Characterizing Fairness for 3G Wireless Networks

Vaneet Aggarwal, Rittwik Jana, Jeffrey Pang, K. K. Ramakrishnan, N. K. Shankaranarayanan
AT&T Labs Research, Florham Park, NJ, USA

***Abstract*—The end to end system data performance over a 3G cellular network depends on many factors such as the number of users, interference, multipath propagation, radio resource management techniques as well as the interaction between these mechanisms and the transport protocol's flow and congestion mechanisms. Using controlled experiments in a public cell site, we investigate the interaction between TCP and the 3G UMTS/HSPA network's resource allocation, and its effect on fairness in the throughput achieved across multiple (up to 26) TCP flows in a loaded cell sector. Our field measurement results indicate that TCP fairness fluctuates significantly when the air interface (radio link) is the bottleneck. We also observe that TCP fairness is substantially better when the backhaul link (a fixed wired link) is the bottleneck, instead of the air interface. We speculate that the fairness of TCP flows is adversely impacted by the mismatch between the resource allocation mechanisms of TCP's flow and congestion control and that of the Radio Access Network (RAN).**

I. Introduction

There has been tremendous growth in the use of data over wireless cellular networks with the advent of smartphones over the last few years. As in wired networks, the dominant transport protocol used in cellular networks is TCP, which comprises over 95% of flows [1]. Even real-time streaming traffic, such as video, is more often transported over HTTP/TCP than UDP. Hence, there is significant interest in understanding how the allocation and scheduling of cellular resources impacts the behavior of TCP flows.

A major part of this interest is on the downlink in UMTS/HSPA cellular networks, i.e. from the base station to the user equipment (UE), such as a smartphone or a cellular data modem. Radio access network (RAN) resources are allocated to each UE in a highly complex and dynamic manner, taking into account the radio resources available, the signal strength observed by the end-device, outstanding data to other receivers, and other considerations. A primary notion is that there is a "channel" for each UE to which RAN resources are allocated (this is analogous to a virtual circuit-switched network's connection). Hence, recent work (e.g., [2], [3], [4]) has examined the interaction between TCP's flow and congestion control mechanisms, the RAN resource allocation methods, and the wireless channel. These studies observed negative performance impact on individual TCP flows – for example, poor efficiency due to link-layer retransmissions that result in variability in RTT. It is important to recall, however, that TCP has two major goals: efficiency and *fairness*. Previous work focused primarily on how RAN resource allocation impacts the efficiency of flows. Our contributions are observations on the fairness in the performance obtained for competing TCP flows based on an extensive set of measurements on an operational UMTS/HSPA cellular network. We characterize both temporal (long term time average) and spatial (distribution of throughout within a cell or sector) fairness.

Understanding how RAN resource allocation interacts with TCP is not straightforward due to its complexity and because it functions independently at the radio link control (RLC) layer, a protocol layer below IP. For instance, High Speed Downlink Packet Access (HSDPA) has fast scheduling and hybrid ARQ. The RAN resource allocation seeks to provide equitable proportioning of radio resources depending on the channel condition, user location, interference, and scheduling discipline. For each transmission time interval (TTI, typically 2 ms), the scheduler located in the Node B (base station), decides on the users to be scheduled and their corresponding data rates. This is based on the reported channel quality by the UEs, as well as fairness metrics. There is also a radio link control (RLC) flow control mechanism between the RNC and the node B. For the same traffic, we also have TCP's flow and congestion control, which operate on timescales of one or more end-to-end RTTs (typically 100-300 ms in cellular networks), while the HSDPA scheduler algorithms operate on finer timescales of the order of 2ms TTIs.

TCP's congestion control mechanism also seeks to achieve proportional fairness across connections [5], [6]. To evaluate how these two layers interact, we measure how TCP flows perform when the RAN is loaded with a controlled number of high-demand, long-lived flows. We vary the number of sources, although we expect an individual source to be capable of offering a load that is a high fraction of the air interface capacity of a single sector. We conducted measurements in two types of environments: (a) where we were confident that the air interface was the bottleneck and (b) where we were confident that the air interface was not the bottleneck.

