Understanding the Super-sized traffic of the Super Bowl

Jeffrey Erman, K.K. Ramakrishnan AT&T Labs Research, New Jersey, USA {erman,kkrama}@research.att.com

ABSTRACT

Large events like the Super Bowl, where almost 75K attendees congegrate for several hours, poses a significant challenge in the planning, design and deployment of wireless networks. This was one of the first events where the LTE cellular network was available widely, in addition to almost 700 WiFi free hotspots. The Super Bowl in 2013 was also unprecedented because of a stadium-wide power outage for over half an hour. This study is the first to look in-depth at the user behaviours and traffic demand of a large ISP's celluar network at such an unique event.

The findings of this study can be used to guide the design of the communication networks of large venues in the future. There are several key insights from our study of the data collected. First, LTE speeds enable subscribers at venues to stream high-quality video and this can be a significant source of traffic. Second, the configuration of the uplink for such events is key, and a thoughtful approach to the design of applications that use the cloud for storing user data can substantially mitigate the congestion on the resource constrained uplink. Further, while it is tempting to take advantage of multicast on the cellular network (e.g., deploying technologies such as eMBMS in a venue), our results indicate that there is a need to combine multicast with caching to remove the strict requirement of overlap of requests from users to derive that benefit.

Categories and Subject Descriptors

C.2.3 [Computer System Organization]: Computer Communication Networks—*Network Operations*

General Terms

Cellular Networks; Large Venues; Characterization

1. INTRODUCTION

Cellular Data traffic has been growing rapidly, at well over 80% growth in total volume year over year, for the last 2-3 years, with the advent of smartphones [1, 2], tablets and other devices. What was mainly a network that was designed for telephony has now become a major carrier of data traffic. One of the characteristics

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions @acm.org.

IMC 2013 October 23-25, 2013, Barcelona, Spain Copyright 2013 ACM 978-1-4503-0193-9/13/10.

that provides relief is that the number of users simultaneously using a particular cell sector and channel is relatively small (in the few hundreds) at a typical cell site. However, a major challenge for cellular service providers is to manage the traffic demands at special events where several tens of thousands of users congregate at a given time. A salient example is the annual Super Bowl football game in the United States, where approximately 75K attendees descend on a sports stadium for about half a day.

The 2013 Superbowl in Febrary at the New Orleans Superdome is one of the first events to have the unique characteristics of a large number of users accessing a LTE cellular network deployed using a Distributed Antenna System, and combined with a stadium-wide deployment of WiFi hotspots. As reported by one cellular operator, the amount of data carried on the cellular network at the 2013 Super Bowl was 80% more than the previous year, with about 388 GBytes transferred [4]. During a long uprecedented power outage, users 'consumed approximately 10 GB more data than they did during any other hour' of the game [4].

Our in-depth study of the traffic characteristics of LTE cellular network of a large ISP at a large venue sheds light into the differences in usage from a typical cell site. Our results show access to web content (text, images) is still dominant. A major distinguishing factor of the traffic we observed is the amount of uplink data traffic - uplink data almost matched the volume of downlink data traffic. This is driven by the use of applications such as photo and video synching to cloud-based storage and multimedia messaging between users. We also observed that video consumption uses a significant share of the downlink data volume, because of the high bitrates at which video is being streamed to mobile devices despite having only a small number of subscribers concurrently consuming video.

This work can be used to influence future design of communication networks for large venues. Applications are changing how they use the network in dramatic ways - going from being predominantly downstream, to having a comparable amount of upstream traffic (nearly 1:1). This has important consequences in terms of configuration of the uplink for events. By formulating an optimization problem, we are able to show that having these applications becoming delay tolerant, even to a limited extent, can substantially reduce the amount of uplink capacity that needs to be provisioned (by almost a factor of two). Further, one would expect significant concurrency in accessing information across users, with such a large number of people being in one area and being interested in a common theme - Super Bowl information and related media streams. Our observation is that a straightforward deployment of multicast on the cellular network (such as eMBMS), even at a venue like the Super Bowl may not yield substantial benefits. We propose an intelligent approach to combining multicast with caching to remove the

strict dependency on the temporal overlap of requests from users. This helps us to get the benefit of multicast in the cellular network. We also observe large shifts in data volume at cell nodes in various parts of the venue: entrances first; then the seating areas. An intelligent self-organizing network could potentially re-allocate resources and power, depending on the phases of the event at a venue, to dynamically increase capacity to match the changing demand.

