Copyright © IEEE.
Citation for the published paper:

Title:

Author:

Journal:

Year:

Vol:

Issue:

Pagination:

URL/DOI to the paper:

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of BTH's products or services Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by sending a blank email message to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Estimating Performance of Mobile Services from Comparative Output-Input Analysis of End-to-End Throughput

Markus Fiedler, *Member, IEEE*, Katarzyna Wac, *Member, IEEE*, Richard Bults, and Patrik Arlos, *Member, IEEE*

Abstract—Mobile devices with ever-increasing functionality and the ubiquitous availability of wireless communication networks are driving forces behind innovative mobile applications enriching our daily life. One of the performance measures for a successful application deployment is the ability to support application-data flows by heterogeneous networks within certain delay boundaries. However, the quantitative impact of this measure is unknown and practically infeasible to determine at real-time due to the mobile device resource constraints. We research practical methods for measurement-based performance evaluation of heterogeneous data communication networks that support mobile application-data flows. We apply the lightweight Comparative Output-Input Analysis (COIA) method estimating an additional delay based on an observation interval of interest (*e.g.*, one second) induced on the flow. An additional delay is the amount of delay that exceeds non-avoidable, minimal end-to-end delay caused by the networks propagation, serialization and transmission. We propose five COIA methods to estimate additional delay and we validate their accuracy with measurements obtained from the existing healthcare and multimedia streaming applications. Despite their simplicity, our methods prove to be accurate in relation to an observation interval of interest, and robust under a variety of network conditions. The methods offer novel insights into application-data delays with regards to the performance of heterogeneous data communication networks.

Index Terms—mobile application, additional delay, heterogeneous networks, application-level, throughput

1 INTRODUCTION

EMERGING wireless network technologies and miniature personalized networked devices enable provision of new mobile applications that enrich daily activities of their users. These applications aspire to deliver mobile services to users ‘anywhere-anytime-anyhow’ [1] while fulfilling their Quality of Service (QoS) and Quality of Experience (QoE) requirements [2] [3], *e.g.*, low application response times. The success of the mobile service delivery depends heavily on the performance provided by underlying network infrastructures [1], [2], [4]. While users are on the move, these services operate in heterogeneous networking environments and knowledge of the overall performance at a particular user location and time is required by the service to optimize its user’s experience. In particular, for interactive mobile applications that exchange data between spatially dispersed mobile and fixed nodes, the end-to-end data delays and their unexpected increase are critical performance measures [4]. However, these are difficult to measure at application run-time due to limited mobile

device resources and clock synchronization issues [5].

To this end, we research practical methods for measurement-based performance evaluation of heterogeneous data communication networks that support mobile application-data flows. Particularly, we propose methods for an application’s run-time estimation of additional delay, *i.e.*, a stochastic delay exceeding the non-avoidable end-to-end delay, which considers some allowable data communication network’s propagation, serialisation and transmission and delays at network nodes [4]–[6]. Additional delay consists of those delay elements that increase the end-to-end delay of a message. For example, for a queued and pending transmission message, it is a delay occurring due to the fact that the scheduler has not rescheduled the thread handling its actual transmission.

We propose five different methods for additional delay estimation, that are based on *Comparative Output-Input Analysis* (COIA), which builds upon the comparative analysis of application level throughput at the sender and receiver nodes (*e.g.*, a mobile device and an application server). We assume a ‘black box’ view of the application on the underlying heterogeneous network infrastructure, where the application can just observe network behavior without enforcing it, or exploiting network-level measurements and/or performance feedback. We also assume an absence of strong clock synchronization between the sending and receiving nodes. We use the stochastic fluid flow model [7], [8] to analyze application-data traffic at small observation time scales,

- Markus Fiedler and Patrik Arlos are with Blekinge Institute of Technology, 37179 Karlskrona, Sweden, e-mail: Forename.Name@bth.se.
- Katarzyna Wac is with Institute of Services Science, University of Geneva, Switzerland, e-mail: Katarzyna.Wac@unige.ch.
- Richard Bults is with MobiHealth BV, The Netherlands, e-mail: Richard.Bults@mobihealth.com.

e.g. a second, where the choice of a time scale is determined by the mobile service data delivery requirements.

We evaluate the accuracy of the proposed methods under a variety of network operator conditions, based on data traces from an existing mobile application, a health telemonitoring application provided by the MobiHealthTM system [9], [10] and a mobile multimedia streaming application. Especially mobile healthcare applications pose strict application-level QoS requirements, since a patient in an emergency situation may require an immediate system response [11].

This paper is structured as follows. Section 2 provides concepts for modelling of additional delay, while Section 3 introduces the additional delay model and its estimation methods. Sections 4 and 5 provide accuracy evaluation results for these methods with data traces from the existing mobile applications. Section 6 provides implications for methods implementation, while Section 7 concludes upon our research and draws on future work areas.

2 MODELLING OF THE ADDITIONAL DELAY

2.1 The Concept of Additional Delay

Additional delay T^{add} is the amount of delay that exceeds the minimal end-to-end delay T^{min} , while T^{min} is dominated by inherent, allowable data communication network's propagation, serialisation and transmission delays at network nodes [4]–[6], [12]. T^{add} obeys a dynamic random process dominated by transmission conditions as well as competition from other traffic and processes [12]. Assuming that data is sent by the sender application entity at time T^{in} and received by the receiver's entity at time T^{out} , T^{add} is obtained as

$$T^{\text{add}} = T^{\text{out}} - T^{\text{in}} - T^{\text{min}}. \quad (1)$$

Ideally, the additional delay vanishes. However, as it grows, the more pronounced its negative effect on the aggregated end-to-end delay becomes, in particular if it varies significantly. Most interactive applications suffer from T^{add} exceeding certain thresholds; *i.e.*, multimedia data can be considered lost if they miss their delivery deadlines, and user-perceived waiting times might grow beyond patience [13]. Ideally, the additional delay of each and every application-level message should be monitored in order to get a clear view of the disturbances that its delivery process is exposed to. This is not feasible. The mobile node might not be able to trace and timestamp each packet in real-time due to scarce processing resources [4]. There are also limitations regarding the exactness of the timestamps, which might be too limited in resolution, incorrect due to processing times [14], and difficult to compare due to clock synchronization issues [15].

