Aalborg Universitet

AALBORG UNIVERSITY

System Level Analysis of Dynamic User-Centric Scheduling for a Flexible 5G Design

Pedersen, Klaus I.; Niparko, Maciej; Steiner, Jens; Oszmianski, Jakub; Mudolo, Luke; Khosravirad, Saeed R.

System Level Analysis of Dynamic User-Centric Scheduling for a Flexible 5G Design

Klaus I. Pedersen[1,2], Maciej Niparko[1], Jens Steiner[1], Jakub Oszmianski[1], Luke Mudolo[1], Saeed R. Khosravirad[1]

Nokia – Bell Labs[1], Aalborg University[2]

Abstract— **In this paper we present our latest findings on dynamic user-centric scheduling for a flexible 5G radio design, capable of serving users with highly diverse QoS requirements. The benefits of being able to schedule users with different transmission time intervals (TTIs) are demonstrated, in combination with a user-centric multiplexing of control and data channels. The proposed solution overcomes some of the shortcomings of LTE-Advanced in terms of scheduling flexibility and performance. In general it is found that using short TTIs is advantageous at low to medium offered traffic loads for TCP download to faster overcome the slow start phase, while at higher offered traffic loads the best performance is achieved with longer TTIs. Using longer TTI sizes results in less control overhead (from scheduling grants), and therefore higher spectral efficiency. The presented analysis leads to the conclusion that a future 5G design shall include support for dynamic scheduling with different TTI sizes to achieve the best performance.**

I. INTRODUCTION

Research towards a new 5G air interface is gaining further momentum with the recent ramp-up of such activities in 3GPP [1]. The ambitions for 5G are high, aiming for support of a large variety of diverse services and deployment use cases. The considered services range from enhanced mobile broadband (eMBB), over massive machine type of communication (mMTC), and towards ultra-reliable low latency communication (URLLC) [1]-[3]. Simultaneously fulfilling the requirements for a mixture of users with such diverse requirements on the same radio interface is a rather challenging task, given the fundamental tradeoffs between capacity, latency, and reliability on a time-variant wireless channel [4]. As an example, the effective capacity expresses the maximum source data arrival rate that a certain channel process can support, while fulfilling a latency constraint [5]. With no latency constraints, the effective capacity approaches the Shannon capacity, while it decreases with stricter latency constraints. A new 5G air interface design must therefore include a flexible, and highly dynamic, resource allocation framework among the user equipment (UE) entities per cell, in order to best serve them in coherence with their individual quality of service (QoS) constraints.

In this paper we further study the performance of an enhanced scheduling framework for 5G (a.k.a. per-user radio resource allocation). Our focus is on the downlink performance, but some uplink dependencies are also taken into account. We build on the recent study in [6] that offers a flexible frame-structure for dynamic scheduling of users with different transmission time intervals (TTIs). The primary objective of our study is to further validate the hypothesis that support for variable TTI sizes is desirable for 5G. In our effort towards this objective, we analyze the system level performance of a multi-cell, multi-user scenario. As will be shown, the performance depends on many factors, such as the offered traffic load, the inter-cell interference in the system, the related traffic model (and associated payload sizes), etc. The effect of the overhead from having to transmit dynamic scheduling grants to users is taken explicitly into account in our analysis. Given the complexity of the considered problem, we sought to dynamic system level simulations, using proven and commonly accepted models and methodologies in order to obtain results with high degree of realism. Care is taken to generate statistical reliable results that can be used to draw mature conclusions.

The rest of the paper is organized as follows: Section II outlines the proposed dynamic scheduler design, including assumptions for multiplexing of users, scheduling formats, control channels for signaling of scheduling grants, etc. Section III presents the undertaken evaluation methodology, based on state-of-the-art dynamic system level simulations with high degree of realism. The corresponding performance analysis is presented in Section IV, followed by supplementary discussions to put the results into further perspective in Section V. The paper is closed with concluding remarks in Section VI.