2. RELATED WORK

There have been many studies that have focused on understanding and characterizing the usage of smartphone traffic [7–9, 12, 14–16]. In many of these studies the focus has been on the characterisation of 3G traffic. However, a difference is our focus on large venues where LTE is available. Large venues are a unique consideration when the planning and operating wireless networks as the large number of people in the venue places an immense demand on resources. With LTE available, these venues are now capable of supporting even more traffic demanding applications which have a unique usage patterns due to the live nature of the event.

Most related to this work is the recent study by Shafiq et. al. [13] that looks at the performance of the 3G cellular network at two large venues hosting events. They found significant voice and data degradation. They found that due to the high demand for radio resources, pre-connection failures increase 100 to 5000 times compared to before the event. In addition, there are large increases in packet loss and round-trip times for user sessions. In terms of application usage, they found that there is an increased use of social networking. Due to demand on radio resources, they propose a connection sharing method to help aggregate demand and potentially improve performance. A key difference of our work is that we analyze a venue with a LTE network that was deployed, with significantly higher capacity. With LTE, several of the RRC state machine bottlenecks observed earlier are no longer as significant, allowing for more user sessions and higher traffic demands. For instance, we find a high percentage of the traffic is from streaming video with sustained bitrates over 1 Mbps. Another key difference is our focus on studying in-depth the user applications and traffic demands in relation to the events of the game (e.g., during the power outage) and where they take place within the stadium. In part, we find there is a significant demand from photo syncing to the cloud. Especially in the context of the available uplink capacity of cellular networks, we examine the effectiveness of optimizing this type of usage. We propose smoothing the syncing of these so as to significantly reduce traffic peaks which occur during breaks in the game.

3. METHODOLOGY

We first describe the data set collected and analyzed.

3.1 Venue

The New Orleans Superdome was setup with a permanent, high capacity Carrier Neutral Distributed Antenna System (DAS). The DAS system provides 4G LTE coverage with a collection of many strategically placed antennas, to maximize the radio coverage for all areas of the stadium - starting from the concourses, to the seating areas, etc. The DAS is organized as a set of non-overlapping cells to maximize the wireless capacity available in the stadium [4,5].

In addition to DAS-based LTE coverage, the Superdome also had over 700 WiFi access points fully deployed to provide significant capacity for offloading from the cell networks. The network supported 802.11n, 802.11a, b, and g in the 2.4 GHz and 5GHz bands.

According to a report [6], the average signal strength was -60 dB in the seating areas.

3.2 Dataset and Analysis

For this study, we passively measured the LTE traffic from the Superdome of a large US-based wireless provider. While we would have liked to also measure the traffic of the WiFi network to compare the usage of both cellular and WiFi, we were unable to do this and hope to do so in future efforts.

The data collection point was at the S1-U interface between the Evolved Node B (eNodeB) and the Serving Gateway (S-GW) for all eNodeB's serving traffic at the venue. The S1-U interface carries all the user plane traffic between the User Device (UE) and the cellular network for LTE. At the GPRS Tunneling Protocol (GTP) layer, each eNodeB can be identified specifically by an IP address. Different UE sessions can be distinguished using the separate Mobile Subscriber IP (MSIP) address assigned to each GTP session. Session terminations and hand-offs can be identified when the GTP tunnel identifiers for a specific MSIP change. While it can be used to differentiate unique UE sessions and count the number of simultaneous subscribers accessing specific data, the MSIP address cannot be used to map back to a specific subscriber or device, thus preserving user privacy.

We analyzed the HTTP header information of each request and correlated this together with flow records. The HTTP headers (and lower layer protocol headers) were collected for every request and response, without resorting to any sampling. For flow records of other protocols, in particular HTTPS, the DNS responses providing the IP addresses associated with each request/response were used to identify the content provider.

The privacy of subscribers was preserved as the study focused on aggregate statistics across all devices.