A certain amount of T^{add} variation, also denoted as jitter, is expected for packet-based data delivery. Each link and node leaves its 'footprint' on the inter-packet-timing within a packet stream, which means that the

timely behaviour of traffic at the receiver does not match the one at the sender. For instance, round-trip delays over an empty UMTS network (*i.e.*, when no other traffic exists on the network) usually exhibits jitter of ± 10 ms [16]. These variations are not necessarily problematic from the viewpoint of the application. However, values of T^{add} in a range of one to several hundred milliseconds are perceptible [13], especially for interactive applications such as gaming [17]. These T^{add} values stem, *e.g.*, from congestion within access or core networks, or from temporarily bad radio network conditions, implying the need to resend data that was lost or corrupted. From the end-user point of view, the latter looks like a sudden loss of capacity, followed by a burst of data arriving at the receiver with a much smaller spacing in time than they were sent [18], [19]. It is thus important to capture and handle additional delays exceeding the expected variation thresholds. In the next subsection, we present a model that is able to provide this information.

2.2 Application Flow Model

A model that has shown to be capable of discerning between less critical delays on a packet level and more critical delays on a burst level is the *fluid flow model* [7], [8]. So far, it has been used for analysis of multiplexing in fast packet-switched networks and the consequences of temporary mismatches between capacity demand and availability, leading to considerable queuing on time scales beyond the packet scale. This model is based on the analysis of instantaneous workload instead of modelling the packet occurrence and length processes. In other words, the packets' intensity flow is quantified by the *data rate* $R(t)$, also denoted as *throughput* and measured in bits per second (bps), Bytes per second (Bps) or packets per second (pps). In general, $R(t)$ is a function of time and can be defined either on a continuous time scale as a derivative of the *cumulative workload* $W(t)$, passing a point of reference, by time:

$$R(t) = \frac{d}{dt} W(t), \quad (2)$$

or on a discrete time scale as the workload observed at a point of reference during averaging interval i of duration ΔT , divided by ΔT :

$$R_i = \frac{W(i\Delta T + T_0) - W((i-1)\Delta T + T_0)}{\Delta T}. \quad (3)$$

The time counting is started upon the occurrence of the first packet being observed at the *inlet* (sender) or the *outlet* (receiver) of the network [19]. We denote the corresponding start times as T_0^{in} and T_0^{out} , respectively. A packet occurring at an arbitrary time t^{in} at the inlet and at time t^{out} at the outlet contributes its workload to the overall workload in the intervals with numbers

$$i^{\text{in/out}} = \left\lceil \frac{t^{\text{in/out}} - T_0^{\text{in/out}}}{\Delta T} \right\rceil. \quad (4)$$

Obviously, the same packet can appear in different intervals at inlet and outlet, resulting in a non-vanishing additional delay at the time scale ΔT as described in the next subsection. The throughput time series at the network inlet and outlet are denoted by $\{R_i^{\text{in}}\}_{i=1}^n$ and $\{R_i^{\text{out}}\}_{i=1}^n$, respectively. They can be obtained from bit-, Byte- or packet counting at each ΔT , followed by calculating (3). Compared to the effort related to tracing each and every timestamp of a packet observed at a point of reference, our approach is *lightweight*.

2.3 Comparative Output-Input Analysis and Equivalent Bottleneck

As outlined in Section 2.2, the comparison of the packet delivery processes at the *outlet* of the network with that at the *inlet*, together with subsequent analysis, abbreviated as COIA, allows for the quantification of the additional delay T^{add} . COIA builds upon the comparative analysis of the application-level throughput as observed at the network inlet and outlet, along a methodology presented in [16] and in absence of perfect clock synchronization of the sender and receiver. Initially, we assume lossless data delivery in order to highlight the properties and in particular the precision of COIA, and we will relax this assumption later.

We have already applied COIA to the analysis of throughput time series and related summary statistics. Amongst others, we have shown a classification of bottleneck behaviour based on the comparison of throughput averages, standard deviations and histograms between inlet and outlet [18], [19]. We have illustrated that if the standard deviation of a throughput time series increased between inlet and outlet, the network in-between acts as a shared bottleneck introducing additional delay. Thus, throughput time series and related summary statistics are lightweight descriptions of the data flow at the burst level. As they capture essential properties of the data flow, they can be used as *Reduced Reference Metrics* (RRM). In particular, RRM can be exchanged between inlet and outlet (*i.e.*, sender and receiver) in order to allow for runtime classification of bottleneck characteristics. Here, we investigate to which extent COIA based on throughput time series allows for an estimation of the additional delay T^{add} . To this end, we consider the end-to-end path as one *equivalent bottleneck*, whose content at the end of interval i is described by [18]:

$$X_i = X_{i-1} + D_i \Delta T, \quad X_0 = 0. \quad (5)$$

X_i describes the amount of data that is still in transit at the end of interval i . As synchronization happens on the first packet, there is no content at the beginning of a session, *i.e.*, $X_0 = 0$. D_i is a throughput difference between inlet and outlet, called *drift* and defined as

$$D_i = R_i^{\text{in}} - R_i^{\text{out}}. \quad (6)$$

The drift is a central parameter in fluid flow modelling [8], describing the rate at which the content increases or decreases, *e.g.*, $D_i > 0$ means $X_i > X_{i-1}$ and

$D_i < 0$ means $X_i < X_{i-1}$ if $X_{i-1} > 0$, while vanishing drift $D_i = 0$ implies a constant content $X_i = X_{i-1}$. If the time series are equal, *i.e.*, $\{R_i^{\text{in}}\}_{i=1}^n = \{R_i^{\text{out}}\}_{i=1}^n$, there is neither drift ($D_i = 0$) nor content ($X_i = 0$), the equivalent bottleneck remains empty and the network is considered to be *transparent* at the time scale ΔT . In the special case $\{R_i^{\text{in}}\}_{i=1}^n = \text{const}$, the variations in $\{R_i^{\text{out}}\}_{i=1}^n$ reflect the variations in D_i and thus of the buffer content of the equivalent bottleneck as such.

The fluid flow model assumes a fluid particle flow of constant intensity during the averaging interval. As long as a packet that was sent in interval i is received inside the same interval, its additional delay is invisible. However, as soon as a packet belonging to interval i traverses an interval boundary and is received in consecutive interval $j > i$, its additional delay becomes visible. Thus, the proposed methods rely on intervals bounds in order to quantify how much data from interval i is delayed to consecutive interval(s). Let us illustrate this by assuming $D_i > 0$, $D_{i+1} = -D_i$ and $X_{i-1} = 0$: During interval i , less traffic leaves the outlet than what was delivered at the inlet. This particular amount $X_i = D_i \Delta T$ is still in transit on the end of interval i and is thus delayed. In the next interval $i + 1$, the queue gets empty again ($X_{i+1} = 0$). Intuitively, we estimate an additional delay in the order of ΔT across the equivalent bottleneck.