II. DYNAMIC SCHEDULING FRAMEWORK

A. Fundamental user multiplexing

The study builds on the frame structure proposed in [6], which allows orthogonal time-frequency multiplexing of users per cell as illustrated in Fig. 1. Users are multiplexed on a grid of time-frequency tiles, corresponding to a minimum time-duration of one subframe, and one physical resource block (PRB) in the frequency domain. Whenever a user is scheduled by the base station, it is informed through a corresponding scheduling grant. The scheduling grant is sent on the so-called in-resource control channel (CCH) that appears at the start of the transmission for the user (marked with dark blue in Fig. 1). The CCH conveys relevant scheduling and link adaptation information such as the time-frequency resource allocation for the user, the used modulation and coding scheme (MCS), hybrid automatic repeat request (HARQ) information, etc. It should be noted here that the multiplexing of control and data channels is much more flexible for the considered 5G case, as compared to current LTE. Notice that the proposed 5G solution adopts a *user-centric* design paradigm, where control and data are multiplexed within resources per user. LTE relies on a *cell-centric* design approach, where control and data are only time-multiplexed within each subframe per cell [7]-[9]. Resulting in a commonly shared resource pool for LTE CCH scheduling grants to all users in a cell.
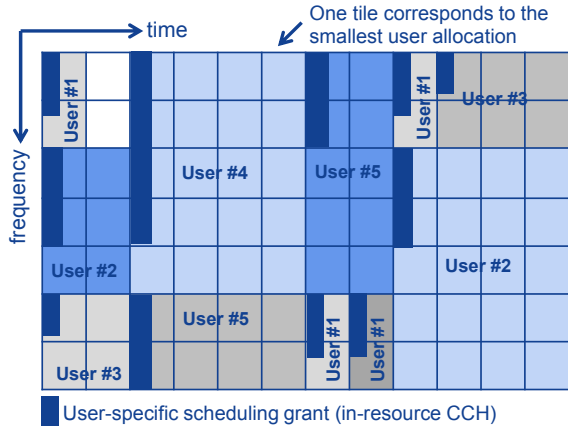
Fig. 1: Sketch of flexible time-frequency mux of users.

The minimum TTI size that a user can be scheduled with, corresponds to one subframe in this study. Scheduling users with a short TTI is attractive from a latency perspective, but comes at a cost of higher relative CCH overhead. Contrary, scheduling a user with a longer TTI size corresponding to an integer number of subframes reduces the relative CCH overhead at the expense of higher latency. As discussed in [6], this framework offers the possibility to adjust the resource allocation per link in coherence with the individual user's service requirements (e.g. operating eMBB users with low CCH overhead, while scheduling URLLC users with short TTIs and higher CCH overhead).

B. Scheduling grant adaptation

The in-resource CCH carries nearly the same information as the LTE physical downlink control channel (PDCCH) [7]-[9], and the corresponding enhanced PDCCH (E-PDCCH) [10]-[11]. Given that observation, we assume that the in-resource CCH for 5G will have comparable link level performance and type of link adaptation (LA) functionality. That means we assume QPSK modulation for the in-resource CCH, tail-biting convolution encoding, using a minimum of 36 resource elements (REs) for transmission. One RE corresponds to one subcarrier symbol. Depending on the user's experienced channel quality, additional repetition encoding is applied in the form of aggregation levels 2, 4, or 8. Hence, the base station will monitor the channel quality indicator (CQI) feedback from the UEs, and adjust the resources for the in-resource CCH – just as is the case for PDCCH link adaption functionality in LTE [7]-[9]. Table 1 summarizes the required number of REs for the in-resource CCH depending on the experienced signal-to-interference-and-noise-ratio (SINR) to achieve a block error rate (BLER) of less than 1%. The SINR depends on both the user's distance from its serving cell, as well as the level of other-cell interference that depends on the traffic load. Referring to Fig. 1, the use of LA for the CCH results in different amount of used radio resources for transmission of scheduling grants for the users.

Table 1: LA table for the CCH carrying the scheduling grant.