3.3 Identification of Upstream Applications

As we see in the subsequent sections, there is a significant amount of upstream traffic at the venue. It is useful to understand the particular applications that generate this traffic.

The multimedia messaging service (MMS) protocol makes use of HTTP(s) protocols to send and receive messages. These flows can be identified by the specific Hostname, and the Content-Type identifiers associated with the MMS service. Based on these signatures, we counted the total number of overall messages sent and received.

Traffic associated with Photo and Document synchronization to the cloud was also identified during the course of our analysis. Based on IP addresses and the associated DNS hostnames, we found these specific HTTPS flows going to subdomains at 'amazonaws.com' and 'windows.net' associated with the Photo and Document cloud sync service.

4. RESULTS

We present in this section an in-depth characterization of the traffic we measured at the Super Bowl.

4.1 Traffic Volumes

The overall traffic volume of the LTE traffic before, during and after the Super bowl is shown in Figure 1. For proprietary reasons, absolute values have been replaced with relative measures on the v-axis.

During the 2.5 hours before the start of the event, traffic volume steadily increases. Once the game began, traffic volume partially subsides, but picks up again towards the end of the second quarter of the game. During the half-time show, there is a substantial

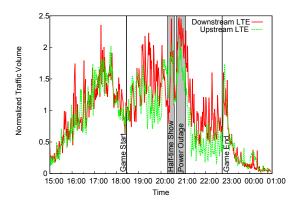


Figure 1: LTE Traffic volumes at Super bowl

spike in usage, subsiding again at the start of the 3rd quarter. But, just after the beginning of the 3rd quarter, there was a partial power outage in the stadium that lasted 34 minutes. The interval of power outage (which did not affect the cellular infrastructure) saw the traffic volume peak, with the long break in play. At the peak usage, the data volume is 7 times the volume before the game. After the power outage was over, the game resumed. Due to the excitement of the game in the 3rd quarter, we see that traffic usage drops off substantially. Finally, the traffic builds again at the end of the game.

Figure 2 shows the total number of unique UE devices active during each minute. An interesting observation is the number of users still accessing the cellular network even with the in-building WiFi network being available. In this figure, we can also more clearly discern the individual breaks in play for television timeouts etc. (such as at 18:45).

One immediate observation is that the traffic volumes and especially the number of active users varies widely with the interest in the game at hand. This leaves many peaks and valleys over short periods of time (i.e., over a couple minutes time period). The peakto-valley ratio over the whole game was 6.7. Even just considering the 1st and 2nd quarters, the peak-to-valley ratio was 3.1. This could be taken advantage of when scheduling uploads of images, messages or downloads of content. Having such applications capable of tolerating delays of a small number of minutes can substantially enhance system capacity. Concomitant with this is the need for the network to provide the necessary information to such delay tolerant applications to take advantages of uncongested periods in the network over short time scales.

4.2 eNodeb Volumes

In Figure 3 the traffic volume is shown for selected eNodeB's which cover (primarily) the different areas of the stadium. The traffic usage in the different areas as expected is not uniform over time. For instance, the entrances and concourses experience different traffic patterns than cells covering the seating sections of the game. This has strong implications on how DAS networks are designed, with LTE self-organizing networks [3]. The radio power and code allocations can be dynamically adjusted for each of the sectors covering the different sections based on this change in usage patterns to optimize overall resource utilization in the network (we do not attempt to cover the specifics of how the physical layer is managed in this paper).

An interesting distinction between subfigures (b) and (d) is the level shift in the traffic volumes for the eNodeB that cover the Middle and Upper sections of the stadium seating area. This level shift is due to the streaming video generated by an individual session. An important take-away is that even in LTE cell networks that have

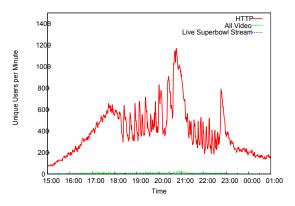


Figure 2: Number of HTTP users per minute

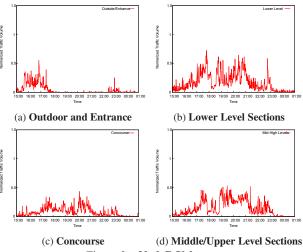


Figure 3: eNodeB Volumes

higher aggregate capacity for an individual cell site than the previous generation 3G networks, a small number of high bandwidth flows can consume a substantial portion of that capacity over short time scales. Therefore, the impact of such flows on the other users needs to be carefully considered.