Consider now one packet sent at a uniformly distributed time during interval i and received at a uniformly distributed time during interval $i + 1$. We arrive at a triangular delay distribution over $[(i-1)\Delta T, (i+1)\Delta T]$ with its median and average at ΔT . The latter is consistent with the estimation discussed in the preceding paragraph. However, the uncertainty of our estimation amounts to $\pm \Delta T$, which implies that the choice of the time interval obviously has an impact on the error margins. This will be illustrated in Sections 4 and 5.

So far, we have implicitly assumed that the contents in the equivalent bottleneck have to be *causal* in the sense that – starting from $X_0 = 0$ – the accumulated amount of traffic at the outlet cannot exceed the accumulated amount of traffic at the inlet:

$$\sum_i R_i^{\text{in}} \Delta T \geq \sum_i R_i^{\text{out}} \Delta T \quad \forall i. \quad (7)$$

The causality principle (7) prevents the content of the equivalent bottleneck from becoming negative ($X_i \geq 0 \forall i$).

Any violation of (7) is easily detectable through $X_i < 0$, and it has to be corrected as it implies the risk for erroneous estimations of T^{add} . A potential reason for negative content in the equivalent bottlenecks may be a desynchronization between the measurements on inlet and outlet, entailing $X_0 > 0$. Indeed, a correction is easily performed by

$$X_0 = X_0 + |X_i| \quad \forall i : X_i < 0, \quad (8)$$

and estimations of T^{add} for intervals ahead of i have to be recalculated if necessary.

In case of loss, X_i in (5) includes even the loss L_i encountered in interval i . In fact, COIA considers lost traffic to remain in transit forever. Reflecting further on the above example, we now assume that the traffic stemming from $D_i > 0$ is lost. This would imply $D_{i+1} = 0$ and $X_j \geq D_i \Delta T \forall j > i$. Actually, there remains a residual content of $L_i = D_i \Delta T$ in the equivalent bottleneck, which appears the same way as a permanent delay of ΔT . Equation (5) does not allow to distinguish between these situations; the only indication for loss might be the appearance of a positive trend in the values X_i .

In contrast to violations of the causality principle, loss L_i needs to be discovered on a higher layer (e.g. through missing sequence numbers) and corrected through

$$X_i = X_{i-1} + D_i \Delta T - L_i \quad (9)$$

in order to not to introduce a permanent positive bias of the estimation of T^{add} .

2.4 Related Work

Related work areas attempt to model and estimate additional delay for data exchanged in a distributed data communication systems. Namely, with regards to modelling efforts, we recognize that the equation (5) can be seen as the fluid flow version of Lindley's recursion formula [20]; it expresses the waiting time at the end of an interval as a function of the waiting time at the beginning of the interval and the number of arrivals or departures during that interval. Based on the Lindley's formula, [21] investigated delay in a constant link capacity network, while our observations relate to the user-perceived capacity of mobile links.

The probably most cited paper in the domain of fluid flow modelling and analysis is [8], providing a closed-form analytical description of the buffer content of a fluid buffer of an unlimited size, fed by homogeneous on-off sources. However, the authors of [8] do not model the output process of the equivalent bottleneck; this is rudimentarily done in [22]. In [23] we have presented a simple yet effective analysis of the bit rate distribution arising from the Anick-Mitra-Sondhi-type equivalent bottleneck [8] and thus describe the effect of the bottleneck in terms of changes of bit rate histograms. The main idea in [23]—deriving information on the bottleneck behavior from a comparison of bit rate histograms at inlet and outlet of a bottleneck—was demonstrated in our previous work through a measurement study of video conferencing traffic [18] and subsequently a measurement study in mobile networks [19]. Both references implicitly used the COIA method.

In principle, the COIA method builds upon Little's Law [24], which however is very general and estimates the total end-to-end delay spent by a packet in the system (*i.e.*, including transmission, propagation, queuing and additional delays). The COIA methods use refined versions of Little Law to estimate merely additional delay values for a packet from a sender or receiver viewpoint. The authors of [25]–[27] have used Little Law

as a base for modeling the total delay of packets in their systems. However they assume system characteristics that are unattainable in reality, and they do not focus on different delay views, as we propose. Namely, in [25], the authors use Little's Law to estimate the total time packets spend in the system, assuming that the number of nodes in a network, a total number of packets exchanged by nodes and total bandwidth available for each node are *a priori* known. They estimate delays for video, audio, voice and messaging data packets and then use the estimated delay values in their proposal on traffic priority schemata. Similarly, the authors of [26] used Little's Law to model delay of MAC-level frames exchanged between nodes connected via a WLAN network. They assume exponential frame inter-arrivals and service times. The authors of [27] model a mean total delay spend by a frame in a flow, assuming a fair scheduling at mobile MAC-layer for different flows and fair capacity sharing at base stations by a mobile network operator.

With regards to modelling or measurements of additional delay, sometimes also denoted by authors as an overall queuing delay, we consider related work on different protocol stack layers for (mobile) nodes, from the MAC-layer, via IP up to the TCP/UDP transport layer. For example, the authors of [28] propose an enhanced MAC layer, *i.e.*, a 'Virtual MAC', that continuously and passively monitors interference at the MAC layer, interprets frames and derives estimates for each frame's queuing delay and, based on these estimates, differentiates servicing of frames depending on their application flows, *e.g.*, real-time voice or non real-time data. The authors of [29] focus on IP-level, per-datagram delay estimation for 2.5G/3G mobile operator networks, assuming a detailed knowledge on statistical characteristics of the radio channel and parameters of the MAC-layer quality control techniques used in the network, and the size of a transported IP datagram. Based on the network's delay estimation, they propose an adaption of size of the transported IP datagrams.

The authors of [30], [31] focus on TCP-level messages delay estimations. Namely, [30] model the messages' queuing delay distribution (using a finite state machine) for Internet traffic, however, assuming a message being sent only every RTT. The authors of [31] propose an estimation of TCP-level RTT for web-based browsing (voice and data traffic) assuming the node in a 3G network knows its link utilization level. The authors of [32] propose measurements of queuing delays experienced by VoIP UDP-level messages via observing their inter-packet time at the receiver. They assume that the first packet in the VoIP application flow has been received without (or with a minimum) queuing delay, and that a baseline inter-packet time can be derived from it. However, this inter-packet monitoring method is resource-intensive and error-prone in mobile devices, due to clock's synchronization and resolution issues, especially for intense flows, where many packets are being sent or received in a second. In contrary, our

methods quantify system behavior at the given time interval, balancing these factors. Moreover, we propose methods for receiver's, as well as the sender's view.