Our contributions are as follows: First, we use real-life controlled experiments with an operational 3G UMTS/HSPA cellular network with multiple datacards and smartphones (total up to 26) to characterize the fairness across TCP flows as a function of increasing load and varying traffic dynamics across 1, 2, and 3 cell site sectors. Second, we investigate how fairness is impacted if the radio air interface is the bottleneck, and compare this with the case when the backhaul is the bottleneck. Our findings indicate that there is greater variability in fairness when the air-interface is congested as opposed to the backhaul being congested. We believe this is because TCP's congestion control takes precedence over the RAN resource allocation algorithms when the backhaul is the bottleneck, thus leading to better fairness.

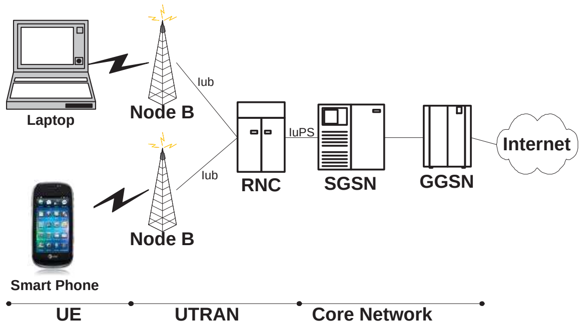

Fig. 1. UMTS network architecture

II. Related Work

The performance of TCP over wireless networks has received extensive study over the past two decades. In this section, we review related work that has studied TCP performance issues in 3G and 4G cellular data networks. Our main contribution over previous work is a better understanding of TCP fairness in an operational cellular network under load based on careful measurements with multiple UEs.

Chan and Ramjee [2] presented a parametric model to estimate TCP throughput that accounts for delay and rate variations inherent in 3G networks. They propose that an Ack Regulator deployed in RNCs could improve overall TCP throughput. Chan and Ramjee [7] later proposed a window regulator to account for channel variations and a heuristic PF to better account for a mixture of short and long flows. In both these works, they show in simulation that their proposed algorithms could improve overall TCP throughput but do not explicitly address the issue of fairness.

Halepovic et al. [8] conducted bulk TCP transfer experiments over an operational WiMAX network to compare the performance of different TCP variants. Claypool [9] conducted TCP tests over an operational 1xEVDO cellular network. Both these works observed peculiarities for TCP performance, but only utilized a single wireless device in their experiments. Thus, they did not study fairness under load. Tan and Lau [10] performed experiments on multiple operational 3G UMTS networks. They found that known theoretical models can not predict the actual capacity of a cell given empirical measurements for standard parameters. Our work complements theirs by showing that theoretical models for TCP fairness may not hold when subject to the complexity of the RLC protocol and wireless channel fluctuations.

III. System Description

As illustrated in Figure 1, the UMTS data network consists of three subsystems: User Equipments (UE), UMTS Terrestrial Radio Access Network (UTRAN), and the 'packet core' (also called the Core Network (CN)). UEs are user mobile handsets or laptops with 3G modems. The UTRAN represents network components between the UE and the CN. It interacts with the UE by a radio interface protocol stack that includes the physical layer, the MAC layer, the RLC (radio link control) layer, as well as a portion of the network layer. Most UTRAN features (part of the MAC layer, RLC and above) are implemented at the Radio Network Controller (RNC), while the base station (Node B) handles a limited set of functions of the MAC and the physical layer. The RNC controls the operation of multiple Node Bs, managing resources such as allocating capacity for data calls, and providing critical signaling functions (e.g., for call set-up), plus the switching and routing functionality. The SGSN provides routing, mobility/session management and user authentication and authorization. Finally, the UMTS core network is connected with an external packet-switched network such as the Internet via a gateway (GGSN). Data users typically access content along the following network path (UE-nodeB-RNC-SGSN-GGSN-Internet).

