Theory and Practice of RRC State Transitions in UMTS Networks

Pekka H. J. Perälä[1], Antonio Barbuzzi[2], Gennaro Boggia[2], and Kostas Pentikousis[1]
[1]VTT Technical Research Centre of Finland, Kaitovayla 1, FI-90571 Oulu, Finland, email: firstname.lastname@vtt.fi
[2]DEE - Politecnico di Bari, 70125, Bari, Italy, e-mail: g.boggia@poliba.it; a.barbuzzi@poliba.it

***Abstract*—3GPP has been enhancing UMTS networks by issuing new standard releases recurrently. However, the RRC state transition model has remained rather unchanged through several releases, namely, Rel. 99, Rel. 05 and Rel. 06. We review RRC state transitions and study them in practice using our novel 3G Transition Triggering Tool (3G3T), focusing on the network configuration parameters that prompt these transitions. We employ 3G3T in a landmark measurement study involving four UMTS networks in three countries and validate our findings using a state-of-the art proprietary measurement tool. Our results show that 3G3T is able to discover RRC configuration parameters without operator involvement or cooperation. We observe significant differences in UMTS network configurations, which directly impact end-user data service performance. We find that, in practice, the behavior of public UMTS networks cannot be solely described based on the theoretical constructs found in the literature, and that tools such as 3G3T are necessary in order to obtain a complete picture. The results presented in this paper aim to assist simulationists in developing better models for UMTS networks. Moreover, we expect that operators can use 3G3T to configure their networks more efficiently. Finally, our methods and algorithms should be of great interest to application developers for mobile broadband networks based on 3GPP standards.**

I. Introduction

Third generation (3G) cellular networks have already gained significant customer base in many countries. The emphasis placed on gradually rolling out features by 3GPP (the evolution from GSM to EDGE to 3G/WCDMA and now HSPA) and enhancing network capacities has played a key role in such wide-spread deployment. In particular, networks based on 3GPP standards have already gone through several updates with respect to handling packet data connections. For example, Rel. 05 introduced High Speed Downlink Packet Access (HSDPA), a major enhancement to the downlink channel, with nominal peak data rates of 14.4 Mb/s. In the following release (Rel. 06), uplink packet data connection was upgraded as well with the introduction of High Speed Uplink Packet Access (HSUPA). When used together, these enhancements form a technology referred to as High Speed Packet Access (HSPA), which is already in service in 51 networks worldwide and deployment is on-going in a further 22 networks as of this writing [1]. Holma and Toskala [2, 3] provide an excellent overview of 3GPP technologies.

Despite the aforementioned improvements with respect to 3GPP access capacity, the Radio Resource Control (RRC) state transition model (illustrated in Fig. 1) has remained essentially unchanged through the currently-deployed 3GPP releases (Rel. 99, Rel. 05 and Rel. 06). RRC state transitions, the focus of this paper, are fundamental in allowing the network to balance radio resources between users and enable energy-efficient operation when no user data is pending.

UMTS terminals (typically referred to as "user equipment" or UEs) have basically two operational modes: *idle* and *connected*; see Fig. 1. When UE is powered on, it enters idle mode, i.e., it is attached to the network but is not actively engaged in data transfers. When an RRC connection is established, the terminal switches to connected mode, in which UEs can be in any of the following four RRC service states: Cell_DCH, Cell_FACH, Cell_PCH and URA_PCH. In the Rel. 99 Cell_DCH state, user data are transferred through a dedicated channel (DCH). If the network supports HSPA, High Speed Downlink Shared Channel (HS-DSCH) and Enhanced DCH (E-DCH) may also be used for downlink and uplink transmissions, respectively. In the Cell_FACH state, data are carried through common channels; typically, Random Access Channel (RACH) for uplink, and Forward Access Channel (FACH) for downlink (DL). In Cell_PCH and URA_PCH states, UEs listen to the Paging and Broadcast Channels (PCH and BCH, respectively), but uplink data transfer is not possible [2, 3].

The transitions between states are dependent on the Buffer Occupancy (BO) level at the Radio Link Control layer. Typically, a transition from Cell_DCH to Cell_FACH takes place when BO is zero and a threshold for *DCH release timer* is exceeded. The transition back to Cell_DCH is carried out, if BO level exceeds a threshold value for a set time, that is, there are data waiting to be transmitted. Moreover, if the period of inactivity in Cell_FACH is long enough (ranging 2-10 s) the UE may change its state to Cell_PCH or URA_PCH. The transition back to Cell_FACH or Cell_DCH is carried out if user activity is detected. The RRC mode may be changed from connected to idle, if the inactivity timer triggers a transition, or the Radio Network Controller is (over)loaded, in which case, the RRC connection is released [2, 4]. Further details about RRC operation can be found in [5].