As a related work we also distinguish the proposal of [33], aiming to use a covert channel, where unused bits of the IP datagram would transfer a coded timestamp from sender to receiver, from which one-way delays and thus additional delays can be estimated. While this approach is passive in the sense that no extra traffic is created, it requires the modification of each and every data packet, which is infeasible given the limited API to the protocol stack, and limited computational resources of a mobile device. On the other hand, [34] proposes to derive one-way delays from flow data available in routers. Both papers address the issue of timestamp accuracy that emerges from different kinds of quantification issues as a key challenge for one-way delay estimations. The work of [6] supports a motivation for our proposal, as it provides a measurement study that compares the performance of 3G and 3.5G networks in both stationary and mobile settings, where it underlines the critical issue of delay spikes, observed in one-way delay measurements in a testbed, and presents their related statistics. Estimation of additional delays is not addressed.

In summary, literature conveys the picture that the challenges of delay spikes, the issue of correct timestamping and the necessity of estimating additional delays are recognised, however, they are not treated in combination for mobile applications. Thus, this article and its predecessor [35] close this gap.

3 ESTIMATION METHODS

This section provides methods for estimations of the additional delay T^{add} based on throughput information at the inlet and outlet of the equivalent bottleneck representing the end-to-end network path. Given the limited communication, processing and storage capabilities of mobile devices, and a need of timely estimations, *simplification* of the estimators is a major point of our concern. In particular, the parameters used for the estimation shall be considered in close time proximity to the current interval in order to make the approach as *stateless* as possible. For interval i , the latter constraint limits the scope to X_{i-1} , X_i , R_i^{in} , R_i^{out} and D_i .

The proposed methods implement the COIA principle and thus require the exchange information about throughput (R_i^{in} , R_i^{out} or D_i) or of buffer level X_i between inlet and outlet (the implications for the implementation of the methods are discussed in Section 6). Therefore, we assume that either

- (a) R_i^{in} can be sent from the sender towards the receiver, where (6) and (5) are calculated; or
- (b) R_i^{out} can be sent from the receiver towards the sender, where (6) and (5) are calculated; or
- (c) the receiver can observe X_i , e.g. from the deviation of a jitter buffer from its reference value, and then calculate D_i from (5) and R_i^{in} from (6); or

- (d) the receiver might estimate the sender's average rate $\mathbb{E}[R_i^{\text{in}}]$ e.g. through the mean of R_i^{out} , which converges to $\mathbb{E}[R_i^{\text{in}}]$ in case of negligible loss, or use a-priori-known rates, e.g. from earlier measurements.

3.1 Sender View Ahead (SVA)

The sender is concerned about whether the data sent in the current interval i has been received without experiencing additional delay. This means that at an arbitrary time, there should not be any (bottleneck) content left in the network. If there would be any content left, it would be visible at the end of the current averaging interval i as $X_i > 0$. Seen from the sender point of view, which is actually expecting a throughput of R_i^{in} , this amount $X_i > 0$ is considered to be late by a time of

$$T_i^{\text{add}} = \frac{X_i}{R_i^{\text{in}}}. \quad (10)$$

As in this method the sender considers its "left-over" for the next interval $i+1$, we call this method *Sender View Ahead*. If the current throughput vanishes, i.e., $R_i^{\text{in}} = 0$, the estimation is undefined.

3.2 Sender View Backwards (SVB)

In another but similar view, the sender is concerned about the impact of "left-over" data from the most recent interval $i-1$ onto the current interval i , while its current expected throughput is R_i^{in} . The *Sender View Backwards* method for the estimation of additional delay is thus defined as

$$T_i^{\text{add}} = \frac{X_{i-1}}{R_i^{\text{in}}}. \quad (11)$$

For constant sender throughput, SVB produces same estimation series as SVA, however moved by one interval, i.e., $\{\hat{T}_{\text{SVB},i}^{\text{add}}\}_i = \{\hat{T}_{\text{SVA},i-1}^{\text{add}}\}_i$. $R_i^{\text{in}} = 0$ leads to an undefined estimation.

3.3 Receiver View Backwards (RVB)

The receiver is concerned about whether the data received in the current interval i has experienced queuing. Such queuing is seen from a non-empty bottleneck at the end of the previous interval $i-1$, i.e., $X_i > 0$. Observing a left-over at the end of the interval $i-1$, the receiver can estimate the transport time needed for this outstanding data based on the current receiver throughput R_i^{out} . The *Receiver View Backwards* method for the estimation of T^{add} is thus defined as

$$T_i^{\text{add}} = \frac{X_{i-1}}{R_i^{\text{out}}}. \quad (12)$$

Replacing R_i^{out} by the nominal link capacity, this method becomes the one used in [20]. Vanishing output $R_i^{\text{out}} = 0$ yields an undefined estimation.

3.4 Maximum of SVA, SVB and RVB (MAX)

Due to the different points of view and depending on throughput and buffering content values, the methods SVA, SVB and RVB are likely to provide different estimations of the additional delay. A pragmatic approach consists of using the most pessimistic estimation, which is the maximum of the three estimations obtained with SVA, SVB and RVB, given as

$$\hat{T}_i^{\text{add}} = \max \left\{ \hat{T}_{\text{SVA},i}^{\text{add}}, \hat{T}_{\text{SVB},i}^{\text{add}}, \hat{T}_{\text{RVB},i}^{\text{add}} \right\}. \quad (13)$$

Undefined values in the argument of the maximum operator are ignored; in the worst case, (13) cannot provide any estimation.

3.5 Mean Sender View Ahead (MSVA)

As a variant of SVA method, this method is designed for a receiver that may estimate the sender's average rate $\mathbf{E}[R_i^{\text{in}}]$ instead of the actual value of R_i^{in} as described above, as a basis for estimation of the bottleneck content X_i . The *Mean Sender View Ahead* method for the estimation of T^{add} is defined as:

$$T_i^{\text{add}} = \frac{X_i}{\mathbf{E}[R_i^{\text{in}}]}. \quad (14)$$

4 VALIDATION: HEALTH TELEMONTORING

In this section, we investigate to which extent the proposed methods are able to estimate T^{add} based on time-synchronized, lossless data traces collected along the execution of a mobile health application. In addition, we evaluate the methods' accuracy under a variety of network conditions.

4.1 Setup

The MobiHealthTM system [9], [10] has been developed and used in the EU-FP6 MobiHealth project and it enables real-time telemonitoring of vital signs (*e.g.*, ECG) and context (*e.g.*, location) of the mobile patients. A patient wears a Body Area Network (BAN) with an a Mobile Base Unit (MBU) as a central unit, acquiring data from wireless sensor system(s), processing it (*e.g.*, deriving heart rate) and sending to a backend-system (BEsys) in, *e.g.*, a hospital. The end-to-end communication path is heterogeneous and includes, *e.g.*, 2.5G or 3G access network provided to a patient. Tele-monitoring service is supported by the proprietary TCP/IP-based *MSP-Interconnect Protocol* (MSP-IP) [36] and conforms to the Jini Surrogate specification [37]. The application-data delivery is both lossless and in-order but may suffer from additional delays. For performance evaluation of the MobiHealthTM system please refer to [16], [38].