In 3G UMTS/HSPA networks, proportional fair (PF) scheduling is typically employed at the Node B to schedule downlink flows among different users. Specifically at each time slot, also known as TTI (typically 2 ms), the PF scheduler schedules the user with the largest $\frac{R_i}{A_i}$ where R_i is the instantaneous achievable rate by user i and A_i is the average rate of user i. The average (a moving average) is computed over a time window. The basis of PF scheduling is to achieve high overall cell throughput while at the same time maintaining proportional fairness among all users. Figures 2 and 3 show the flow control and the downstream and upstream traffic management model for a 3G UMTS environment.

3G services on a HSDPA network do not reserve a fixed amount of bandwidth on the Iub interface (backhaul link between Node B and RNC). A flow control algorithm [11] is implemented in the RNC and the Node B which ensures that the Iub interface is optimally used when it is the bottleneck. For the downstream link, the backpressure flow control algorithm working between the Node B and the RNC keeps the backhaul link congestion level low. The two possible cases for congestion in the RAN are a) constrained Iub interface or b) congested Uu (air interface, or radio link). For the latter, fair allocation of the radio resources depends on the vendor's scheduler specifics in the Node B (typically a proportional fair algorithm), and the flow control algorithm seeks to maintain the fairness imposed by the scheduler. However, if the Iub interface (backhaul) is the bottleneck, air-interface resources may be optimally utilized. Finally, one point to note is that the RNC controls the state of the UE which can switch between various idle and active states which affect performance [3]. For our experiments however, since we use long-lived TCP flows (source transmits as much as the window will allow), we believe that the state of the UE at the RNC will be in the Dedicated Channel (DCH) state which provides the highest level of performance.

IV. Experimental setup

In this section we detail our experimental setup. We conducted our experiments on a public 3-sector UMTS/HSPA cell site between midnight and 5AM on two different weeknights in a space of 2 weeks, when the site was likely to be relatively idle. By situating the devices close to the cell tower

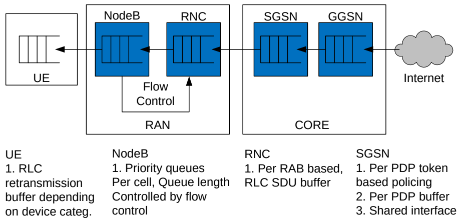

Fig. 2. Downstream traffic management

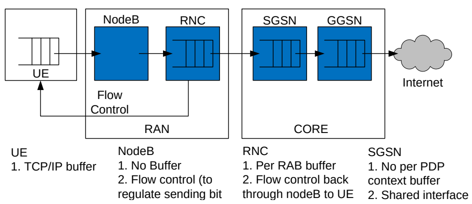

Fig. 3. Upstream traffic management

(approximately 300 feet from the base of the cell tower), we were able to secure a relatively strong signal strength in line of sight. We were careful to ensure that devices were connected to the particular target sector for the test scenario. We ensured that all the UEs were locked to a particular carrier frequency throughout the experiments. We ensured that the UEs were in the middle of each of the sectors we measured, and that the radio environments for each of the UEs was as similar as possible, while maintaining a half-wavelength separation between UEs. During each measurement, the UEs were stationary and in a vehicle with line of sight to the base antenna. The received signal strengths were all very strong. While we cannot avoid differences in radio conditions at each UE, the experimental setup is representative of different users who happen to be in the same place. The objective was to measure performance as experienced by typical users with typical UEs in a public network rather than a laboratory setting.

For our group of UEs, we had 16 UMTS/HSPA Sierra Wireless laptop USB data cards as well as 10 Samsung Galaxy-S Android smartphones. The datacards, smartphones, and network were all capable of HSDPA 7.2 Mb/s. The datacards and network were capable of HSPA 2.0 Mb/s, while the smartphones were capable of UMTS 384 kb/s uplink. The uplink rate on the smartphone was high enough to avoid any constraints on our downlink-based multi-user TCP performance tests. For ease of operations, each of our 4 test laptops was running four Linux virtual machines (VM) and each VM was connected to a dedicated data card. We took care to ensure that the use of VMs did not skew or bias our observations with respect to fairness. Our approach was to perform tests by loading a single sector first, followed by two sectors and finally all three sectors. Note that while carrier locking is possible on the UEs, it is not possible to prevent ordinary UEs from handing off to an adjacent sector. In each experiment, we ran a TCP throughput test (downlink) that lasted 300 seconds and checked that the device was on the same sector at the beginning and end of the experiment. It is unlikely that the device would have been handed-off to the other sector in the middle of the experiment. All devices were synchronized to start at the same time and the results were logged on the device at a one-second granularity.