SINR threshold	Scheduling grant CCH overhead
-5 dB	8x36 = 288 REs
-2.2 dB	4x36 = 144 REs
0.2 dB	2x36 = 72 REs
4.2 dB	1x36 = 36 REs

C. Scheduling format

Next we consider the possible scheduling formats that shall be supported. The resource allocation to a user should be sufficiently large to contain both the in-resource CCH as well as a reasonable size payload and reference symbol overhead. The payload corresponds to at least one medium access (MAC) packet data unit (PDU), containing data from higher layers, MAC header, and potentially also MAC control elements. As per the results in Table 1, the in-resource control channel may occupy up to 288 REs in worst case, and a reasonable assumption for reference symbol overhead is on the order of 10% for 2x2 MIMO [10]. Given these conditions, the proposed minimum scheduling formats are summarized in Table 2 for different TTI sizes. As the physical (PHY) layer numerology for 5G have not yet been fully decided, we use the LTE numerology as a reference here. That is, 15 kHz subcarrier spacing, 14 symbols per 1 ms, and a PRB size of 12 subcarriers corresponding to 180 kHz. For a more exhaustive study on 5G PHY numerology options we refer to [12], and also the examples in [6], where settings allowing TTI sizes of 0.2 ms and 0.25 ms are presented.

The results in Table 2 shows that the assumed minimum frequency domain allocation for a user scheduled with a short TTI size of only 0.14 ms (corresponding to 2 OFDM symbols) is 1.4 MHz in order to be able to carry both data and control. That essentially limits the ability to gain from radio channel-aware frequency domain scheduling. On the other extreme, scheduling a user with a longer TTI allows allocation of less frequency domain resources (i.e. subcarriers). Thus, a narrowband low cost mMTC device could be scheduled with a 4 ms TTI size on only a single PRB, while a user with URLLC is best served with a short TTI on a larger bandwidth to meet the latency requirements. Notice from Table 2 that scheduling a user with the short TTI size of 0.14 ms on single frequency domain block of 8-PRBs is only possible for CCH aggregation level up to 4. Hence, a user that requires CCH aggregation level 8 will have to be scheduled on at least two contiguous frequency domain blocks of 8-PRBs.

Table 2: Scheduling formats when assuming LTE PHY numerology.

TTI size	Frequency domain scheduling block size (subband size)	Resource elements (REs) per block size
0.14 ms	8 PRBs (1440 kHz)	192
0.5 ms	4 PRBs (720 kHz)	336
1.0 ms	3 PRBs (540 kHz)	432
2.0 ms	2 PRBs (360 kHz)	576
4.0 ms	1 PRB (180 kHz)	576

D. Interaction with transport layer

The optimal TTI size and scheduling algorithm is highly application and transport layer dependent. In this study, we

analyze the performance for the internet transport protocol (TCP), including the effects of the end-to-end flow control mechanism as applied for this protocol. That means including the well-known slow start TCP procedure. In line with [6], our hypothesis is that during the initial data transmission session, the end-user-experienced performance is primarily determined by the RTT due to the slow start. Therefore, it would be advantageous to first perform scheduling with short TTIs, followed by longer TTI sizes when reaching steady state operation. Hence, for large file size eMBB users, using a long TTI size is generally expected to be beneficial, while small file size downloads are best served with short TTIs as those are dominated by the slow start TCP procedure. In this study we validate this hypothesis, using the generally accepted TCP Reno model [13]. On a related note, the study in [14] finds that use of queue maximum weight (Q-MW) scheduling, tailored specifically to TCP dynamics, is attractive by giving higher priority to TCP flows whose queue at the base station is very small in order to encourage data at a faster rate.

III. EVALUATION METHODOLOGY AND KPIS

A. Methodology and default assumptions.

Extensive dynamic system-level simulations are conducted, following the methodology typically used in 3GPP. The default simulation parameters are summarized in Table 3. The time-resolution of the simulator is on OFDM symbols. Whenever a user is scheduled, the SINR at the receiver is calculated for each subcarrier symbol, followed by mapping to the mutual information effective SINR metric (MIESM) that is used to determine if the transmission is correctly decoded [15]. Dynamic link adaptation for both data channel transmissions and the in-resource CCH (Table 1) is assumed. Closed-loop single-user MIMO with 2x2 is assumed in the downlink. The listed transmission formats in Table 2 are supported. The default packet scheduling algorithm is proportional fair, as we are mainly interested in studying the effect of using different TTI sizes. Asynchronous HARQ with soft combining is modelled for erroneously decoded packets. A dynamic birth-death traffic model is applied, where the user arrival is according to a homogeneous Poisson process with arrival rate λ. Users connect to the cell corresponding to the highest received power. There is a finite payload of B bits for each call for the downlink, leading to termination of the call when successfully delivered. The offered downlink traffic equals $\lambda \cdot B$. Modelling of TCP follows the Reno model [13]. When a TCP packet (with maximum segment size – MSS - of 1500B) is generated at the traffic source, it is subject to a core network (CN) latency of 2 ms before arriving at the base station. The corresponding TCP acknowledgement (Ack) from traffic sink (UE) in the uplink is transmitted with the same TTI size as in the downlink. Conveying the TCP Ack from the base station to the traffic source is again subject to the CN latency.