For our studies, we exploit data traces derived from telemonitoring of a (hypothetical) cardiac patient living at the campus of the University of Twente (the Netherlands), using MBU¹ with 3G-UMTS access network of

Vodafone [16] and BEsys being a high-performance server in the university network. We encapsulated application-data in fixed size (524 B) TCP message payload, and controlled the send rate to 7, 8, 9, 11 or 12 packets/second; *i.e.*, resulting in an uplink rate of 32.6 to 55.9 kbps (given the the overhead of 58 B per packet²). We have also controlled the application and the TCP buffer sizes, along the combinations: 64/64 KB; 32/64 KB and 32/32 KB. The MBU and BEsys clocks were synchronized using Simple Network Time Protocol (SNTP)³ over an external, dedicated Ethernet connection. We timestamp each sent and received application-data packet with a potential inaccuracy of ± 20 ms [14]. At the sender side (*i.e.*, MBU), we collect throughput time series $\{R_i^{\text{in}}\}_{i=1}^n$ at the ingress boundary of the TCP socket (*i.e.*, the network inlet, after the 'send' function), and collect $\{R_i^{\text{out}}\}_{i=1}^n$ at the egress boundary of the TCP socket (*i.e.*, network outlet, after the 'receive' function) at the receiver side (*i.e.*, BEsys). Five different data rates, three combinations of buffer sizes, five replications of each experiment and one trace affected by a hardware crash, left us with 74 traces on which we conducted a validation of the proposed additional delay estimation methods.

4.2 Illustrative Example

We present an example of the causal relation between network (*i.e.*, the equivalent bottleneck) behaviour and the telemonitoring application-data traces for 8 pps (125 ms inter-packet time), and application and TCP buffer sizes of 32 KB, respectively. In this example, the equivalent bottleneck content increases over a period of two seconds and then rapidly decreases (*i.e.*, the queue releases). We investigated how the proposed methods estimate the additional delay. First, we present in Fig. 1 the behavior of sender and receiver and its relation to the value of T^{add} , captured by the data series "Actual delay". The *x*-axis displays the time intervals under consideration, and the *y*-axis shows its corresponding T^{add} value as derived from the timestamps in the sender (MBU) and the receiver (BEsys) application-data traces. Each data point represents a packet as registered on the receiver side.

The packet sent in interval 119 experiences the maximal additional delay of $T^{\text{add}} = 712$ ms, which amounts to almost six nominal inter-packet times (Fig. 1). We can see that this packet is suddenly released at the receiver together with five subsequent packets being queued. This behaviour, which is quite common for mobile links, can be explained as follows: The first packet was corrupted or lost while being in transit and is retransmitted, while the subsequent packets are held until the retransmission of the first packet was successful.

Fig. 2 presents the data rate at (a) inlet (MBU) and (b) outlet (BEsys), and (c) the bottleneck content for intervals

1. Asus laptop (MPIII 1 GHz proc., 640 MB RAM, WinXP OS), using a Nokia 6650 phone as a USB modem to 3G-UMTS network.

2. MSP-IP (10 B), TCP/IP (40 B) and PPP (8 B) overhead.

3. SNTP: <http://www.ntp.org/ntpfaq/NTP-s-def.htm>.

Fig. 1. Estimated additional delays for the sender and receiver view for $\Delta T = 100$ ms.

Fig. 2. Equivalent bottleneck: (a) inlet, (b) outlet and (c) content.

as in Fig. 1. The sender has a regular pattern of sending data in four out of five intervals, which matches the ratio between ΔT and the inter-packet time. The reception is however quite bursty with no data being received during intervals 119 to 125. During the interval 126, the receiver gets hold of all six outstanding packets. Afterwards, it continues receiving a data stream during intervals 127 to 129. The content of the equivalent bottleneck shows the increase and a release of the data.

Fig. 1 presents estimated additional delay values for the five proposed methods. The methods SVA and SVB provide increasing estimations of T^{add} , reaching values of 600 ms as they react upon the growing bottleneck content. SVB with its backwards point of view is typically one interval “late” with its estimations. A vanishing input in intervals 114, 119, 124 and 129 makes the results undefined, cf. (10) and (11). When the bottleneck grows, the RVB method performs very poorly if there is no observed data at the receiver, e.g. in intervals 119 to 125;

thus, the result of formula (12) is undefined. Only in the interval 127, method RVB provides T^{add} estimations that relate to the measured values.

For interval 124, as there is data neither sent nor received, none of the methods SVA, SVB or RVB are able to estimate the additional delay. In the other intervals, the MAX method provides estimations of an increasing value, as it uses the maximum value of SVA, SVB and RVB, where especially the first two react upon the growing bottleneck content.

The MSVA method delivers quite conservative estimations. It takes the estimated sender behaviour instead of the real one into account. Hence, it assumes that 420 B are sent in each interval of $\Delta T = 100$ ms. The method accounts this data for bottleneck content, which results in an over-estimation of additional delays. Similarly to SVA, this method reacts upon the growing bottleneck content. Just before releasing the queue, MSVA estimates $\hat{T}_{\text{MSVA},126}^{\text{add}} = 819$ ms, which is to be compared to real value of $T^{\text{add}} = 712$ ms and the SVA-based estimation of $\hat{T}_{\text{SVA},126}^{\text{add}} = 600$ ms. For the interval 124, MSVA is the sole method able to provide an estimation of $\hat{T}_{\text{MSVA},124}^{\text{add}} = 619$ ms.

4.3 Accuracy of the Estimations

Having examined in details an illustrative example in the previous section, we now present the cumulative results for the accuracy of estimations of additional delays along all 74 application-level traces collected by the telemonitoring application that serve as points of reference. From the sender (MBU) and receiver (BEsys) timestamps, we derive for each trace T^{min} which is defined as the minimum delay value ever occurred in that trace. Then, for each packet p of this trace, we derive its (measured) additional delay T_p^{add} . The *maximum additional delay* ever occurred in the trace is denoted as:

$$T_{\text{max}}^{\text{add}} = \max_p \{T_p^{\text{add}}\} = \max_p \{t_p^{\text{out}} - t_p^{\text{in}} - T^{\text{min}}\}. \quad (15)$$

Then, for each method M , we calculate its maximal estimated additional delay in the trace as

$$\hat{T}_{M,\text{max}}^{\text{add}} = \max_i \{\hat{T}_{M,i}^{\text{add}}\}. \quad (16)$$

The relative estimation error is defined as

$$e_M = \frac{\hat{T}_{M,\text{max}}^{\text{add}} - T_{\text{max}}^{\text{add}}}{\Delta T}. \quad (17)$$

We have observed in all traces that a typical value of $T_{\text{max}}^{\text{add}}$ is found in the order of magnitude of 300 ms, with some exceptions reaching up to 712 ms as shown before. Fig. 3 presents the cumulative distribution function (CDF) of the relative estimation error e_M of all five methods SVA, SVB, RVB, MAX and MSVA for $\Delta T = 100$ ms, 300 ms and 1 s, respectively. The x -axes display the estimation error in % of ΔT .