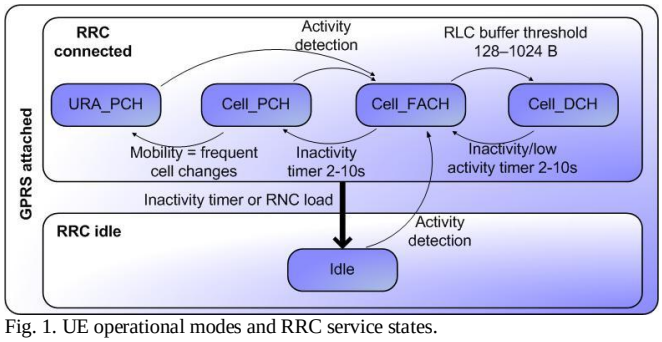

Fig. 1. UE operational modes and RRC service states.

TABLE I. POWER CONSUMPTION IN RRC STATES

RRC state	Power consumption [mA]
Cell_PCH/URA_PCH	< 5 mA
Cell_FACH	100 – 150 mA
Cell_DCH	200-300 mA

RRC state transitions have direct impact on end-user experienced performance. For example, transition from Cell_FACH to Cell_DCH mandates setting up a DCH and possibly HSPA-related channels, a process that requires a few seconds to complete. This may deteriorate user experience when using, for example, a web browser, which typically includes long periods of inactivity ("think time"). Another important aspect is battery consumption in each state. Indicative power consumption values per RRC state are given in Table I, as per [2]. The trade off is as follows: the longer UE remains in Cell_DCH, the more power is consumed; network resources are reserved for longer, but the user receives higher nominal data rates.

This paper contributes towards a better understanding of UMTS networks in practice in three distinct ways. First, we present a tool called 3G3T (3G Transition Triggering Tool), which can be used to trigger RRC state transitions, to measure one-way delay, and to determine the delay due to paging procedures and/or channel setup. A rudimentary first version of 3G3T was recently used to discover network parameters and demonstrate the feasibility of the approach [6]. In this paper, we present the extended capabilities 3G3T and validate it with the proprietary Nemo Outdoor air interface measurement tool (www.nemotechnologies.com). Second, we use 3G3T in a comparative study carried out in four live UMTS networks in three different European countries. Our study is, to the best of our knowledge, the most extensive in the published literature concentrating on studying RRC state transitions. It is also the most up-to-date as we empirically study all currently-deployed 3GPP releases (Rel. 99, Rel. 05, and Rel 06). Third, we show how the operator network settings may differ drastically from each other and that, in fact, one cannot always infer network behavior based on the available literature. 3G3T is so far the only tool that can be used to bridge this gap, especially since operators are typically not willing to share this type of information.

The rest of the paper is organized as follows. Section II presents salient related work on the area. Section III overviews the methodology and tools used in the experiments and Section IV presents the results from our experiments. Finally, Section V concludes the paper and outlines future work items.

II. RELATED WORK

The empirical and analytical performance evaluation of 3G networks is an area with several active research groups worldwide, so that several aspects of UMTS network operation and performance have been treated in the peer-reviewed literature. For instance, work has been made on the perception of the QoS experienced by end users; the so called Quality of Experience (QoE) [7] influenced by latency, delay spikes, packet loss probability and connectivity gaps. These aspects have to be taken into account by network operators in managing existing or when planning new 3G (and beyond) services [8].

User experience depends on the performance of protocols and applications employed for services. For this reason, the performance of TCP applications over 3G networks have been investigated both theoretically and experimentally; see [9]-[13] and references therein. All these studies use simulation or fixed values for describing the RRC state transition process. Early on it was pointed out that signaling delays have a significant impact on network performance [12]. Recently, RRC inactivity timers were shown to influence user-perceived throughput, which, for interactive traffic, can be up to four times lower than expected [13]. In the process, several tools have been developed aiming at evaluating the performance of live 3G networks and estimating QoE [14, 15].

One aspect highlighted by several papers is that network parameter setting for state transition (in particular, inactivity timers) can heavily affect QoE [16]. In particular, simulations were employed to analyze the impact of the inactivity timer on various system performance metrics [17]. RRC transition parameters influence UE energy consumption [18] as well. We highlight that the values of timers governing channel switching and, in general, the radio management activities are known only to operators and typically differ from one network to another (as it will become evident from our empirical measurements later in this paper). Extensive measurement studies have shown that performance depends on the 3G network considered. That is, operators have the tools to customize their network configurations [19]. Barbuzzi et al. present in [6] a first attempt to infer network parameter settings without any cooperation from network operators, but only the downlink direction and the HSDPA variant are considered. In this paper we build upon [6] and introduce 3G3T, validating it with Nemo Outdoor. We also examine transport channel transition, an issue not examined earlier.