B. Key performance indicators (KPIs)

The primary performance metric is the end-user experienced data rate, calculated as the average experienced data rate for each user to download the file size. Based on user-throughput samples from a large number of users, the empirical cumulative distribution function (cdf) is created. In order to gain further insight for the TCP performance, the round trip time (RTT) of TCP packets is monitored, defined as the time from the server generates the TCP packet until the corresponding Ack is received. In line with the definition in RFC6298, the smoothed RTT is reported. Finally, the cost of scheduling users in terms of the in-resource CCH overhead is monitored. The relative CCH overhead is defined as the number of REs for the CCH as compared to the total amount of allocated REs for the user per scheduling instant.

Table 3: Summary of default simulation assumptions.

Description	Assumption
Environment	3GPP Urban Macro (UMa); 3-sector base stations with 500 meters inter-site distance. 21 cells.
Carrier	10 MHz carrier bandwidth at 2 GHz
Numerology	TTI sizes: 0.14 ms, 0.5 ms, 1 ms. Other PHY numerology settings in line with LTE.
Scheduling grant	In-resource control channel (CCH) scheduling grants with dynamic link adaptation as in Table 1.
MIMO	Single-user 2x2 closed loop MIMO with dynamic rank adaptation.
CSI	Periodic CSI every 5 ms, including LTE-alike CQI, PMI, and RI.
Data channel MCS	QPSK to 64QAM, with same encoding rates as specified for LTE. BLER target for first transmissions: 10%
Reference scheduler	Frequency-domain proportional fair (PF)
HARQ	Asynchronous HARQ with soft combining. Maximum 4 HARQ retransmissions Minimum HARQ retransmission delay: 8 TTIs
RLC	RLC Acknowledge Mode (AM) Maximum five retransmissions
Basic traffic model	Poisson arrival process Finite buffer file size (50 kB, 500 kB file size)
Transport layer	TCP Reno model, RFC 5681 TCP MSS: 1500B Initial TCP Window: 3xMSS SSThreshold: 45xMSS=67.5kB One-way core network delay: 2 ms

IV. PERFORMANCE ANALYSIS

A. Scheduling overhead under full buffer conditions

The first set of performance results in Fig. 2 shows the cdf of the experienced CCH overhead. These results are reported for a simple full buffer model in order to illustrate the cost of dynamic scheduling. As the simulated network consists of 21 cells, the cases with 21 and 210 users corresponds to having on average 1 and 10 users per cell, respectively. In line with our expectations, the relative CCH overhead increases with the number of average users per cell, as well as when shortening the TTI size.

The results in Fig. 2 clearly shows how the in-resource CCH overhead smoothly scales in coherence with associated scheduling demands such as number of users, TTI size, etc. This is one of the advantages of the considered 5G concept [6] due to the flexible division between CCH and Data channel resources. LTE have a much more rigid time-division multiplexing between CCH and Data, leading to scheduling