We developed custom tools based on Iperf to inject load on the downlink. Since all tests lasted for a duration of 300 seconds, they were long enough for the (300-second) average TCP throughput measurements to overcome the initial effects which include TCP's slow start protocol. To capture the transient behavior, we also looked at the throughput achieved over each 1 second time interval. Iperf is a commonly used networking tool that can create both TCP and UDP data streams for measuring the throughput of the network. The Iperf code was ported to run on the Linux VMs. We wrote a custom application similar to Iperf to run on the Samsung Android phones.

V. Experiment Results

We present measurement results that highlight the temporal and spatial fairness of the TCP throughput achieved across long flows from a server (on the wired network) to wireless end-devices on one or more sectors at the cell site. We mostly present results from the 16 datacard measurements and supplement them with results from the 10 smartphones. Scenario 1 represent our first set of results which comprises of Scenarios 1a, 1b, and 1c with traffic loading for 1, 2, and 3 sectors respectively. In Scenario 1a, we first characterize how fairness varies as the number of UEs in one sector increases; in this scenario, the air-interface (radio link) is the performance bottleneck. Next, we show how fairness improves as we load 2 sectors (Scenario 1b) and all 3 sectors (Scenario 1c). In Scenario 1c, the total load is high enough to congest the backhaul capacity, which was dimensioned to be less than the sum of the 3 radio sector capacities; in this scenario, the air interface is not congested, and thus the backhaul is the bottleneck. Then, in Scenarios 2a, 2b, 2c, we show results when the backhaul link is considerably over-provisioned so that it is no longer the bottleneck. There is high variability in fairness when the air-interface is the bottleneck.

We compute fairness using the familiar square root of the average square deviation from average [12]. The fairness index is dimensionless, conveying that a truly fair system will have zero deviation from the average and the index will be zero. In our definition of fairness, the fairness index increases as the system becomes more unfair. We choose this metric as it is sensitive (in terms of the actual value of the index) to the level of unfairness, in comparison to the fairness index proposed by [13]. If a system allocates resources to K contending UEs such that the kth UE receives an allocation x_{nk}, then this fairness index $f_I(x)$ is given by

$$f_I(x) = \frac{\sum_{n=1}^{N} \frac{1}{\bar{x}_n} \sqrt{\sum_{k=1}^{K} (x_{nk} - \bar{x})^2 / K}}{N} \quad (1)$$

where x_{nk} denotes the kth sample point at time n. N is the total number of time points and K is the number of samples at time n (e.g. typically $N = 300$ for a 300 second experiment, K = # of devices participating in an experiment).

Scenario 1a (single sector, air interface bottleneck): Figure 4 shows the average fairness index over 300 seconds v. the number of UEs in Scenario 1a. The fairness decreases (higher values of fairness index) as the number of UEs in a cell sector increases. Since the UEs have long-lived high-demand flows, and the network does not artificially limit the peak rate to the UEs, the air-interface can become congested as the number of heavy users grows beyond a small number (about 3 to 4). While some level of difference in throughput may be expected (e.g. if the UEs face variations in radio conditions), the dynamic variation in fairness over the 300 seconds is of particular interest. This information is not easily available without controlled tests such as ours. Figure 5 shows the fairness index plotted against time, parameterized by the number of UEs. As the number of UEs sharing the sector increases from 3 to 7 to eventually 15 active "heavy" users, the unfairness increases dramatically. In fact, in the 15-UE case, we observed the fairness index goes up to 1.5 at certain times (e.g., at 250 seconds). At this point, the difference between the throughputs (averaged over a 1 second interval) of the TCP flows is more than a factor of 5. We note that when things become unfair for a particular UE, they stay unfair for many seconds. The RAN PF scheduler operates at a time scale of 2 ms each, and seeks to achieve fairness over a time window relatively short when compared to TCP, which also seeks to provide fairness over a multiple of Round Trip Times (RTT)s. We expect that over a time window of one second both of these mechanisms should achieve a reasonable level of fairness, and certainly over the length of the 300 second experiment. Thus we believe that the throughput across all competing TCP flows should be fair.