The aforementioned references point to the importance of understanding RRC state transition properties as well as the interest in being able to trigger them based on different parameters (in particular, inactivity timers), even when there is no cooperation from the operator side. To the best of our knowledge there are no tools that analyze this aspect in a generic HSPA network, and no previous study considered in a comprehensive manner the measurements conducted in both uplink and downlink directions. This paper bridges this gap.

III. INTRODUCING THE 3G3T MEASUREMENT TOOL

We have developed an active measurement tool (3G3T) which enables the discovery of RRC state transition parameters, one-way delay measurement without expensive clock synchronization and determining the delay taken by different RRC procedures (e.g. channel establishment, paging). Clock synchronization is not needed, since only one computer is used to send and receive our measurement traffic. 3G3T, written in Python, is responsible for the initial setup, measurement execution, traffic trace capture and storage, and display of all measurement data. It permits execution of a set of experiments with various parameters on remote hosts that may

be reached using SSH connections. The measurement setup used in this paper is illustrated in Fig. 2.

In our typical usage scenario, a laptop computer with 3G3T is equipped with both a 3G data card and a Fast Ethernet card, and is connected, simultaneously, to the public Internet and to the 3G network under test. 3G3T injects data traffic of different characteristics into the network, receives the traffic in the other network interface and records the timestamp, one-way delay and Inter-Departure Time (IDT) for each packet. IDT is the delay between sending two consecutive probe packets. By increasing the IDT slowly we are able to detect the expiration of an inactivity timer as increased delay.

The characteristics and the parameters of the sending pattern are defined in the 3G3T configuration file. In order to be able to use different packets types (ICMP or UDP, for example) and to avoid packets passing through the loopback interface, packets are sent using raw sockets or the *WinPcap* library (http://www.winpcap.org/). Furthermore, 3G3T may be used to start a test from the Cell_DCH state; a burst of packets is sent at the beginning of the test until the connection is in Cell_DCH mode. All packets departing from or arriving at both interfaces are captured for later processing. In particular, the capture timestamps are used to calculate the one-way delay of each packet. For each packet, the one-way delay and the interval from the previously sent packets can be dynamically displayed in a scatterplot.

As already stated, if we exclude cell-reselections (due to mobility, for example) and contention (due to arrival of higher priority traffic, for instance), the UE state depends only on the traffic on the user plane. Thus, the characteristic of the traffic generated and received by the mobile device is directly coupled with RRC state transitions. This is ensured with 3G3T, which recognizes background traffic and discards results affected by traffic that is not generated by it.

The work presented herein makes use of a novel tool, called 3G3T, that was written based on the earlier experience presented in [6]. One of the many enhancements to the previous work was a Windows version of the tool in addition to the Linux scripts. This permitted the simultaneous use of 3G3T with Nemo Outdoor. Another was Network Address Translation (NAT) and firewall traversal. This was very important for this work, because many UMTS operators assign a private IP address to the user or limit the inbound traffic with a firewall, which effectively hinders sending packets to DL direction. Data display, trace formats and general usability were improved as well.

IV. Measurement Results

In this section we first provide a brief description of the measurements setup, then we proceed to investigate the RRC state transition behavior, examining for the networks under test at first the transition from Cell_DCH to Cell_FACH and then the transition from Cell_FACH to onwards state.

A. Measurement Setup

Our measurement campaign was conducted on four different networks, two of them, Net#1 and Net#2, are operated by two Finnish operators in Oulu. Net#3 belongs to a major Austrian operator in Vienna. Net#4 is located in Bari, Italy.

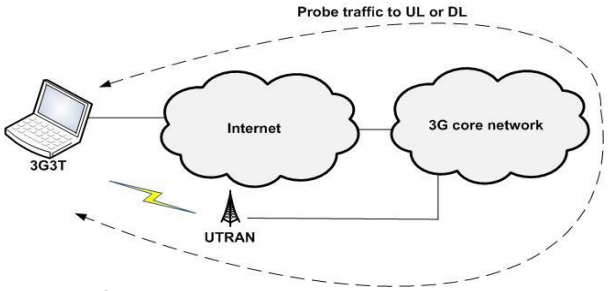

Fig. 2. A typical 3G3T measurement setup.