4.3 Application Mix at Venues

We next look at which applications and content providers were used by subscribers. In Table 1 we find the main applications being used are HTTP-based web browsing, streaming video over HTTP and email. This is not wholly unexpected based on the plethora of other studies that report the same characteristics of smartphone traffic.

Table 2 shows which content providers were most accessed. We found surprisingly two cloud providers account for over 22.1% of the overall traffic. Upon further investigation, these were providing photo and document syncing to popular cloud services. The second largest provider is a sports website that provided a live video stream of the Super Bowl (discussed more below). Social networking and photo sharing sites were also popular. It is interesting that when summing over the top 20 providers, over 5% of the traffic was related to downloads for speed tests. In comparision to Shafiq et. al [13], there is a substantial difference in usage between the previous 3G results and our LTE results. Both syncing to the cloud and the live streaming were not previously observed.

We have not shown the content provider time series due to space considerations, but during the power outage, in additon to uplink

Table 1: Traffic by Application Class

Application	% Overall Traffic
Web Browsing (HTTP)	56.5%
Streaming (Over HTTP)	19.6%
HTTPS	13.0%
Email (SMTP(s), POP3(s), IMAP(s))	5.9%
Other (VPN, VoIP, etc)	5.0%

Table 2: Top 10 Content Providers

Content Provider	% Overall Traffic
Cloud Provider	16.6%
Sports	15.0%
MMS Message	9.6%
App Store	6.4%
Social Nework	7.8%
Cloud Provider 2	5.5%
User-Gen Video	4.7%
Email	2.3%
Search	2.3%
Speed Testing	1.6%

sources (see next section), there was a considerable increase in the use of social networking and photo sharing sites. But, there was not a corresponding increase in video traffic.

4.4 Video Consumption

We noted that video traffic accounted for 19.6% of the overall traffic at the Super bowl. We now investigate the sources, and the number of subscribers consuming video content.

Figure 4 shows the traffic volumes of the total downstream LTE traffic, the volume for all video traffic delivered over HTTP across all content providers, and the video traffic that related to a specific content provider that was providing a live stream of the Super Bowl. We can see that the live stream of the Super bowl accounts for 72% of the total video traffic delivered during the game. This live stream was using an Adaptive Bitrate Protocol (which adapts to the available capacity of the channel to the end-user) to deliver the content and was available at bitrates ranging from 50 Kbps to 2 Mbps. Table 3 shows the distribution of the bitrates that subscribers accessed. This evokes the key question of - why would users access a live stream at the Super Bowl while they are at the Super Bowl venue? One potential explanation would be that the live stream offers alternative views not available for the user to view on the in-stadium large monitors that display the game live. We investigated this and found that only 6.1% of the video chunks requested were for alternative views, while the remaining 93.9% were for the main television broadcast feed. Another possibility could be that these streams were delivered to users trying to get a better view of the game online. Upon futher investion of the User-Agent fields we found that over 99% of the live video stream requests were from tablet devices.

Another aspect we wanted to explore is – due the substantial video usage during the game for the same type of content would multicast would be an appropriate solution? An emerging capability for LTE cellular networks is eMBMS. Figure 5 shows, for both overall video and for the live Super Bowl streams, the number of simultaneous users during each minute. For video, we see that there are a maximum of 26 simultaneous users and for the Super Bowl

Table 3: Bitrates of super bowl stream (Kbps)

50K	200K	350K	550K	800K	1400K	2000K
0.0%	1.3%	0.7%	3.8%	5.1%	24.0%	65.2%

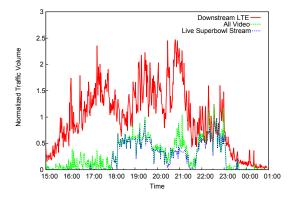


Figure 4: Video traffic during Super bowl

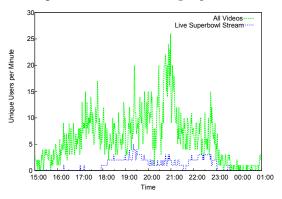


Figure 5: Concurrent users for all video and live streams

stream, a peak of 6 simultaneous users (on average less than 1 per cell). Thus, for the Super Bowl stream, a cellular operator would have to evaluate if it would be cost effective, both in terms of dedicating radio resources for multicast and the investment needed for deployment of the eMBMS solution, when such a small number of users concurrently access the stream.