We also believe that the RAN PF scheduler treats each type of device somewhat differently, and therefore examined the fairness in the throughput achieved across each class of device. To illustrate the range of throughputs, Table I shows the average throughput achieved across the UEs during another experiment, where we had a total of 26 flows, with the 16 datacards and the 10 Android phones. The first column for Scenario 1a shows the average throughput achieved by the 16 laptop datacard UEs (over the 300 seconds), and there is considerable variability in their throughput (almost a factor of 2). Similarly, across the phones in Scenario 1a, there is variability in the average throughput achieved across the phones (almost 25% difference between the lowest and the highest). As such, our observation is that the overall fairness achieved when the air interface is bottlenecked is poor, and the interaction between the RAN PF scheduler and the TCP congestion control mechanisms is less than ideal. We speculate that there is a mismatch between the time scales over which TCP (acting on an end-end basis) and the cellular PF scheduler at the Node B (operating at 2 msec timescale) allocate resources across the flows. In addition, loss and variable delay further contribute to the unfairness across flows. We seek to understand this detail in our current work by inspecting Wireshark traces of the collected measurements reported here.

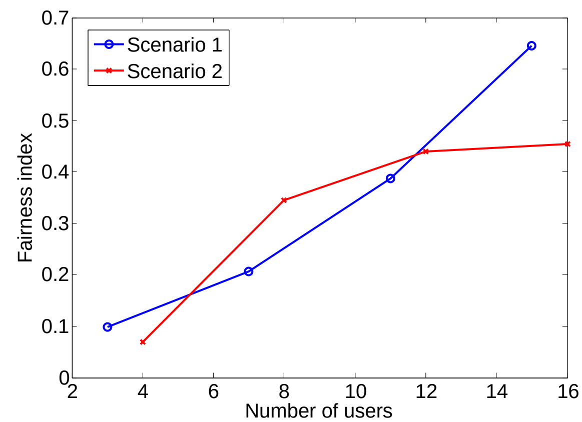

Fig. 4. Average Fairness index v. no. of UEs, Scenarios 1a and 2a, 1 sector, air interface bottleneck (lower index = better fairness)

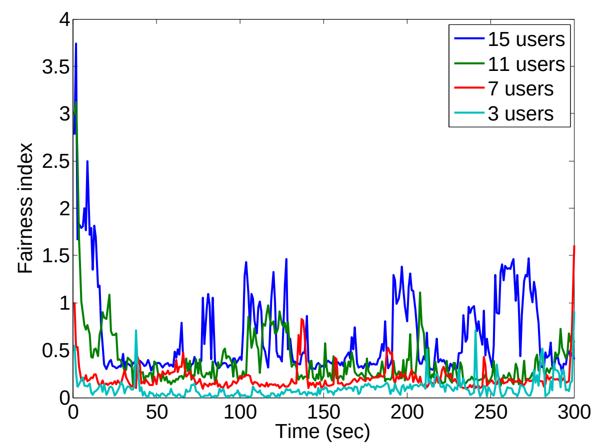

Fig. 5. Fairness index v. time for Scenario 1a (air-interface bottleneck)

Scenario 1b (2 sectors) and Scenario 1c (3 sectors, backhaul bottleneck): We then distributed the 16 UEs across 2 and 3 sectors of the cell site to examine the impact on fairness across the flows. With load across all 3 sectors, we were confident from the results that we exceeded the capacity of the total backhaul. With the backhaul link being the bottleneck (and experiencing congestion), we had an opportunity to examine the fairness across the TCP flows. To highlight the spatial fairness, we estimated the overall average fairness achieved across all the TCP flows for the maximum user case of 16 UEs. The first column of Table II shows that average fairness improves (from 0.65 to 0.40 to 0.16) as we load 1, 2, and 3 sectors in Scenarios 1a, 1b, and 1c. The standard deviation of the fairness also improves going from 1 sector to 3 sectors. Figure 6 shows the fairness index v. time for Scenarios 1a, 1b, and 1c with 16 UEs distributed in 1, 2, or 3 sectors respectively. We note that there is significantly better