Traceroute was used to determine the length of paths between the two interfaces; for Net#1 there were 7 hops, for Net#2 14, for Net#3 12 and for Net#4 17 hops. Several measurements were carried out with 3G3T in these networks in order to discover the parameters that determine RRC transition behavior. The empirical results presented in this section were repeated several times in different times and days from stationary locations in order to rule out the effects of other users and mobility. For the tests, we have used UDP probe packets with 200 byte payloads.

Four different 3G data cards were used for our tests: an Option FUSION+ HSDPA, an Option GT MAX 3.6, Option GT MAX HSUPA and Option GT 3G Quad. The first two cards are Rel. 05-compatible (HSDPA), the third card is Rel. 06-compatible (HSPA-capable). The fourth card is only Rel. 99-compatible (WCDMA). Our measurement setup throughout this study was essentially the one illustrated in Fig. 2.

B. Cell_DCH to Cell_FACH Transitions

In the first measurement with 3G3T we used the HSPA-compatible data card to trigger a transition from Cell_DCH to Cell_FACH in Net#1. The results for the downlink are shown in Fig. 3 as a function of the packet IDT. The one-way delay, *d*, (shown on the vertical axis) is the period between the time that a probe packet is transmitted and the time that it is received on the receiving network interface. As IDT increases, the transition to the Cell_FACH state is carried out when the period of inactivity exceeds the release timer threshold as configured by the operator. This was observed as a change in delay: in Cell_DCH the average DL delay ($d_{DCH\ DL}$),is 44 ms, but in Cell_FACH the delay (d_{FACH}) increases to 140 ms. This behavior was verified with Nemo Outdoor: the data card was indeed in Cell_DCH state with both HS-DSCH and E-DCH active before the transition. In the transition, the packet technology changed from HSPA to UMTS FDD and the transport channels carrying user data changed to FACH for downlink (DL) and RACH for uplink (UL) respectively. In the case of Net#1, the threshold of the release timer ($T_{DCH,\ off}$) was approximately 2.8 s. Few packets experience a delay between 1.4 and 1.7 s (mean 1.6 s). These packets experienced a longer delay because they arrived during the state transition; hence they were buffered and retransmitted as soon as the state transition procedure was successful. The average delay of these packets ($D_{FACH,\ ON}$) can be used as an heuristic estimate of the procedure completion time.

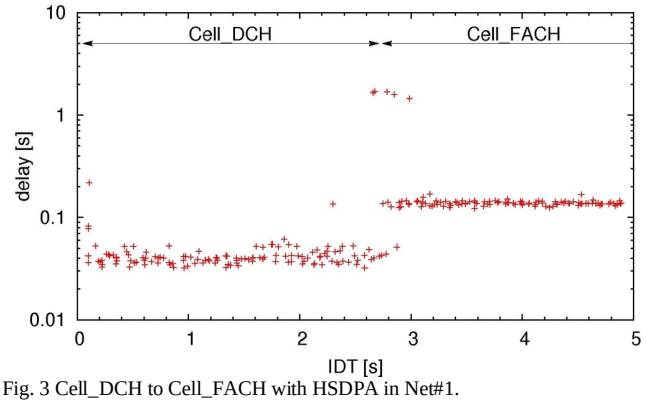

Fig. 3 Cell_DCH to Cell_FACH with HSDPA in Net#1.

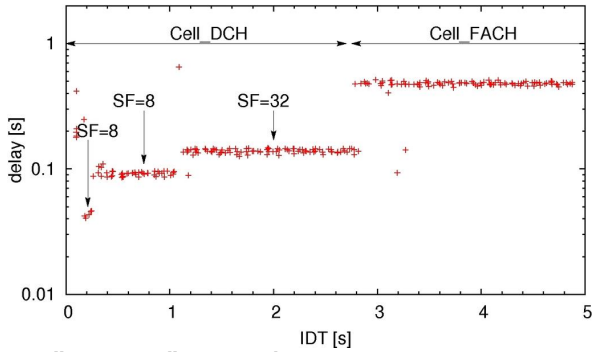

Fig. 4. Cell_DCH to Cell_FACH with HSUPA in Net#1.

A similar experiment was conducted for the uplink direction as well. Fig. 4 illustrates that in this case the - measurement results showed much more dynamic behavior when compared with the DL measurements shown in Fig. 3. After further experiments with Nemo Outdoor it became clear that the first steps related to the network assigning a channel with lower capacity to the user due to inactivity, a result of the increasing IDT. The capacity of the channel is directly coupled with a Spreading Factor (SF), which is the ratio of the chips to the baseband information rate. In the beginning, the SF for uplink is 4, but when the IDT becomes larger, SF increases leading to larger one-way delays. Similar dynamic capacity allocation behavior was also observed in the uplink of Net#3 that assigned two different SFs. On the other hand, only one SF was observed when conducting similar measurements for Net#2 and Net#4.

