed by the length of the match, offset in current packet, and offset in cached packet (all encoded in 9, 10, or 11 bits depending on their size.)

3.3 Bandwidth Savings

We define the *bandwidth savings* as the ratio between the volume of traffic saved after applying RE and the volume of traffic before RE. The bandwidth saved by each user depends on both the volume of redundant content that RE finds and the number of continuous repeated sequences of bytes (matches). This is regulated by parameters p and w of the **modp** algorithm. Recall that p determines how often we verify whether a fingerprint leads to a match and w represents the size of the smallest match. A small p has the potential to find more matches, but also strains the match-finding data structure.

We implement **modp** and run it for all users in each data set, for both uplink and downlink traffic. We vary the value of p from 1 (all fingerprints considered) to 256 (one in every 256 fingerprints considered). We present results for $p = 32$ because it generated the highest savings. We fix w to 8 because that is the size of the smallest possible match that would not waste bandwidth. We set the cache size to 64,000 packets. Unless otherwise noted, we use the same parameters throughout the rest of the paper.

Figure 4 (left) shows the cumulative distribution of bandwidth savings for all users in each data set. Uplink traffic is more redundant than downlink traffic. This is due to the increased number of identical or near-identical TCP ACKs and client requests sent on the uplink. WSP-1 has a higher overall redundancy than both WSP-2 and WSP-3. We speculate that this is because mobile data applications were not as diverse at the collection time of WSP-1 (October 2007).

3.4 The Effect of Loss

Loss in wireless is the result of transmission errors over the air interface, insufficient buffers at end points, or congestion. In Section 2, we showed that it is not unusual for wireless networks to experience high packet loss rates. Next, we study the effect of packet loss on bandwidth savings using simulations driven by our three data sets and live experiments on a US-based CDMA 1x-EvDO network.

3.4.1 Trace-based Simulations

To study the effect of loss under a controlled environment, we deploy RE on two computers connected through a wired

network. One computer acts as the sender and the other as the receiver. We replay the traffic for each of the k mobile users in the trace between the sender and the receiver while simulating a lossy connection, where each packet is lost with probability pl .

We capture the effect of loss on RE with two measures: bandwidth savings and RE loss. We use the bandwidth savings (defined in Section 3.3) to capture the volume of traffic saved in the presence of loss. To capture the correctness of RE, we define the *RE loss* as the fraction of packets that cannot be decoded by the receiver because previous packets, on which they depend, are missing from the cache. Figure 4 (middle) shows the distribution of RE loss for all uplink traffic in WSP-3. Even when the loss rate is 0.01 (≈ 1 in 100 packets is lost), 25% of the users cannot decode 20% of the packets they receive. Table 2 shows the results for all data sets.

We present the bandwidth savings for the uplink WSP-3 traffic in Figure 4(right) and for all data sets in Table 2. For the packets that receivers can decode, the bandwidth savings drop significantly (compared to a scenario where there is no loss): only around 20% of users see savings of more than 20% in uplink WSP-3 (compared to almost half when there is no loss). We also observe that the bandwidth savings and RE loss change with the loss ratio for WSP-1 but not for WSP-2 and WSP-3 (except for the RE loss rate for WSP-2 when $pl = 0.10$). We believe this is due to the lower volume of traffic that each user in WSP-1 sends or receives. In conclusion, packet loss affects both the performance and correctness of RE in wireless cellular networks.

3.4.2 Live deployment

To show the effect of loss on bandwidth savings under live network conditions, we deploy RE on a 3G-enabled laptop, as described in Section 2.3. The laptop tries to compress all outgoing packets, regardless of the destination, then tunnels them to a server on our campus network which decodes and forwards them to their destinations.

We run several experiments where we transfer, using UDP, parts of the Enron e-mail corpus [8] over the 3G connection. We transfer email because it is known to have higher repeated content. Each transfer takes 60 seconds and we send about 2MB of data. We demonstrate the effect of lost bandwidth savings due to loss, therefore we do not make any attempt to recover lost packets (*i.e.*, we assume that they could be retransmitted, without being compressed, by a higher layer).

Figure 5 shows the results of a representative experiment. The line labeled “expected savings” represents the evolution in time of the bandwidth savings expected by the sender, assuming no packet losses. The line labeled “actual savings” corresponds to the savings computed by the receiver from the packets it can decode (*i.e.*, bandwidth saved in the presence of packet loss). The average loss rate during the experiment, as measured with *iperf* was around 0.25. Although RE should be able to save around 63% of the bandwidth at the end of the experiment, due to loss, it saves much less—only 2%. This number is optimistic since it does not include the bandwidth lost by retransmitting the lost packets.