Fig. 3. TCP: Empirical CDF of the relative estimation error e_M for $\Delta T \in \{(a) 100 \text{ ms}, (b) 300 \text{ ms}, (c) 1 \text{ s}\}$.

For $\Delta T = 100 \text{ ms}$ (Fig. 3.a), the estimation error for SVA ranges from $-1.2\Delta T$ to $1.1\Delta T$, for SVB from $-1.5\Delta T$ to $1.4\Delta T$ and for MAX from $-1.2\Delta T$ to $1.4\Delta T$. In most cases, the error is bounded by $\pm\Delta T$. RVB underestimates the additional delay by $-5\Delta T$ to $-0.5\Delta T$. This is due to its inability to trace growing content in the equivalent bottleneck, if there is no data being received (cf. Section 4.2). The MSVA displays error values between $-0.3\Delta T$ to $2.2\Delta T$; it has a clear tendency to overestimate the additional delay (cf. also Section 4.2).

For $\Delta T = 300 \text{ ms}$ (Fig. 3.b), the estimation error for SVA ranges from $-0.8\Delta T$ to $0.7\Delta T$, for SVB from $-0.9\Delta T$ to $0.8\Delta T$ and for MAX from $-0.9\Delta T$ to $0.8\Delta T$. The majority of absolute errors is found in the range of $-\frac{1}{2}\Delta T \leq e_M \leq \frac{1}{4}\Delta T$. RVB exhibits errors in the range of $-1.7\Delta T$ to $0.2\Delta T$; it again shows a tendency to under-estimate the additional delay. The MSVA has an error range of $-0.4\Delta T$ to $0.7\Delta T$, with a less pronounced tendency to overestimate the value because $\mathbb{E}[R^{\text{in}}]$ is closer to the actual R^{in} values on this time scale. For any of the methods, the relative error has decreased significantly at this time scale, whose choice coincides with the typical value of $T_{\text{max}}^{\text{add}}$ in the order of 300 ms.

For $\Delta T = 1 \text{ s}$ (Fig. 3.c), the most precise estimation methods are SVA and MSVA, with relative estimation errors ranging from $-0.6\Delta T$ to $0.04\Delta T$ (SVA) and to $0.12\Delta T$ (MSVA). Yet, 99 % of the e_M values for both methods are found in range of -10% to 4% (SVA) and to 12% (MSVA). The reason for such a high accuracy of the latter method is as follows. In the telemonitoring applica-

tion, the sender is precisely timing the one-second time intervals during which it sends an integer number of packets. Hence, $\mathbb{E}[R^{\text{in}}] = R^{\text{in}}$ and that results in a precise estimation of the additional delays. This particular case shows that the MSVA method can be beneficial if ΔT corresponds to an integer multiple of the nominal inter-packet time. From Fig. 3, we also conclude that SVB and RVB and hence MAX have a tendency to overestimate the additional delays. The estimation error ranges from $-0.6\Delta T$ to $0.8\Delta T$.

In general, the best estimations are delivered by the SVA method, with a tendency to underestimate the additional delay, and by the MSVA method if the averaging interval matches the periodicity interval of the traffic. On the time scale $\Delta T = 100 \text{ ms}$, which is considerably smaller than the majority of additional delays in our traces, the estimation error is more or less limited by this time scale, with exception of the method RVB. Indeed, in most cases, we perceive errors that are smaller than the time scale itself. Obviously, the methods even allow for observation of additional delays that are smaller than the time scale on which the measurements are carried out. This is observed because packets are occasionally “pushed” from one interval to the subsequent one, thus the methods can detect it. If ΔT is small, changes in T^{add} are significant, and the shift of a single packet by a few milliseconds can easily yield a much larger predicted additional delay and thus a significant relative error. When the interval size increase, it is still possible to detect a number of packets being delayed from the

one interval to the other and estimate an additional delay, which is a fraction of the time scale ΔT itself. At the same time, the longer averaging time also helps to reduce the relative contribution of a potentially shifted packet to the estimation error.

4.4 Discovery of Additional Delay Spikes

In this section we investigate the accuracy of the proposed estimation methods for T^{add} to indicate *additional delay spikes*, in a trace. A spike occurs when T^{add} is higher than a certain threshold defined based on application's requirements for delays. For example, as illustrated in Fig. 1, the value of 712 ms is a T^{add} spike, for the threshold of 300 ms, derived from the health telemonitoring application-level delay requirements.

In that particular example, the spike extended over several consecutive observation intervals ΔT and was thus easily located. However, a spike might also occur inside an arbitrary interval. Consequently, X_i is not affected, and because of this, the spike goes unnoticed. Such an uncaught spike is henceforth called *false negative*. On the other hand, due to the approximation nature of the proposed methods, their additional delay value might overstep the pre-defined threshold for a spike, while the actual additional value delay might not. We denote this case as *false positive*.

Let us define a set of approximation performance parameters for an observation time scale ΔT as follows:

- σ^{real} as the total number of spikes (so-called true positives) existing in a set of traces;
- σ_M^{ind} as the number of spikes indicated by method M ;
- σ_M^{FP} as the number of false positives indicated by method M ;
- σ_M^{FN} as the number of false negatives indicated by method M ;
- $\sigma_M^{\text{TP}} = \sigma^{\text{real}} - \sigma_M^{\text{FN}} = \sigma_M^{\text{ind}} - \sigma_M^{\text{FP}}$ as the number of true positives indicated by method M ;
- $\gamma_M = \sigma_M^{\text{TP}} / \sigma^{\text{real}}$ as the spike hit ratio for method M .

We consider a threshold for spikes of 300 ms as derived from health-telemonitoring application requirements and we find $\sigma^{\text{real}} = 45$ spikes in the investigated set of 74 traces. For $\Delta T = 100$ ms, $\Delta T = 300$ ms, and $\Delta T = 1$ s, Table 1 presents the total number of spikes indicated by the methods; the numbers of false positives; the numbers of false negatives; and the hit rate for each combination of method and time scale, respectively.