In the case of Net#1 (Fig. 4), the measured $T_{DCH,\ off}$ for uplink was very close to the one measured for downlink. Thus, we inferred that the release timer threshold was the same for both uplink and downlink. Two packets in Fig. 4 have delay values similar to the ones measured in Cell_DCH state although the IDT is already above $T_{DCH,\ off}$. These packets are considered outliers and the lower delay values are likely caused by traffic other than our probe packets triggering a Cell_FACH to Cell_DCH transition. Figures 3 and 4 also point out the different capabilities of uplink and downlink transport channels in Cell_FACH state: d_{FACH} is 140 ms, while d_{RACH} is 470 ms, or more than three times larger.

After concluding the experiments with the HSPA-capable card, we used the Option GT 3G Quad card which is only capable of accessing WCDMA in Net#1 in order to find out if

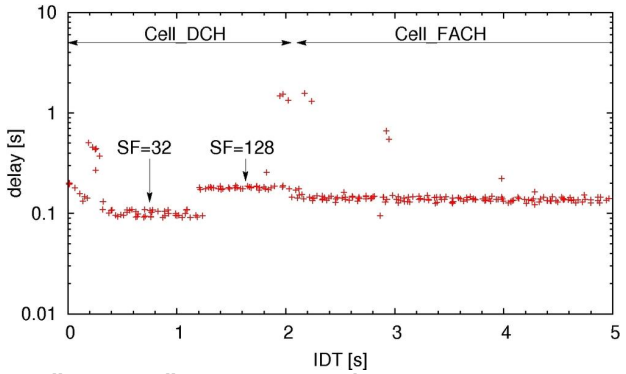

Fig. 5. Cell_DCH to Cell_FACH transition with WCDMA DL in Net#1.

the observed behavior was 3GPP release dependant. First, the WCDMA downlink was measured and is illustrated in Fig. 5. Several differences are evident. $T_{DCH,\ off}$ is approximately 2.1 s, a value approximately 0.8 s lower than the ones recorded in the HSPA measurements. Furthermore, similar dynamic channel allocation as seen for HSUPA in Fig. 4 was observed in downlink as well. This suggests that dynamic capacity allocation is implemented only for scheduling of dedicated channels. On the contrary, HSDPA user plane transport channel is shared among all cell users, which requires different scheduling policies from those used for dedicated channels. One particularly interesting finding is that a dedicated channel with SF=128 is actually slower than the service provided by FACH. Similar measurements were carried out in the uplink direction and the results were very similar to the ones presented in Fig. 4, but are omitted due to space constraints. The only discernible difference was that the release timer threshold was the same as for downlink, namely 2.1 seconds. This suggests that the release timer parameter configuration is packet technology specific. The lower value for the release timer in WCDMA is likely due to the need to release DCH faster than the shared HS-DSCH channel.

We note the special start-up behavior in Figures 3-5 and 8. We attribute it to buffering occurring at the beginning of each measurement. This is akin to the first connection goodput phenomenon first reported in [12]. When the channels are allocated, the buffers are emptied and several packets are received in order, which results in high delays for the packets that were sent first.

Up to now, we have focused on the transition from Cell_DCH to Cell_FACH, which is triggered by timers. On the other hand, the transition from Cell_FACH to Cell_DCH is triggered based on the level of buffer occupancy. The threshold was measured by using a traffic generator to send 56-byte UDP payloads through Net#1. The transition back to Cell_DCH was reached with a transmission rate of 7 packet/s. In terms of layer 3 throughput, the threshold was ~4 kb/s.

C. Cell_FACH to Power-conserving State Transitions

The second set of measurements was conducted to trigger the transitions from Cell_FACH state to states with less power consumption (Cell_PCH, URA_PCH, and idle mode). The IDT of consecutive packets was increased in steps of 0.5 s. In

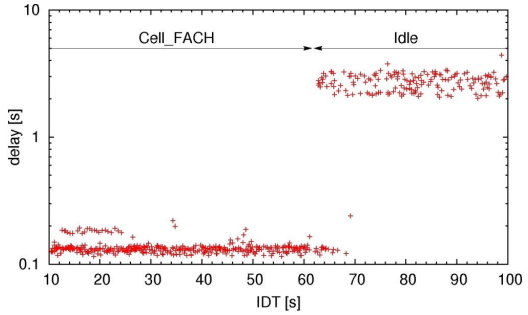

Fig. 6 Cell_FACH to idle mode transition for HSDPA in Net#1.