CCH overhead of either 7%, 14% or 21% for the PDCCH [10]. In fact, the limited flexibility of LTE is reported to result in PDCCH blocking in [9] under realistic traffic and QoS conditions, meaning that radio resource are left unused due to lack of resources for sending scheduling grants. In a related LTE-Advanced study [11], non-negligible CCH blocking for the E-PDCCH was also observed, although lower than for the baseline PDCCH. For the proposed 5G design with the in-resource CCH, the problems of CCH blocking (as known from LTE) are significantly reduced.
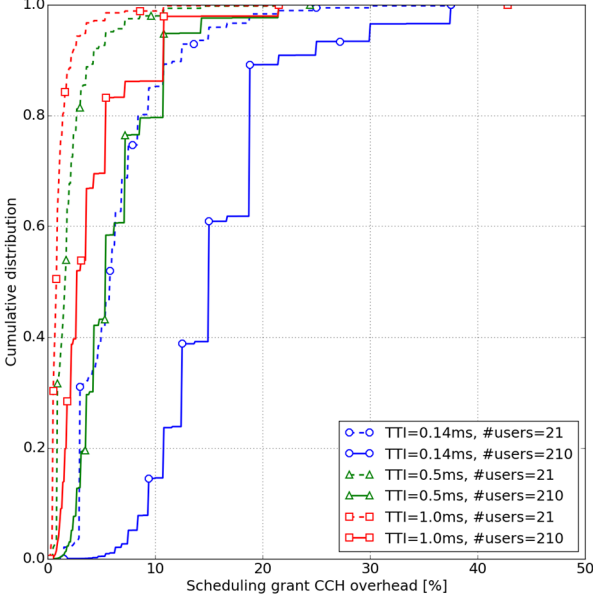

Fig. 2: Cdf of the relative overhead of in-resource CCH transmissions for full buffer cases.

B. Performance for download of small file sizes

For the remaining results of the paper, the default dynamic traffic model with TCP is assumed. The results presented next are for the small file size cases (50kB). Results are presented for different offered load levels, and hence different levels of inter-cell interference in the network. Fig. 3 shows the cdf of the experienced end-user throughput, including cases with different TTI sizes. At the low offered load, there is approximately 50% gain in the median experienced end-user throughput from using the smallest TTI sizes. The gain originates from the lower latency of using a short TTI that is especially advantageous for the small file sizes where the slow start TCP is dominant. Secondly, due to the low offered load, users experience only modest level of inter-cell interference, as hence the overhead from the in-resource CCH scheduling overhead is not that significant. On the other hand, at the high offered loads, the results in Fig. 3 shows a reduced benefit from using short TTI size, and in 50% of the cases worse performance than using the long TTI size. The latter is mainly due to the following two factors: (i) higher level of interference causing a need for more resources for in-resource CCH, and (ii)

at high offered traffic momentary queuing delays at the base station also start to occur.
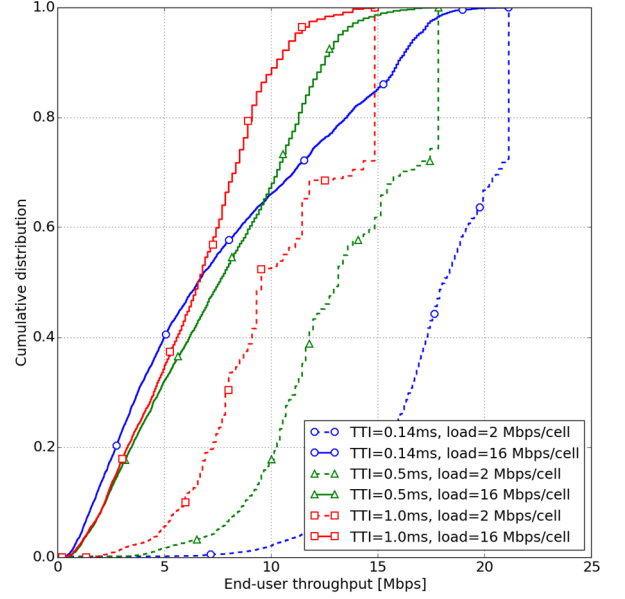

Fig. 3: Cdf of end-user throughput for download of small file sizes (50kB) under different load conditions and TTI sizes.