4.5 Uplink Traffic

Access networks have traditionally seen a dominant amount of downstream traffic demand compared to the upstream demand. Currently on cellular networks, we have typically observed the ratio of downlink to uplink traffic to be in the order of 9:1. This is driven by primarily by content consumption by users. As such,the resources allocated for downstream versus upstream has also been correspondingly unbalanced. However, compared to the typical 9:1 ratio, the ratio of downstream to upstream, as observed at the Super Bowl in Figure 1 was near 1:1, which has significant implications to the cellular service provider.

The initial expectation was that the increased uplink traffic may be primiarily from social networking. However, in contrast to previous studies, we found the two largest sources of upstram traffic (which incidentally occured most during the peak demand interval that was during the power outage) was the background 'syncing' of pictures to the cloud and photo and video messaging through MMS. Demand for both of these services are highlighted in Figure 6

5. CELLULAR NETWORK CAPACITY MANAGEMENT AND SCHEDULING

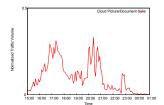
The achievable aggregate capacity of radio access network (RAN) in a cellular network is determined by a large variety of parameters and configuration settings. Some of these are specific

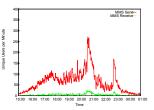
to individual vendor implementations of the RAN equipment (e.g., the base station, eNodeB, in an LTE network) and some are based on how a cellular network operator configures an individual cell site. The achievable capacity is quite different between the downlink and the uplink, because of the combination of considerations - the expected amount of downlink traffic versus the uplink traffic; the channel access protocol characteristics for uplink access for the mobile device to transmit data versus the protocol used for downlink transmissions from the eNodeB to the mobile; the overheads associated for the downlink vs. uplink protocols; the impact of contention and interference from transmissions by other mobile devices; and finally the scheduler characteristics. While a detailed treatment of each of these aspects is not feasible, we will provide a very high level overview to enable the reader to appreciate the considerations involved. In our context, it is of particular importance on how much capacity is available for the uplink versus the downlink and the overheads involved.

The first consideration is the amount of capacity that can be allocated based on the 'baseband capacity' that is available. Typically the hardware capacity sets a maximum throughput for downlink vs. uplink, and the ratio is often between 3:1 to (possibly) 4:1. This is certainly well in line with the observations of the proportion of traffic seen in a typical cell site, where there is is far more downlink traffic than uplink traffic. In addition, the cellular operator and equipment vendor may also dimension the actual amount of resources that could be allocated for the uplink. The 'physical resource blocks' (PRB) are dimensioned such that the aggregate amount of power and code resources available for all the traffic from a particular eNodeB available to be allocated to downlink and uplink traffic are managed to meet the traffic demand. Typically, the PRBs allocated also follow a ratio of near 2:1, so as to accommodate peak usage in one or the other direction.

Another important consideration is the design of the scheduler itself. Transmissions are scheduled on a 1 millisecond transmission time interval (TTI) (unlike 2 ms with 3G networks). With Orthogonal Frequency Division Multiplexing (OFDM), the radio resources are allocated for each individual flow in a manner that seeks to balance efficient usage of resources versus achieving fairness across the end user devices that are competing for resources. For each TTI period, the available PRBs along with the available capacity (e.g., downlink baseband capacity) are used by the scheduler to then schedule the transmission of data packets waiting to be sent to a particular end device (in the case of the downlink). In addition, the scheduler ensures that the end device is in a state to receive data. The frequency block that is used for transmission (e.g., downlink) may be of the order of 5 MHz or 10 MHz depending on what the cellular operator has allocated. This block is divided into subcarriers of a certain size. To avoid inter-cell interference, the frequencis selected for transmission may be selected by starting at a random subcarrier and assigning a number of subcarriers for that transmission. For downlink transmissions, only the eNodeB scheduler is involved. With LTE, Multiple Input Multiple Output (MIMO) is used on downlink transmissions. For the uplink transmissions from the mobile end device to the eNodeB, both the UE and the eNodeB are involved in the scheduling, with the buffer status information at the end device being communicated to the eNodeB. Often, on the uplink, MIMO may not be supported in current cellular networks.