Fig. 6 the transition from Cell_FACH to idle state is illustrated. It was discovered that Net#1 did not carry out any transitions to Cell_PCH or URA_PCH states, but after approximately 63 s of inactivity the state was changed to idle. The value of this release timer is denoted hereafter as T_{Idle}. The transition from connected to idle mode requires the release of all RRC connections. Thus, whenever user data was sent in idle state, the system triggered a state transition to Cell_DCH, which caused rather long delays due to the three-way handshake signaling and establishment of related channels. An estimate for the delay taken by the transition from idle to Cell_DCH (D_{DCH}) is obtained by subtracting the average delay in Cell_DCH state from the idle mode delay, and is on average around 2.7 s. The few packets that are transferred faster after the transition(s) to idle mode start to occur, are sent while in the Cell_FACH state. This is because D_{DCH} is 2-4 s and $T_{DCH,off}$ is approximately 2.9 s. This way the idle release timer of every other packet is started after waiting for these delays first. Thus, the transition from Cell_FACH to idle does not take place with every second packet when IDT is between approximately 63–69 s. Similar measurements were carried out with the GT 3G Quad card and to uplink as well and the results were similar to the ones depicted in Fig. 6 for Net#1, but are omitted in due to space limitations.

Measurements similar to the one in Fig. 6 were also carried out in Nets #3 and #4 with an HSDPA-enabled card. The observed behavior was well in line with our previous experiences in Net#1, although the values of discovered parameters were different. This again highlights the high level of customization on settings done by operators. In Fig. 7 an uplink measurement in Net#3 with IDTs ranging from 20 s to 50 s is depicted. T_{Idle} is nearly 29 s in this case. The line connecting the markers stresses the fact that every second packet is sent in Cell_FACH due to the state transition delay from idle mode to Cell_DCH.

Completely different behavior from the other three networks was observed with Net#2, as shown in Fig. 8. The original goal was to trigger similar state transitions from Cell_DCH to Cell_FACH as was done with Net#1 in Fig. 3. It turned out that the system was in Cell_DCH state until the IDT between probe packets reached 3 s, which was also $T_{DCH,off}$ for Net#2. After that, the system starts to trigger transitions back and forth between Cell_FACH and Cell_DCH states (dynamic region in Fig. 8). Occasionally, a transition to Cell_PCH state occurred as well. Furthermore, each probe packet triggers transition back to Cell_DCH state, which suggests very low capacity in the FACH/RACH channels. After the dynamic region (~3-11 s) the transition pattern for each probe packet is Cell_PCH → Cell_FACH → Cell_DCH → Cell_FACH → Cell_PCH, and so on. The threshold that triggers state transition from Cell_FACH to Cell_PCH is well below 10 s, thus packet delay may cause transition to Cell_PCH, although the IDT value should not yet trigger it. This is observed in the variable delay values of the dynamic region. The test was also done towards the uplink, and the behavior was similar, despite the fact that the uplink channel was Rel. 05 DCH.

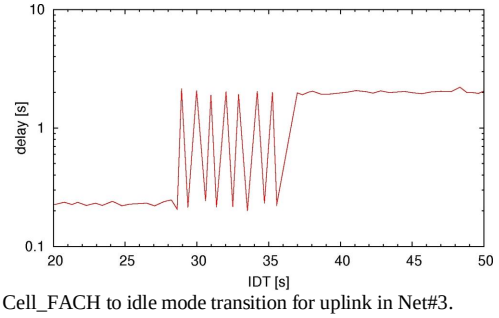

Fig. 7. Cell_FACH to idle mode transition for uplink in Net#3.

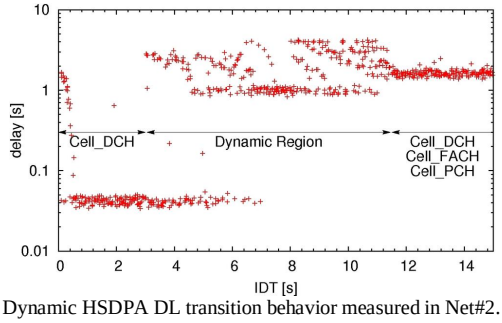

Fig. 8. Dynamic HSDPA DL transition behavior measured in Net#2.