TABLE I
AVERAGE THROUGHPUTS(KBPS) OF EACH UE FOR RADIO BOTTLENECK AND BACKHAUL BOTTLENECK CONDITIONS, SCENARIOS 1A AND 1C

Scenario 1a 1 sector(radio bottleneck)		Scenario 1c 3 sectors(backhaul bottleneck)	
Datacard	Phone	Datacard	Phone
144	205	268	243
146	209	275	260
153	213	280	271
167	216	280	279
173	222	281	280
181	227	281	281
185	234	283	286
190	241	287	292
190	243	288	292
191	256	291	296
192		292	300
206		295	
208		295	
223		296	
278		299	
283			

TABLE II
SPATIAL FAIRNESS FOR ALL UES IN 1, 2, AND 3 SECTORS: f_I IN EQUATION 1, AND STANDARD DEVIATION = $\frac{1}{\bar{x}_n}\sqrt{\sum_{k=1}^{K}(x_{nk}-\bar{x})^2/K}$

Scenario	Scenario 1a,1b,1c		Scenario 2a,2b,2c	
# sectors	f_I	Std(f_I)	f_I	Std(f_I)
1	0.65	0.46	0.46	0.3
2	0.40	0.44	0.36	0.1
3	0.16	0.01	0.56	0.1

fairness and there is a lot less fluctuation as we go from one sector to two and finally to three sectors. This is because we begin to congest the backhaul link.

We can observe the throughputs per UE plotted against time in Figure 7) for each of the 16 datacard UEs in Scenario 1c. Both the PF scheduler as well as the TCP congestion control mechanism should be treating these flows (to the same class device) fairly. We see that all flows achieve nearly identical throughput, when all three sectors are loaded. For another supplementary experiment with 15 laptop datacards and 11 Android phones distributed across the 3 sectors, the average throughputs over 300 seconds for the 26 UEs are tabulated in the second column of Table I (backhaul bottleneck.) It is interesting to note from the table that not only are the throughputs of the 15 datacard TCP flows nearly identical in Scenario 1a, but also the throughput achieved by the phones is also nearly the same. We believe that the primary resource allocation is determined by the backhaul link which is the bottleneck in this case. Thus, there is no adverse interaction between the RAN PF scheduler and TCP, and TCP divides the backhaul bottleneck capacity fairly across all flows (irrespective of even the device type).

Scenario 2 (over-provisioned backhaul): To verify that the unfairness observed was a result of the air interface being the bottleneck, we repeated the same set of measurements from the same points in the various sectors on a different night when the backhaul link capacity to the cell site was substantially