The SVA, SVB, MAX and MSVA methods indicate many spikes and false positives for $\Delta T = 100$ ms, while those numbers decrease for $\Delta T = 300$ ms and 1 s. On the other hand, the number of false negatives rises for these four methods as ΔT grows. Despite those general trends, the numbers for the different methods and time scales differ. As compared to SVA, the three methods SVB, MAX and MSVA indicate more spikes and false positives. For the short time interval $\Delta T = 100$ ms, SVB misses most spikes, while MSVA does not miss any.

TABLE 1
TCP: Performance of additional delay spikes indication by different estimation methods ($\sigma^{\text{real}} = 45$).

	ΔT	SVA	SVB	RVB	MAX	MSVA
σ_M^{ind}	100 ms	125	129	4	147	176
	300 ms	51	91	21	111	79
	1 s	6	19	23	27	13
σ_M^{FP}	100 ms	85	94	0	107	131
	300 ms	23	61	9	72	45
	1 s	2	9	13	15	4
σ_M^{FN}	100 ms	5	10	41	5	0
	300 ms	17	15	33	6	11
	1 s	41	35	35	33	36
γ_M	100 ms	89 %	78 %	9 %	89 %	100 %
	300 ms	62 %	67 %	27 %	87 %	76 %
	1 s	9 %	22 %	22 %	27 %	20 %

The method RVB behaves differently, having a few indications, but no false positives on the short time scale, and rising numbers of those on the longer time scales. This behaviour stems from the difficulty that RVB has with estimating additional delays when the bottleneck is increasing and nothing is received, cf. Fig. 1. The T^{add} for data inside the bottleneck content cannot be estimated by this method. The method gets more accurate as ΔT grows, because there are fewer intervals with vanishing receiver throughput. Still, the number of false negatives is high almost independently of ΔT .

From the methods' hit ratios on the different time scales we observe the same trends as already indicated. The sender-based methods show large hit rates between 78 % (SVA) to 100 % (MSVA) for $\Delta T = 100$ ms. For the time interval matching the delay threshold $\Delta T = 300$ ms, between 62 % (SVA) and 87 % (MAX) of the real peaks are discovered. For this ΔT , also RVB shows its largest hit ratio of 27 %. On the time scale of $\Delta T = 1$ s, all the hit ratios are found between 9 % (SVA) and 27 % (MAX).

Summarizing the findings, we can see that the sender-based methods SVA, SVB and MSVA are superior to the receiver-based method RVB unless the time interval ΔT clearly exceeds the duration of the spike. SVB and MAX are more indicative than SVA. While for small ΔT , the sender-based methods have a tendency to overestimate the delay values (resulting in false positives), all methods tend to not notice spikes in case of the large ΔT , which is somehow expected given the fact that the threshold for spikes is merely 30 % of the observation interval. This behavior stems from the fact that, in order to accurately detect a spike of a given value, ΔT must be of a smaller size than this spike. The larger ΔT , the more spikes smaller than ΔT go unnoticed. The choice of ΔT shall be related to the designated levels of delay spikes to be discovered, when using the methods.

5 VALIDATION II: MULTIMEDIA STREAMING

In this section we provide evaluation results along the same goals and approach as already presented in Section 4, but for UDP-based application traces. Therefore,

in this section, we investigate to which extent the proposed methods are able to correctly estimate T^{add} and delay spikes. We evaluate the accuracy of each method against time-synchronized data traces collected at sender and receiver along the execution of a multimedia streaming application provided to a mobile user.

5.1 Setup

The system used for the evaluation of methods provides a multimedia streaming to and from a mobile user with use of a UDP-based application protocol; the user downloads or uploads some multimedia content. For our studies, we exploit the traces obtained from an application used on a mobile laptop, by a user located in Karlskrona (Sweden) and using his application in one specific location (Blekinge Institute of Technology). The laptop is connected to 3G-UMTS networks of three different mobile operators, two of which share the radio part of their radio access networks.

The application was configured to send data streams using different multimedia coding; resulting in 1 up to 25 packets per second, with packet sizes of 64, 256, 512 or 750 B, being sent separately to the network at inter-packet times from 4 ms to 256 ms. The protocol stack overhead⁴ is 36 B and the overall data rate uploaded by the mobile laptop to the application server ranges from 8 to 360 kbps, while data rate downloaded to the mobile laptop ranges from 8 kbps to 2 Mbps. Instead of collecting the data in the laptop and server, we opted to collect the from the data link layer using Endace DAG cards. The setup is similar to the one shown in [39]. This perfectly synchronised system, with an accuracy of less than 60 ns, provides an ideal basis for our validation endeavour. At the ingress boundary of the UDP socket (*i.e.* at heterogeneous network infrastructure inlet) at the laptop side, we collect throughput time series $\{R_i^{\text{in}}\}_{i=1}^n$, while $\{R_i^{\text{out}}\}_{i=1}^n$ is collected at the egress boundary of the UDP socket (*i.e.* at the network outlet) at the application server side. All the packets included a sequence number that enables us to trace losses and packet reordering. We consider 33 traces as a base for the accuracy validation for the proposed methods. Traces exhibited occasional, random packet losses (0.2%) and no packet reordering. Similarly to the TCP case, we observed the typical value of $T_{\text{max}}^{\text{add}}$ around 300 ms.

As for the TCP-based traces evaluation, we present the methods' accuracies of the estimated maximal delay at different time scales (Section 5.2) as well as the evaluation of spikes of additional delays (Section 5.3).

5.2 Accuracy of the Estimations

The Fig. 4 presents the cumulative distribution function (CDF) of the relative estimation error e_M of all five methods SVA, SVB, RVB, MAX and MSVA for $\Delta T = 100$ ms,

300 ms and 1 s, respectively. The x -axes display the estimation error in % of ΔT .

For $\Delta T = 100$ ms (Fig. 4.a), the relative estimation error for SVA is bounded by $-1.5\Delta T$ to $0.65\Delta T$ and for SVB by $-2.4\Delta T$ to $1.5\Delta T$. For RVB, the error stretches from $-5\Delta T$ to $10\Delta T$. This is due to its inability to trace growing content in the equivalent bottleneck, if there is no data being received, and due to over-estimating delays if very little traffic is received compared to what has been sent. The MSVA has a tendency to overestimate the additional delay. It is caused by two factors, both resulting in an effective sender throughput being smaller than the estimated sender's average throughput. First, the sender does not follow a very regular sending pattern (even at the time scale of 1 s); hence, the estimated sender's throughput may be significantly lower or higher than the instant one. Second, the traces exhibit occasional losses; for such events, the estimated sender's throughput is higher than the instant one.