D. Comparison among networks

The results for Net#3 and Net#4 were very similar to the ones obtained from Net#1. Unfortunately, not all results can be presented here due to space limitations. Hence, we present the obtained averages (with ± the bounds of the 90% confidence interval) of timers and DCH setup delays in Table II. For the timer values a range between minimum and maximum is given. The behavior of Net#3 and Net#4 could not be verified with Nemo Outdoor because of license issues, but due to similarity of the results we are confident that the state transitions work as they do with Net#1.

The averages of measured one-way delay are given in Table III. Net#3 gave lowest delay for HSDPA ($d_{DCH,\,DL}$). On the other hand Net#1, which was the only one capable of HSUPA, has lowest $d_{DCH,\,UL}$. It does not differ significantly from other networks, because the system dynamically assigns lower SF to the user as IDT increases, as was seen in Fig. 4. However, this does not imply that HSUPA cannot perform better than traditional WCDMA [20]. Table III also presents the measured delays for RACH and FACH channels (d_{RACH}/ d_{FACH}) as well as the delay for sending packets when UE is in the idle mode (d_{IDLE}).

On the other hand, Net#2 showed a behavior very different from the other three networks, as discussed earlier. This is

TABLE II – TIMERS AND DCH SETUP DELAYS

	$T_{DCH,\ off}$ [s]	T_{Idle} [s]	$D_{DCH,\ DL}$ [s]	$D_{DCH,\ UL}$ [s]
Net#1	2.66 - 2.99	62.9-63.6	2.67 ± 0.05	1.87 ± 0.01
Net#2	3.02 - 3.07	n/a	n/a	n/a
Net#3	2.74 - 2.91	29.1- 30.2	2.63 ± 0.033	1.88±0.0047
Net#4	1.20 - 1.80	20.4 – 21.3	2.45 ± 0.019	1.73 ± 0.058

TABLE III – MEASURED ONE-WAY DELAYS

	$d_{DCH,\ DL}$ [ms]	$d_{DCH,\ UL}$ [ms]	d_{FACH} [ms]	d_{RACH} [ms]	$d_{IDLE,\ DL}$ [s]	$d_{IDLE,\ UL}$ [s]
Net#1	43.6 ± 2.8	129 ± 8.2	138 ± 1.10	440 ± 2.7	2.71 ± 0.05	2.0 ± 0.01
Net#2	48.9 ± 8.3	109±13	n/a	n/a	n/a	n/a
Net#3	33.7 ± 0.23	147 ± 2.5	136 ± 0.47	445± 0.59	2.67± 0.033	2.03 ± 0.0047
Net#4	46.0 ± 1.26	93.9 ± 1.3	152± 0.58	479± 1.6	2.50 ± 0.019	1.82 ± 0.058

due to the completely different configuration of RRC state transition parameters. In fact, Net#2 was the only one that carried out transitions to Cell_PCH state, but not to idle mode if the packet connection was not released by the user. The network triggered state transitions aggressively with IDTs between 0-15 s. This calls for a trade off: longer delays are introduced aiming to achieve energy savings. In practice, however, a Net#2 user often has to wait several seconds before any data is transmitted, even if actual inactivity is merely a few seconds, since a transition to Cell_DCH has to be performed always. A user browsing the web is therefore bound to experience performance degradation due to these transitions. The UE and the network also incur more signaling overhead in Net#2 due to the frequent transitions.

V. CONCLUSION

This paper presented and validated a novel active measurement tool called 3G3T. We successfully used 3G3T to determine RRC state transition parameters, measure one-way delay without expensive clock synchronization equipment and determine channel setup and paging delays of UMTS networks without any operator involvement or cooperation. We validated correct operation using Nemo Outdoor and used 3G3T in four different public UMTS networks in three countries to empirically investigate through a large scale measurement study RRC state transitions across three 3GPP releases. The results presented will be of great value to other researchers in this field, not least, simulationists and application developers and even network operators, to some extent.

We found that three of the networks studied had similar transition behavior, although the examined parameters had different values. These networks did not carry out any transitions to Cell_PCH or URA_PCH states, but after a period of inactivity the mode was changed directly to idle. This was rather surprising, since standard literature points to a transition to Cell_PCH or URA_PCH instead. On the other hand, Net#2 did carry out RRC state transitions as expected, but transition threshold configuration actually led to highly transitory operation. In short, we observe that, in practice, advanced cellular networks operate differently from what a thorough literature survey would suggest. Although all networks follow 3GPP specifications, reliably predicting their exact behavior beforehand is difficult, if at all possible, and tools such as 3G3T can be valuable, especially as operators typically do not share network configuration information. 3G3T can be used to extract information about the network parameters, as was shown in this work. However, one should bear in mind that the parameters may vary across operators or across different cells of an operator.