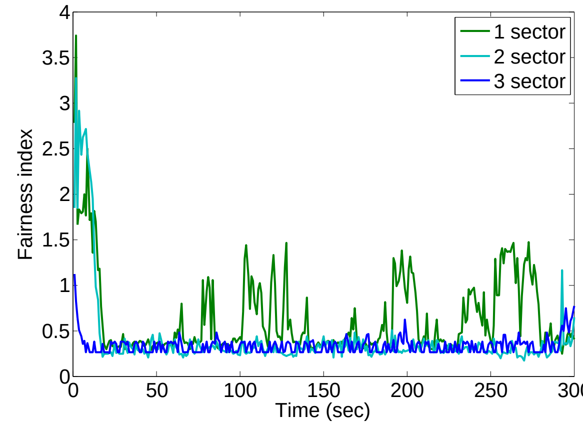

Fig. 6. Fairness index v. time for Scenarios 1a,1b,1c, 16 UEs

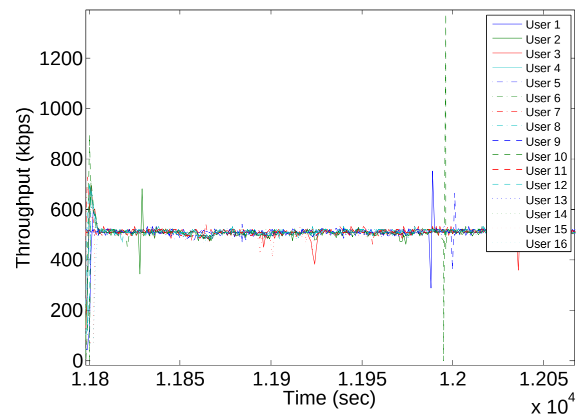

Fig. 7. Throughput v. time for Scenario 1c, 3 sectors, 16 UEs

higher and over-provisioned. We observed that the TCP results with the single-sector Scenario 2a were essentially similar to Scenario 1a. To avoid being repetitive, we do not include detailed results beyond the fairness index v. UEs in Figure 4 and the average fairness in Table II.

We now increase the load by distributing the 16 datacard UEs across 1, 2, and 3 sectors. Since the backhaul was over-provisioned, the air interface was the bottleneck independent of the number of sectors loaded. Unlike the trend from Scenario 1a to Scenario 1c when the backhaul became the bottleneck, with an over-provisioned backhaul, fairness in fact generally worsens (going from 0.46 for 1 sector to 0.56 with 3 sectors, from Table II). We speculate that when the air interface is the constraint, the radio network's internal flow control and PF scheduling algorithm operating over a much finer time scale (every 2 msec) interacts with the TCP congestion control mechanism in an undesirable manner. The dynamics over time are evident in Figure 8 which shows the fairness vs. time for Scenarios 2a, 2b, 2c with 1, 2, 3 sectors respectively. We observe that the fairness index (note the baseline level) goes up as we load up more sectors (indicating more unfairness) while the variability decreases. This is in direct contrast to Scenario 1 as shown in Figure 6.

Average throughputs in kbps of 16 UEs in Fig. 9															
251	411	458	476	490	498	548	548	550	783	802	816	1359	1378	1383	1392

TABLE III
UNFAIRNESS IN THE AVERAGE THROUGHPUT ACROSS 16 UEs IN SCENARIO 2C

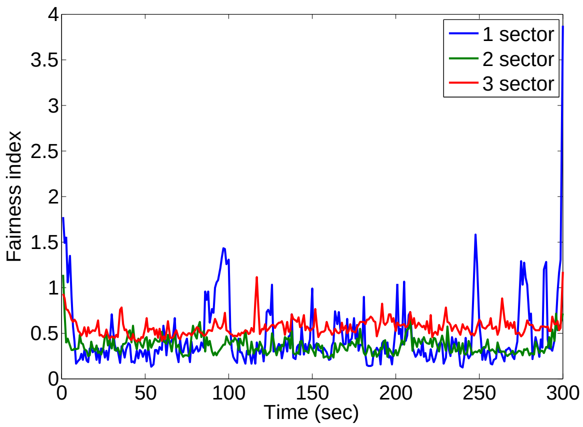

Fig. 8. Fairness v. time for Scenario 2a,2b,2c (over-provisioned backhaul), 16 UEs