4. LOSS RECOVERY

Packet loss can compound when RE is used: missing packets will prevent receivers from being able to decode later

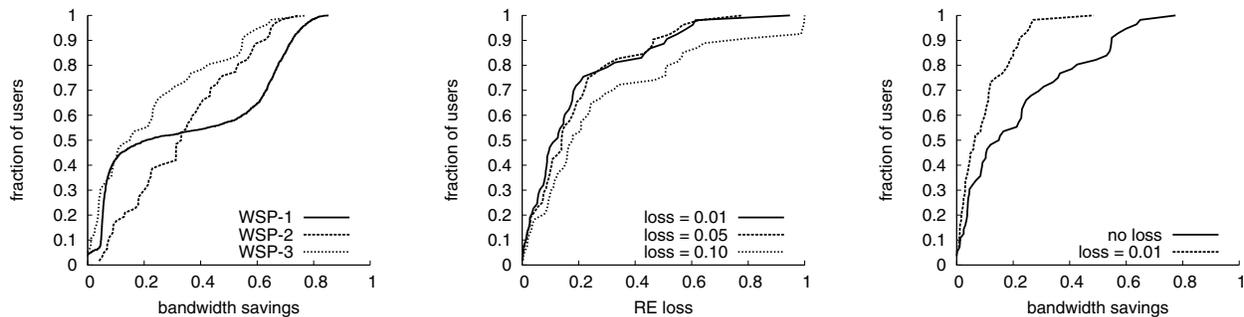


Figure 4: (left) Cumulative distribution of bandwidth savings without packet loss for uplink traffic, all data sets; (middle) Cumulative distribution of the fraction of packets that receiver cannot decode (RE loss), for various packet loss rates, uplink WSP-3; (right) cumulative distribution of bandwidth savings with and without packet loss, uplink WSP-3.

Data set	dir	no loss		loss					
		BW savings	BW savings			RE loss			
			$pl=0.01$	$pl=0.05$	$pl=0.10$	$pl=0.01$	$pl=0.05$	$pl=0.10$	
WSP-1	down	0.50	0.46	0.04	0.00	0.00	0.00	0.78	0.95
	up	0.65	0.63	0.20	0.03	0.00	0.00	0.33	0.44
WSP-2	down	0.22	0.03	0.02	0.02	0.45	0.48	0.49	
	up	0.46	0.15	0.14	0.10	0.23	0.26	0.55	
WSP-3	down	0.08	0.01	0.01	0.01	0.27	0.26	0.30	
	up	0.36	0.14	0.14	0.08	0.05	0.07	0.09	

Table 2: 75th percentile values for bandwidth savings and RE loss for each data set with and without loss.

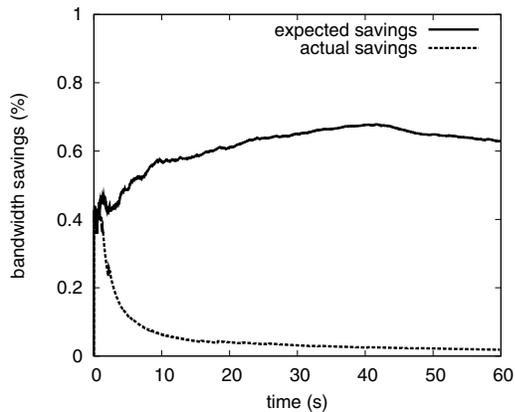


Figure 5: Expected bandwidth savings vs actual bandwidth savings when running RE in the presence of packet loss on a real 3G CDMA wireless network.

packets, which will then be dropped. We term this *RE loss*. Although we cannot prevent the original packet from being lost, we may be able to prevent the encoding from using lost packets and thus reduce or eliminate RE loss. For this, we must synchronize the sender’s and receiver’s caches. We consider two general solutions: (1) retransmit lost packets so that the receiver can decode future packets encoded based on them, and (2) detect lost packets and prevent the sender from encoding future packets based on them.