C. Performance for download of larger file sizes

We next present similar results for download of the larger file sizes, where the effect of slow start TCP is expected to be less dominant. Fig. 4 shows the corresponding end-user throughput statistics. In line with our hypothesis, these results show superior performance for using the longer TTI as compared to the shorter TTI at high offered load. However, at the low offered load, using the short TTI size of 0.14 ms still offers improved performance.
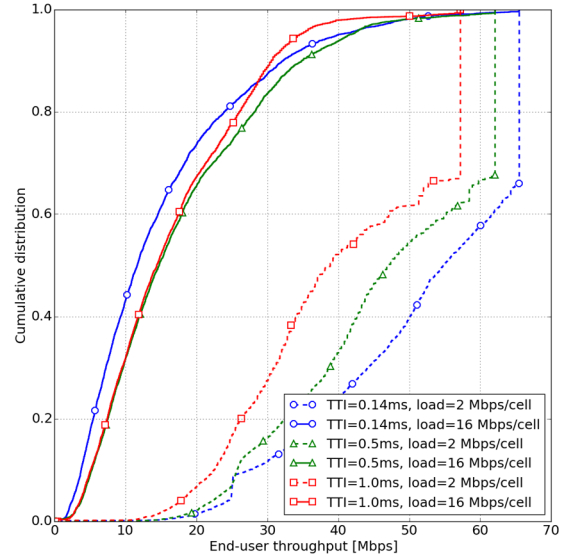

Fig. 4: Cdf of end-user throughput for download of large file sizes (500kB) under different load conditions and TTI sizes.

D. Supplementary performance statistics for TCP cases

In order to gain further insight, additional performance statistics is presented in the following. Fig. 5 shows the cdf of the CCH overhead for different file sizes, TTI lengths, and offered loads, respectively. For the short TTI size (high offered load, large file size), the CCH overhead is observed to vary from just 3% up to nearly 30% for some cases. The large dynamic range in the required CCH resources for expedition of scheduling grants emphasize the importance of having a highly flexible 5G design with efficient user-centric multiplexing of control and data. In line with our expectations, the CCH overhead is reduced with approximately a factor seven when increasing the TTI size from 0.14 ms to 1.0 ms. At low offered loads, the relative CCH overhead is reduced to take only modest values, even for the cases with short TTIs.
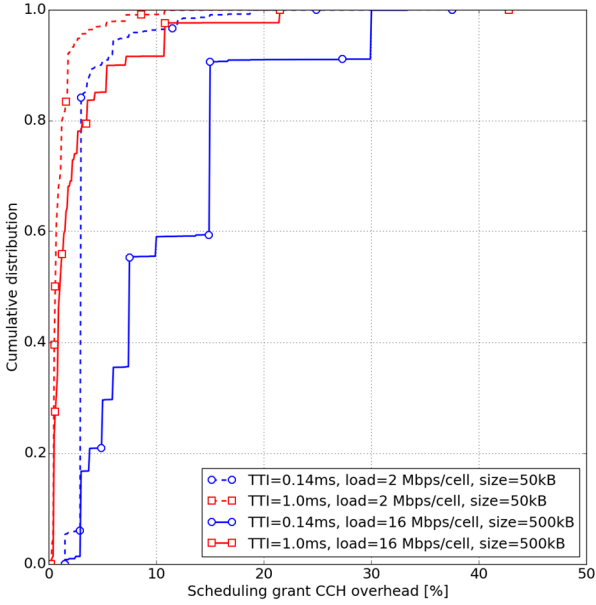

Fig. 5: Cdf of CCH overhead for different file sizes, TTI sizes, and offered load.

Fig. 6 shows the median (i.e. the 50-percentile) end-user throughput, TCP RTT, and scheduling grant CCH overhead versus the offered load for the large file size. These results shows how the optimum TTI size varies depending on the offered load. At the low offered loads, using the short TTI size of 0.14 ms results in good performance (i.e. highest end-user throughput and lowest TCP RTT), while at higher offered loads using the short TTI size results in loses. The latter is a result of higher CCH overhead, resulting in lower spectral efficiency, and therefore causing higher base station queuing delays, causing increased TCP RTT. Thus, the results in Fig. 6 clearly illustrates the benefits of being able to dynamically adjust the TTI size, rather than operating the system with a fixed TTI size.
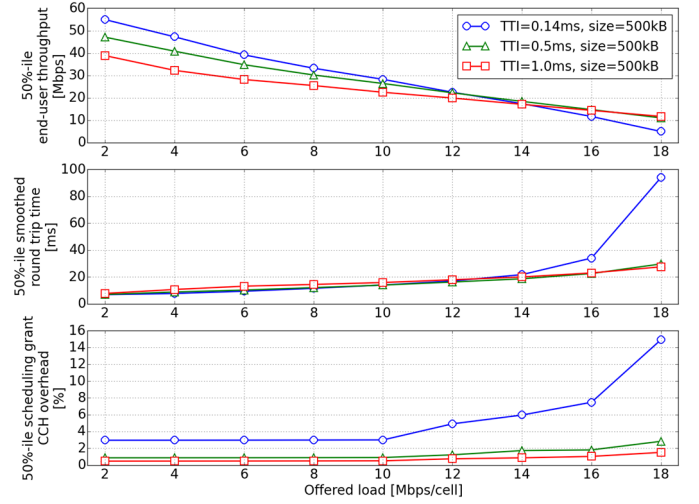

Fig. 6: Median experienced end-user throughput, TCP RTT, and CCH overhead versus offered load per cell.