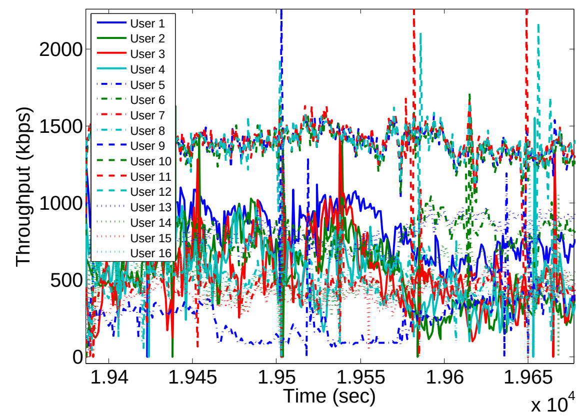

Fig. 9. Throughput v. time for Scenario 2c, 3 sectors, 16 UEs

We can observe the throughputs per UE plotted against time in Figure 9 for each of the 16 datacard UEs in Scenario 2c. We observe the dramatic variability in the throughput observed across the flows. We do observe that the throughputs of flows for one sector are clustered together, but there is significant differences across flows in the other two sectors. There are also considerable differences between the throughputs of flows across sectors (overall, there is a range from 1.5 Mbps to about 250 Kbps, a factor of 6 difference). This difference observed in the time series is eventually reflected in the average throughput (over the 300 second experiment) being also quite different (251 Kbps to 1392 Kbps, see Table III).

Contrasting the results observed between Scenario 1c (3 sectors, backhaul limited) and Scenario 2c (3 sectors, air interface limited), we conclude that there is a need for improvement in the resource allocation mechanism in the air interface and the end-to-end TCP flow and congestion control mechanisms. We are continuing to examine more detailed information, and anticipate further work (both in terms of measurement and analysis) to more precisely understand the cause of this unfairness.

VI. CONCLUSIONS

We have conducted controlled experiments in an operational 3G UMTS/HSPA cell site by injecting multiple (up to 26) long-lived high-demand TCP traffic flows to study the amount of unfairness in the observed end-to-end throughputs. We show that unfairness increases as the number of high-demand users within a sector increases. Our results also indicate that TCP's fairness fluctuates significantly when the air interface is the bottleneck. And, under similar conditions at the same cell site, the performance of TCP in terms of fairness is substantially better when the backhaul link (a fixed wired link) is the bottleneck instead of the air interface. We speculate that the fairness of TCP flows is adversely impacted by the mismatch between the resource allocation mechanisms of TCP's flow and congestion control and that of the radio access network (RAN).

REFERENCES

[1] A. Gerber, J. Pang, O. Spatscheck, and S. Venkataraman, *Speed testing without speed tests: estimating achievable download speed from passive measurements*, Proc. of IMC, 2010.
[2] M. C. Chan and R. Ramjee, *TCP/IP Performance over 3G Wireless Links with Rate and Delay Variation*, Proc. of MOBICOM, 2002.
[3] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen and O. Spatscheck, *TOP: Tail Optimization Protocol for Cellular Radio Resource Allocation*, 18th IEEE International Conference on Network Protocols, Kyoto, Japan, October 2010.
[4] K. Mattar, A. Sridharan, H. Zang, I. Matta, A. Bestavros, *TCP Over CDMA2000 Networks: A Cross-Layer Measurement Study*, Proc. of Passive and Active Measurement Conference, April 2007.
[5] V. Jacobson, *Congestion avoidance and control*, Proc. of SIGCOMM, 1988, pp.314-329.
[6] F. Kelly, *Fairness and stability of end-to-end congestion control*, European Journal of Control 9, 2003, pp.159-176.
[7] M. C. Chan and R. Ramjee, *Improving TCP/IP Performance over Third Generation Wireless Networks*, Proc. of IEEE INFOCOM, 2004.
[8] E. Halepovic, Q. Wu, C. Williamson, and M. Ghaderi, *TCP over WiMAX: A Measurement Study*, Proc. of IEEE MASCOTS, 2008.
[9] K. Claypool, L. Lee and G. Ratner, *Characterization by measurement of a CDMA 1xEVDO network*, Proc. Of Wireless Internet, 2006.
[10] W. L. Tan, F. Lam and W. C. Lau, *An Empirical Study on 3G Network Capacity and Performance*, Proc. of INFOCOM, 2007.
[11] 3GPP, TR 25.902 version 6.0.0 (UMTS Iub/Iur congestion control).
[12] P. L.Palyi, S. Racz and S. Nadas,*Fairness-Optimal Initial Shaping Rate for HSDPA Transport Network*, Proc. of ICCS, 2008.
[13] R. Jain, D.M. Chiu, and W. Hawe,*A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems*, DEC Research Report TR-301, 1984.