Every TTI (i.e., 1 ms), the RAN scheduler determines the resources to be assigned to each end device that has data to transmit. A higher scheduling priority provided to an end device gives it a higher probability to be able to receive (from the eNodeB) or transmit data. The allocation of resources is on a per end device basis. The end device that is provided the highest priority is selected first





(a) Cloud Picture/Doc Uploads

(b) MMS Messages

Figure 6: Main Upstream Sources

Table 4: Upstreaming Applications

Application	% Uplink Traffic
Photo/Document Sync to Cloud	29.2%
MMS	12.7%
Social Network	7.8%
Email 1 (HTTPS)	3.0%
Seach	2.4%
Email 2 (HTTPS)	1.3%
Others	36.1%

(ignoring issues related to the different 'bearers' associated with the end device). For the uplink, the scheduling is also done on an end device basis, with the grant provided to the end device to transmit data based on its own local scheduling policy. For both the downlink and uplink, the hybrid ARQ-based retransmissions are given higher priority over new data packet transmissions.

The cellular operator typically configures the proportion of the downlink bandwidth to be used by the downlink scheduler, which then assigns physical resource blocks in the frequency domain, on a subframe/TTI basis. This essentially determines the downstream peak throughput and capacity. Uplink transmission scheduling is slightly more complex and transmissions are managed carefully by the eNodeB upstream scheduler to avoid interference. For uplink transmissions, the end device sends a scheduling request to the scheduler in the eNodeB to indicate that it has data in its buffer and requests an opportunity to transmit in an upcoming subframe. When the end device is scheduled in a sub-frame, the uplink scheduler at the eNodeB transmits a scheduling grant to the end device on a control channel, indicating the resource blocks and transport format to use for uplink transmission. Together with the uplink data, the end device transmits buffer status information to the uplink scheduler which uses this information along with the channel quality information to assign an appropriate number of resource blocks to the device in the uplink.

The scheduler in the eNodeB follows a policy that balances efficiency with fairness. The fairness criterion generally used is to achieve some level of proportional fairness, with a weight used to achieve efficient use of radio resources (power and code). The weight favors those end devices at the center of the cell or at the edge of the cell depending on whether the cellular operator wishes to provide a greater weight for efficiency or fairness.

For all these reasons, we see a need for applications to be aware of their use of uplink bandwidth. We examine ways of optimizing its use for some of the applications being used in the venue setting.

6. STRATEGIES FOR BETTER CONTENT DELIVERY IN VENUES

Based on our observations in the previous section, we now discuss some possible strategies for optimizing content delivery at venues.

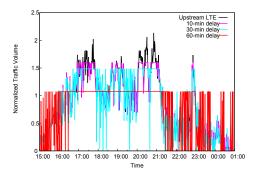


Figure 7: Peak reduction by delaying uplink traffic during congestion

6.1 Trading-off delay to mitigate congestion

One possible approach to mitigate the load from uplink traffic is to make applications congestion-aware and delay tolerant, so that load is shifted away from times when the network is congested. In that case, upload traffic would be only slowed under congestion.

To evaluate the benefit of having such congestion aware and delay tolerant applications, we formulated a linear program to examine the reduction in the peak uplink traffic for different amounts of delay tolerance and solved it using CPLEX. Let v_t denote the current upload rate. We define x_t as the upload rate after waiting a maximum of d. The objective is to find a solution x_t that minimizes the peak traffic during the time window T.

$$\min \max_{t \in T} x_t \tag{1}$$

s.t.
$$\sum_{t \le s \le t+d} x_s - v_t \ge 1, \, \forall t$$
 (2)

$$\sum_{t \le s \le t+d} x_s - \sum_{t \le s \le t+d} v_s \le 0, \,\forall t \tag{3}$$

$$\sum x_t - \sum v_t = 0, \,\forall t \tag{4}$$

$$0 < x_t, \, \forall t$$
 (5)

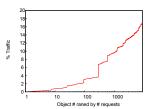
Constraint 2 ensures that all data was delivered within the deadline. Constraint 3 ensures that data is not delivered before it originally arrived. Constraint 4 ensures the total delayed data delivered is equal to original amount. Constraint 5 ensures x_t is a positive value.