We are working to further improve 3G3T to discover network parameters and plan to co-operate with network operators in performance and security issues posed by RRC state transition triggering.

ACKNOWLEDGMENT

This work was supported by Celtic Easy Wireless II project and the Future Internet program of TIVIT (Finnish Strategic Centre for Science, Technology and Innovation in the field of ICT), which are in part funded by Finnish Funding Agency for Technology and Innovation (TEKES), and by PS-121 project (Italy, Apulia Region), as well as COST action IC0703.

REFERENCES

[1] GSM World, HSPA networks, http://hspa.gsmworld.com/networks.
[2] H. Holma and A. Toskala, *HSDPA/HSUPA for UMTS*, John Wiley & Sons Ltd., 2006, pp. 29-30.
[3] H. Holma and A. Toskala, *WCDMA for UMTS* , John Wiley & Sons Ltd., 2000, pp. 136-137.
[4] G. Gómez and R. Sánchez, *End-to-End Quality of Service over Cellular Networks*, John Wiley & Sons Ltd., 2005, pp. 150-152.
[5] 3GPP, "Radio Resource Control (RRC) protocol specification" *3GPP specification of release 6 – TS 25.331*, v6.2.0, Dec. 2008.
[6] A. Barbuzzi, F. Ricciato, and G. Boggia, "Discovering parameter setting in 3G networks via active measurements," *IEEE Comm.. Letters*, vol. 12, no. 10, pp. 730-732, Oct. 2008.
[7] C. Gomez, M. Catalan, D. Viamonte, J. Paradells, and A. Calveras, "Internet traffic analysis and optimization over a precommercial live UMTS network," *Proc. IEEE VTC*, Stockholm, Sweden, May 2005.
[8] Y.-M. Li and Y.-S. Yeh, "Service quality's impact on mobile satisfaction and intention to use 3G service," *Proc. HICSS*, Hawaii, USA, Jan. 2009.
[9] M. Catalan, C. Gomez, D. Viamonte, J. Paradells, A. Calveras, and F. Barcelo, "TCP/IP analysis and optimization over a precommercial live UMTS network," *Proc. IEEE WCNC*, New Orleans, USA, Mar. 2005.
[10] K. Pentikousis, "TCP in wired-cum-wireless environments," *IEEE Commun. Surveys & Tutorials*, vol. 3, no. 4, 4th Quarter 2000, pp. 2-14.
[11] W. L. Tan and O. Yue, "Measurement-based Performance Model of IP Traffic over 3G Networks," *Proc. IEEE TENCON*, Melbourne, Australia, Nov. 2005.
[12] K. Pentikousis, M. Palola, M. Jurvansuu, and P. Perala, "Active goodput measurements from a public 3G/UMTS network," *IEEE Commun. Letters*, vol. 9, no. 9, Sep. 2005, pp. 802-804.
[13] M.Catalan, C Gomez, A. Calveral and J. Paradells, "Web Optimization in Real UMTS Networks with Channel Switching Mechanisms", *Proc. IEEE VTC*, Marina Bay, Singapore, May 2008.
[14] A. Diaz, P. Merino, A. Gil, J. Munoz, "x-AppMonitor μAgent: a tool for QoS measurements in cellular networks," *Proc. ISWCS*, Valencia, Spain, Sep. 2006.
[15] J. Prokkola, M. Hanski, M. Jurvansuu, M. Immonen, "Measuring WCDMA and HSDPA Delay Characteristics with QoSMeT," *Proc. IEEE ICC*, Glasgow, Scotland, June 2007.
[16] L. de Bruynseels, "Tuning Inactivity Timer Settings in UMTS," Comm-Square Tech. Rep., 2005.
[17] M. Chuah, W. Luo, and X. Zhang, "Impacts of inactivity timer values on UMTS system capacity," *Proc. IEEE WCNC*, Orlando, FL, USA, Mar. 2002.
[18] C.-C. Lee, J.-H. Yeh, and J.-C. Chen , "Impact of inactivity timer on energy consumption in WCDMA and cdma2000," *Proc. of WTS*, Pomona, CA, USA, May 2004.
[19] W. L. Tan, F. Lam, and W. C. Lau, "An Empirical Study on the Capacity and Performance of 3G Networks ," *IEEE Trans. on Mobile Computing*, vol. 7, no. 6, Jun. 2008, pp. 737-750.
[20] J. Prokkola, P. Perälä, M. Hanski and E. Piri, "3G/HSPA Performance in Live Networks from the End User Perspective", *Proc. IEEE ICC*, Dresden, Germany, June 2009.