Fig. 7 shows similar results, but here illustrated for the 5%-ile cell-edge outage users. That is, the users experiencing the 5%-ile lowest end-user throughput, and the 95%-ile highest CCH overhead and RTT. Similar trends as for the results in Fig. 6 are observed. However, for the cell-edge users, the CCH overhead is clearly more dominant, as well as the RTT. Hence, these results clearly show that the optimum TTI size depends not only on the file size and offered load, but also on the users experienced radio channel quality.
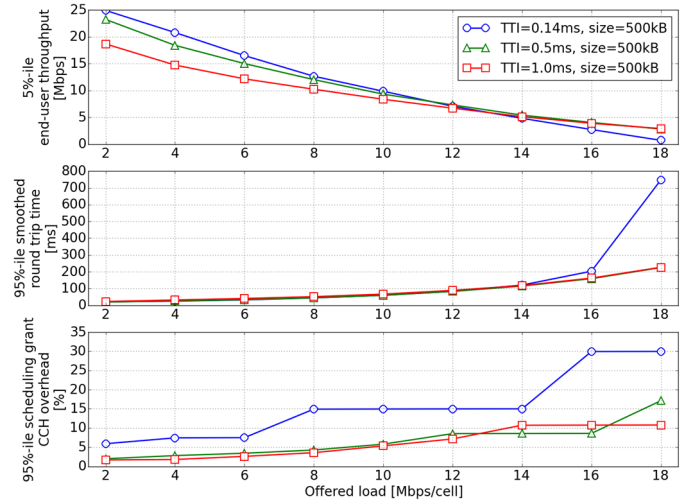

Fig. 7: Experienced end-user throughput, TCP RTT, and CCH overhead vs. offered load for the 5%-ile cell-edge outage users.

The RLC retransmission probability was also monitored in the simulations, typically showing low values. At the higher offered loads of 14 Mbps, the RLC retransmission probability was found to equal 0.1% to 0.3%, with the lowest value when scheduling with the short TTI size. Among others, the latter is contributed by the shorter HARQ RTTs for the short TTI size. At the low to medium offered loads, the RLC retransmission probability was found to some-times increase to 0.5% for the cases with small file sizes and short TTIs. The former is partly

due to the larger interference fluctuations that makes it more challenging to conduct accurate link adaptation.

V. Discussions

The presented results in the previous section clearly shows how the use of different TTI sizes result in the best performance, depending on the offered load conditions, whether slow start TCP effects are dominating, and the users' radio conditions. These findings are observed even though the conducted system level performance analysis have limited variations in the traffic characteristics. When further broadening the scope towards more diverse services (and requirements) in 5G, the benefits of supporting variable TTI size will be even more significant. Low cost mMTC devices supporting only narrow bandwidth operation will require long TTIs, while URLLC devices call for short TTIs to fulfill stricter latency requirements. Multiplexing of eMBB, mMTC, and URLLC on a wideband 5G calls for a highly flexible scheduling framework, capable of supporting large dynamic ranges in the number of simultaneously scheduled users, as well as highly diverse resource allocations. The presented user-centric concept (as inherited from [6]) meets such demands by multiplexing control and data on a per-user basis, rather than on a per-cell basis. The latter overcomes the problem of control channel scheduling grant blocking as known from LTE and LTE-Advanced ([9], [11]) that can lead to sub-optimal usage of radio resources.

Moreover, use cases with the user datagram protocol (UDP) are also of relevance [16]. UDP uses a simple protocol with a minimum of protocol overhead, and no flow control as is the case for TCP. Hence, for file download using UDP, the same benefits of using short TTIs as observed for the TCP cases (with higher layer flow control) reduces. Details on the performance of scheduling with different TTI sizes for UDP are for further study.