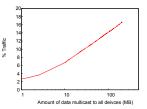
Figure 7 shows the results of solving the LP for the peak uplink rate. The original normalized peak capacity requirement was 2.14. With just a 5-minute delay tolerance, the peak is reduced by 25% to 1.6 (not shown), with 10-minutes it reduced by 30% to 1.49, at 30-minutes by 40.1% to 1.26. Finally at 60-minutes of delay tolerance, the peak reduced by 49.7% to 1.08. Thus, with even a relatively small amount of application delay tolerance, there can be a substantial reduction in the peak capacity requirement - significantly mitigating uplink congestion at these venues.

6.2 Can we exploit multicast in LTE?

Intuition would suggest that a venue would see considerable commonality in accessing content and temporal locality in accessing that content. Thus, using multicast should be effective in saving capacity.

A proposed technology for multicast for LTE is eMBMS. Evolved Multimedia Broadcast and Multicast Services (eMBMS) is used for transmission of the same content to multiple receivers.





(a) Object Rank

(b) Bytes Saved by Multicasting

Figure 8: Benefit of multicasting popular content and caching

It can be used for streaming live content or multicasting common content of interest to multiple users. It can be used across multiple cells, with what is known as a 'single frequency network' (SFN) where the transmission across cells is synchronous using a common carrier frequency. Having a single frequency network operation avoids inter-cell interference. eMBMS and unicast channels are multiplexed in the time domain, thus avoiding the long-term dedication of radio resources for multicast. Configuration of LTE cells participating in eMBMS (MBSFN areas) is determined by operations and management. The configuration of radio resources used for eMBMS can be determined dynamically, within a wide range (0.3% to 60%) of downlink unicast capacity.

Live content (e.g., live TV broadcast of the event to spectators) could be a candidate for eMBMS. But there is, on average, only one viewer of the live content per cell. The peak usage (measured over 1 minute time scales) was about 6 simultaneous viewers across all the cell sites provisioned at the stadium. Thus, one would conclude that there is not a significant benefit from multicasting the live TV broadcast to the participants at the Super Bowl.

In contrast to the viewers of live video, there were a large number of users (over 30K) accessing the network for information, with a peak number of simultaneous users of over 1K. A substantial amount of data was consumed by users at the event - especially web content. The question is: is there temporal locality in accessing this content and can we multicast it effectively? As one might expect, given the diversity of the content, there may not be enough web content simultaneously accessed to make multicast of such information to be effective either. There is a significant skew in popularity of information and content accessed by the users in the venue, as has been observed by others. Figure 8 (a) shows that the top 100 objects, ranked by number of requests, constitute a 2% share of the downstream bytes, and the top 1000 take up almost 10%. Previous studies have shown that, despite the lack of very tight temporal locality of access, using the end-system storage as a cache, can effectively relax the need for such tight temporal locality to deliver the information through multicast [10, 11]. By delivering the most popular content by multicast to the end-system caches (with a time to live lasting the length of the event at the venue), we can save a significant portion of the downstream bandwidth. It can also reduce the peakedness of the traffic on the downlink. Therefore, we explore how a small cache on the end-devices can save on downlink capacity, and make eMBMS-based multicast more effective. Figure 8 (b) shows that by multicasting popular content and having just a 10 Mbyte cache on the end-systems can save about 6.7% of the total downstream transfer, and a 100 Mbyte cache can save almost 14% of the transferred bytes (we need to suitably factor in applications such as "speed test", which download random bytes). In any case, we believe that multicast with eMBMS on the cellular network can be quite effective in reducing capacity, if used intelligently. It would require quickly identifying popular content and exploiting the cache on the end-systems to multicast and pre-place popular content.