4.1 Loss recovery schemes

Existing RE schemes for wired networks [20] propose to retransmit lost packets. To detect loss, the receiver can wait for a packet that it cannot decode and request the

missing packet from the sender or rely on sequence numbers assigned by the sender and request retransmissions of packets with missing sequence numbers. The potential of a retransmission-based scheme is high: it could offer bandwidth savings similar to when there was no loss. However, a retransmission-based RE scheme will likely waste significant additional bandwidth in a high-loss environment such as cellular networks.

Another solution for reducing the RE loss is to detect and blacklist lost packets such that they cannot be used in future encodings. We consider two detection methods: *ack snooping* and *informed marking*. Ack snooping relies on the fact that in TCP, all packets correctly received are acknowledged. The scheme requires the sender to monitor all traffic from the receiver and blacklist any packet not acknowledged after a predetermined interval. Ack snooping quickly detects loss without explicit feedback from the receiver and can reduce the RE loss rate to 0. However, it can generate false negatives by mistakenly considering slow packets lost. Moreover, it does not work with transport protocols that do not acknowledge received packets such as UDP, which is common in cellular wireless networks (see Table 1).

In search for a more general loss detection technique, we propose *informed marking*, where each receiver signals the sender whenever it cannot decode a packet due to a missing packet from its cache. The receiver sends a control packet with the hash of the missing packet and the sender blacklists the corresponding packet in its own cache; in future encodings, the sender ignores repeated content matched to any blacklisted packet. Informed marking reduces the RE loss rate to at most the value of network loss. This is because receivers need to receive one packet that they cannot decode to be able to notify the sender. Unlike ack snoop-

Data set	dir	retransmission	ack snooping			informed marking		
		$pl=0.01, 0.05, 0.10$	$pl=0.01$	$pl=0.05$	$pl=0.10$	$pl=0.01$	$pl=0.05$	$pl=0.10$
WSP-1	down	0.53	0.53	0.53	0.53	0.52	0.52	0.53
	up	0.65	0.65	0.64	0.65	0.65	0.65	0.65
WSP-2	down	0.22	0.22	0.19	0.19	0.22	0.20	0.20
	up	0.46	0.46	0.45	0.45	0.44	0.43	0.37
WSP-3	down	0.08	0.07	0.07	0.07	0.07	0.07	0.07
	up	0.36	0.36	0.37	0.37	0.35	0.35	0.34

Table 3: 75th percentile values for bandwidth savings for loss recovery methods based on retransmission, ack snooping, and informed marking (IM). In simulations, RE loss is 0 for retransmission and ack snooping and at most the network loss for informed marking.

ing, informed marking works with any type of traffic and is flexible because it does not introduce feedback overhead when there is no loss or when lost packets are not used in the encoding.

We evaluate the benefit of the three loss recovery schemes in simulation, by computing the bandwidth they save. We consider an ideal scenario where lost packets are immediately retransmitted or detected. Table 3 presents the 75th percentile values for bandwidth savings on the three data sets. We make two observations. First, all three schemes perform similarly and are able to remove the effect of loss rates on bandwidth savings observed when no recovery scheme is used (Table 2). Second, we obtain similar savings regardless of the network loss rate. The results show that it is more important to prevent senders from encoding based on lost packets than to retransmit the lost packets; RE loss has a much bigger effect on the bandwidth savings than network loss. By eliminating or limiting the RE loss, we are able to obtain savings similar to when there is no network loss at all.

Because it is more general and it does not introduce unnecessary feedback overhead, we implement informed marking and deploy it in a real cellular wireless network. We present the results of the deployment next.

4.2 Deployment of informed marking

In our simulations, receivers react immediately when they cannot decode a packet and instantaneously notify the senders. It is possible that, in a real wireless network, by the time a sender receives a request to blacklist packet A, it has already encoded another packet with pointers to A. To understand the extent to which this affects the bandwidth saved, we use the same experimental run as in Section 3.4.2. In Figure 6, we show the actual savings with informed marking and the expected savings (note that the expected savings are smaller than in Figure 5 due to packet blacklisting). The actual savings are about 25% smaller than expected.

The difference between expected savings and actual savings would be explained if there are many matches between packets sent very close to each other (*i.e.* traffic has temporal locality). Anand *et al.* have already shown that, for enterprise traces, 60-80% of matches are between packets in 10% of the cache. To verify whether such behavior exists in our data sets, we define the distance between two packets as the difference between their indices in cache. We compute the median distance for each user and find that 75% of all users have a median distance of at most 5. This means that half of their packets are encoded using recent packets.