VI. Concluding remarks

In this paper we have presented a flexible scheduling framework that allows efficient multiplexing of users with highly diverse service requirements, experiencing different radio channel conditions. The proposed solution relies on an approach, where the resources for the user-specific control channel (CCH) scheduling grants and the corresponding data transmission are multiplexed on a per-user basis. We refer to this as a user-centric design paradigm, which deviates significantly from the cell-centric LTE design. The proposed 5G framework allows dynamic scheduling of the users with different TTI sizes. The former is an important functionality as the presented performance results show that the optimal TTI size depends on numerous cell-specific and user-specific factors, and hence must be dynamically adjusted to achieve the best performance. Using short TTI sizes offers latency benefits, but at the same time comes at the cost of higher CCH overhead. As a few examples, our results demonstrates benefits of using short TTI sizes for TCP use cases to quickly overcome the slow start phase at low to medium offered load levels, while the system is best operated with longer TTI sizes for high offered traffic loads; i.e.

optimizing for spectral efficiency to minimize queuing delays at the base station nodes.

In our future research, we will further study efficient scheduling of eMBB, mMTC, and URLLC, including derivation of more sophisticated scheduling algorithms.

References

[1] 3GPP Technical Report 38.913, "Study on Scenarios and Requirements for Next Generation Access Technologies", March 2016.

[2] IMT Vision – "Framework and overall objectives of the future development of IMT for 2020 and beyond", International Telecommunication Union (ITU), Document, Radiocommunication Study Groups, February 2015.

[3] E. Dahlman, et.al., "5G Wireless Access: Requirements and Realization", IEEE Communications Magazine - Communications Standards Supplement, December 2014.

[4] B. Soret, et.al., "Fundamental Tradeoffs among Reliability, Latency and Throughput in Cellular Networks", IEEE Proc. GLOBECOM, December 2014.

[5] D. Wu and R. Negi, "Effective capacity: A wireless link model for support of quality of service," IEEE Transactions on Wireless Communications, vol., 2, no. 4, pp. 630-643, July 2003.

[6] K.I. Pedersen, et.al., "A Flexible 5G Frame Structure Design for Frequency-Division Duplex Cases", in IEEE Communications Magazine, pp. 53-59, March 2016.

[7] D. Villa, et.al., "Performance of Downlink UTRAN LTE Under Control Channel Constraints", in IEEE Proc. Vehicular Technology Conference (VTC-spring), May 2008

[8] F. Capozzi, et. al., "UTRAN LTE Downlink System Performance under Realistic Control Channel Performance", IEEE Proc. Vehicular Technology Conference (VTC-Fall), September 2009.

[9] D. Laselva, et.al., "On the impact of realistic control channel constraints on QoS provisioning in UTRAN LTE", in IEEE Proc. VTC 2009 fall, September 2009.

[10] H. Holma and A. Toskala (Editors), "LTE-Advanced: 3GPP solution for IMT-Advanced", John Wiley & Sons, 2012.

[11] A.K. Talukdar, "Performance evaulation of the Enhanced Physical Downlink Control Channel in a LTE network", in IEEE Proc. PIMRC, pp. 987-991, September 2013.

[12] G. Berardinelli, et.al., "On the Design of a Radio Numerology for 5G Wide Area", in Proc. International conference on Wireless and Mobile Communications (ICWMC), October 2015.

[13] J. Padhye, et.al., "Modeling TCP Reno Performance: A Simple Model and Its Empirical Validation", in IEEE/ACM Trans. on Networking, Vol. 8, No. 2, pp. 133-145, April 2000.

[14] N. Shoejaedin, et.al., "TCP-aware scheduling LTE networks", in IEEE Proc. World of Wireless, Mobile and Multimedia Networks (WoWMoM), June 2014.

[15] K. Brueninghaus, et.al, "Link performance models for system level simulations of broadband radio access systems," in IEEE Proc. Personal, Indoor and Mobile Radio Communications (PIMRC), pp. 2306-2311, Sept. 2005.

[16] A.S. Tanenbaum, "Computer networks", fifth edition, Prentice Hall, 2011.