7. CONCLUSIONS

Venues that host large events like the Super Bowl are a challenge to design and provision to meet user demand. Our analysis of the Super Bowl traffic indicated that with high-speed LTE there is significant traffic demand despite extensive deployment of in-stadium WiFi. Moreover, we found with the shift in application mix, uplink traffic is much more prominent (almost matching downlink). This is significant for the cellular service provider as it is more complex to provision and manage. We suggested an approach to tradeoff delay tolerance for mitigating congestion. While it is tempting to exploit cellular multicast (eMBMS), a straighforward implementation isn't enough. Our proposed approach to use end-system storage as a cache of popular content that is pre-placed through multicast can indeed help achieve the expected benefits of multicast. We also observed that video consumption, despite having only a small number of users, consumes a large share of the capacity. We believe understanding the role of providing a fair share of resources to meet the utility of all the users across the application spectrum, is important. Finally, we observed that usage changes significantly from one part of the stadium to another as the event progresses. Thus, it would be very appropriate to shift radio resources, as demand shifts from one eNodeB to another at the venue.

8. REFERENCES

- [1] AT&T Press Release. http://www.att.com/gen/press-room?pid= 4800&cdvn=news&newsarticleid=30527, February 2010.
- [2] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2010Ű2015.
 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html, February 2011
- [3] 3GPP Self-Organizing Networks. http://www.3gpp.com/SON, 2013.
- [4] Behind the Scenes Look at AT&TŠs Super Bowl XLVII
 Network Strategy. http://www.sporttechie.com/
 2013/04/21/behind-the-scenes-look-at
 -atts-super-bowl-xlvii-network-strategy/,
 April 2013.
- [5] From party to pigskin, our network is ready for this weekendŠs fanfare. http://blogs.att.net/ consumerblog/story/a7787203, February 2013.
- [6] Super Bowl plans to handle 30,000 Wi-Fi users at onceUnd sniff out 'rogue devices'. http://arstechnica.com/information-technology/2013/02/super-bowl-plans-to-handle-30000-wi-fi-users-at-once-and-sniff-out-rogue-devices/, February 2013.
- [7] T. M. T. Do, J. Blom, and D. Gatica-Perez. Smartphone usage in the wild: a large-scale analysis of applications and context. In *Proceedings of the 13th international conference* on multimodal interfaces, ICMI '11, pages 353–360, New York, NY, USA, 2011. ACM.

- [8] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin. A first look at traffic on smartphones. In Proceedings of the 10th ACM SIGCOMM conference on Internet measurement, IMC '10, pages 281–287, New York, NY, USA, 2010. ACM.
- [9] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, and D. Estrin. Diversity in smartphone usage. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, MobiSys '10, pages 179–194, New York, NY, USA, 2010. ACM.
- [10] V. Gopalakrishnan, B. Bhattacharjee, K. K. Ramakrishnan, R. Jana, and D. Srivastava. CPM: Adaptive VoD with Cooperative Peer Assist and Multicast. In *Infocom'09*, May 2008.
- [11] V. Gopalakrishnan, R. Jana, K. K. Ramakrishnan, V. Aggarwal, R. Calderbank, and F. Yu. The Effectiveness of Intelligent Scheduling for Multicast VoD. In ACM MM'09, October 2009.
- [12] G. Maier, F. Schneider, and A. Feldmann. A first look at mobile hand-held device traffic. In *Proceedings of the 11th* international conference on Passive and active measurement, PAM'10, pages 161–170, Berlin, Heidelberg, 2010. Springer-Verlag.
- [13] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang. A First Look at Cellular Network Performance during Crowded Events. In *Sigmetrics* '13, June 2013.
- [14] M. Z. Shafiq, L. Ji, A. X. Liu, and J. Wang. Characterizing and modeling internet traffic dynamics of cellular devices. In Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems, pages 305–316. ACM, 2011.
- [15] I. Trestian, S. Ranjan, A. Kuzmanovic, and A. Nucci. Measuring serendipity: connecting people, locations and interests in a mobile 3g network. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, IMC '09, pages 267–279, New York, NY, USA, 2009. ACM.
- [16] Q. Xu, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman. Identifying diverse usage behaviors of smartphone apps. In *Proceedings of the 2011 ACM* SIGCOMM conference on Internet measurement conference, IMC '11, pages 329–344, New York, NY, USA, 2011. ACM.