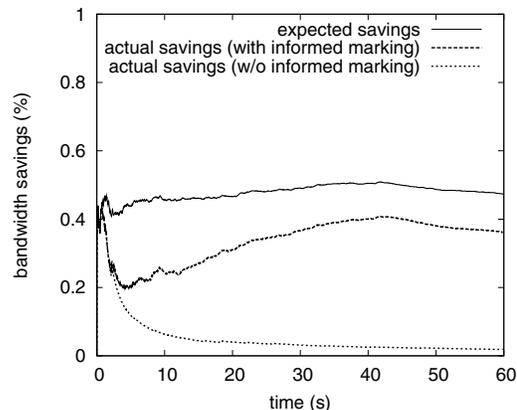


Figure 6: The benefit of informed marking on a real 3G CDMA wireless network.

If these recent packets are lost it is difficult to blacklist them in time to prevent RE loss.

5. CONCLUSIONS

We presented the first study on the effectiveness of traffic redundancy elimination in cellular wireless networks. Using real-world traces from three European and North American wireless service providers, we showed that the majority of mobile users sends or receives traffic with at least 10% redundant bytes. Yet, this redundancy cannot be exploited by redundancy elimination (RE) algorithms due to the high packet losses in cellular networks: even a few losses frustrate RE and prevent bandwidth savings. We proposed informed marking, a simple technique that allows RE to quickly detect and avoid lost packets. We evaluated informed marking both in simulation and in a live network and showed its effectiveness in reducing the effects of loss.

6. REFERENCES

- [1] Third generation partnership project, RLC protocol specification 3G TS 25.322. <http://www.3gpp.org>, 1999.
- [2] B. Aggarwal, A. Akella, A. Anand, A. Balachandran, P. Chitnis, C. Muthukrishnan, R. Ramjee, and G. Varghese. EndRE: An end-system redundancy elimination service for enterprises. In *NSDI*, 2010.
- [3] Alexa. <http://www.alexacom/>.
- [4] A. Anand, C. Muthukrishnan, A. Akella, and R. Ramjee. Redundancy in network traffic: Findings and implications. In *Sigmetrics*, 2009.
- [5] A. Anand, A. Gupta, A. Akella, S. Seshan, and S. Shenker. Packet caches on routers: The implications of universal redundant traffic eliminations. In *SIGCOMM*, 2008.

- [6] R. Chakravorty, S. Banerjee, J. Chesterfield, P. Rodriguez, and I. Pratt. Performance optimizations for wireless wide-area networks: Comparative study and experimental evaluation. In *ACM Mobicom*, 2004.
- [7] M. C. Chan and R. Ramjee. TCP/IP performance over 3G wireless links with rate and delay variation. In *Proc. MobiCom'02*, pp. 71–82, 2002.
- [8] Enron Email Dataset. <http://www.cs.cmu.edu/~enron/>.
- [9] Gartner Top End User Predictions for 2010: Coping with the New Balance of Power. http://www.gartnerinsight.com/download/Predicts2010_NewBalancePower.pdf.
- [10] H. Lin and S. K. Das. Performance study of link layer and MAC layer protocols to support TCP in 3G CDMA systems. *IEEE Transactions on Mobile Computing*, 4(5), 2005.
- [11] Motorola. Long Term Evolution (LTE): Overview of LTE air-interface technical white paper. <http://business.motorola.com/experienceLTE/pdf/LTEAirInterfaceWhitePaper.pdf>.
- [12] Netequalizer. <http://www.netequalizer.com>.
- [13] Packeteer. <http://www.bluecoat.com/packeteer>.
- [14] Peribit. <http://www.juniper.net>.
- [15] M. Rabin. Fingerprinting by random polynomials. In *Harvard University, Technical Report*, pp. TR-15-81, 1981.
- [16] Riverbed. <http://www.riverbed.com>.
- [17] Data service options for spread spectrum systems: Radio link protocol type 3, TIA/EIA/IS-7-7-A-2.10. <http://www.3gpp2.org>, 2000.
- [18] P. Rodriguez and V. Fridman. Performance of PEPs in cellular wireless networks. In *Web Caching and Content Distribution Workshop*, 2003.
- [19] J. Santos and D. Wetherall. Increasing effective link bandwidth by suppressing replicated data. In *Usenix Technical Conferences*, 1998.
- [20] N. T. Spring and D. Wetherall. A protocol-independent technique for eliminating redundant network traffic. In *SIGCOMM*, 2000.