For $\Delta T = 300$ ms (Fig. 4.b), the relative estimation error for SVA is bounded by $\pm\Delta T$ and for SVB by $-1.1\Delta T$ to $0.2\Delta T$. The methods RVB and MSVA methods show a tendency to over-estimate the additional delay for the reasons mentioned in previous paragraphs. In general, for all of the methods, the relative estimation error has decreased significantly at this time scale.

For $\Delta T = 1$ s (Fig. 4.c), all the estimation methods but MSVA sufficiently accurate the e_M values for the SVA, SVB, RVB and MAX methods are in range of $-1.0\Delta T$ to $0.5\Delta T$. At this time scale, the phenomenon of a growing content in the equivalent bottleneck is not as pronounced as at smaller time scales; hence, the RVB method exhibits a high accuracy. Moreover, at this time scale, the phenomenon of the sender's imprecise behaviour influences less the accuracy of the MSVA method.

From all the three figures, we conclude that the relative accuracy of SVA, SVB and RVB and hence even MAX methods increases with the increasing time scale. Yet, especially at small time scales, all these methods have a slight tendency to overestimate the additional delays. This overestimation is in many cases related to the high variability of the throughput being received. Namely, there exist intervals in which very small amount of traffic is received with comparison to what has been sent, but the traffic being in transit does not necessarily experience delays peaks, as the methods attempt to indicate. The latter behavior relates to the fact that methods relay their estimation of T^{add} on counting packets that are pushed from one interval to the subsequent one. The smaller the intervals, the greater the potential errors become, while larger intervals help to average out some issues, however at the price of an increased risk of missing delay spikes.

5.3 Discovery of Additional Delay Spikes

Considering a threshold for spikes of 300 ms, we find 10 spikes in the investigated set of 33 application-level traces. For $\Delta T = 100$ ms, $\Delta T = 300$ ms and $\Delta T = 1$ s,

4. It includes the UDP, IP and PPP overheads.

Fig. 4. UDP: Empirical CDF of the relative estimation error e_M for $\Delta T \in \{(a) 100 \text{ ms}, (b) 300 \text{ ms}, (c) 1 \text{ s}\}$.

Table 2 presents the total number of spikes indicated by the methods; the numbers of false positives; the numbers of false negatives, and the hit rates for each combination of method and time scale, as introduced in Section 4.4.

All methods indicate many spikes for $\Delta T = 100 \text{ ms}$, while the numbers decrease for $\Delta T = 300 \text{ ms}$ and 1 s . The number of false negatives rises as ΔT grows for these methods. The methods do not indicate spikes as the time scale grows, as these peaks happen inside the observation interval. Despite those general trends, the numbers for the different methods and time scales differ. The methods SVA and SVB exhibit the highest accuracy.

The method MSVA behaves in a different way. We observe a lot of spikes, many of which are false positives, especially on the short time scale. In general, the accuracies decrease on the longer time scales. This behaviour stems from the inaccuracy of MSVA throughput estimations when the sender is not adhering to regular sending patterns and when data is occasionally lost.

Looking at the methods' hit ratios on the different time scales, we observe the same trends as indicated in the previous sections. The SVA, SVB, RVB and MAX methods show high hit rates, especially on a scale $\Delta T = 100 \text{ ms}$. The MAX method exhibits an accuracy of 100 %. For the time interval matching the pre-defined threshold $\Delta T = 300 \text{ ms}$, between 50 % (RVB) and 70 % (SVA, MAX) of the real peaks are discovered. Finally, on the time scale of $\Delta T = 1 \text{ s}$, all hit ratios of the methods are 30 %.

As already indicated in Section 4.4, this behavior stems from the fact that, in order to detect a delay spike of a given value, the interval ΔT must be shorter than this

TABLE 2
UDP: Performance of additional delay spikes indication by different estimation methods ($\sigma^{\text{real}} = 10$).

	ΔT	SVA	SVB	RVB	MAX	MSVA
σ_M^{ind}	100 ms	11	12	10	17	166
	300 ms	14	8	11	15	45
	1 s	5	5	5	5	27
σ_M^{FP}	100 ms	2	3	3	7	157
	300 ms	7	2	6	8	38
	1 s	2	2	2	2	24
σ_M^{FN}	100 ms	1	1	3	0	1
	300 ms	3	4	5	3	3
	1 s	7	7	7	7	7
γ_M	100 ms	90 %	90 %	70 %	100 %	90 %
	300 ms	70 %	60 %	50 %	70 %	70 %
	1 s	30 %	30 %	30 %	30 %	30 %

spike, hence making it possible to discover packets that are pushed from one interval to the subsequent. This indicates that the choice of ΔT shall be related to the designated levels of delay spikes to be discovered using the proposed methods.

Summarizing the findings, we can see that the MAX method is the most indicative as it combines strengths of all the underlying methods: SVA, SVB and RVB. While for small ΔT , the sender-based methods have a tendency to under- or over-estimate the delay values (and produce false negatives and positives), all methods tend to overlook spikes in case of large ΔT , which is somehow natural given the fact that the threshold for spikes is merely 30 % of the observation interval. As

already indicated in the Section 4.4, this behavior stems from the fact, that in order to detect a delay spike of a given value, the interval ΔT must be of a smaller length than this spike, hence enabling to discover packets that are “pushed” from one interval to the subsequent one, forming a spike. This indicates that the choice of ΔT shall be related to the designated levels of delay spikes, to be discovered when using the methods.

6 IMPLICATIONS FOR IMPLEMENTATION

In this section we draw conclusions and recommendations for a practical application of the proposed and evaluated additional delay estimation methods for real mobile application scenarios.

6.1 Impact of the Observation Interval

One important practical implication relates to the choice of an appropriate observation interval ΔT at which additional delays are estimated. The choice is limited by a trade-off between feasibility and effort required to instrument the application for throughput measurements at the sender and receiver side, and the requirements of discovering application-data additional delay at the given time scale. In many cases, this additional delay will be related to inter-packet times. We propose that an optimal ΔT shall be defined as a percentage of the required minimal end-to-end delay and, if possible, related to a integer multiple of the designated inter-packet time, in order to improve the accuracy of the proposed methods. If ΔT is too long, the methods loose possibilities to discover additional delays, shown by many false negatives. If ΔT is too short, the system becomes too sensitive, discovering many non-existing additional delay spikes, *i.e.*, many false positives. Such a situation can destabilize a system whose application control loops depend upon these estimations. The most important question one needs to answer when choosing the observation interval ΔT is how much additional delay is tolerated by the application before a control action, such as application adaptation, or a warning towards the user would be required.

6.2 Particularities of Specific Methods

The RVB method exhibited a low performance, *i.e.*, frequent underestimations of the additional delay, especially for a growing bottleneck content where the receiver does not receive any data. To improve this method’s accuracy in practice, one can combine it with other methods and assume that if the receiver does not receive any data, the bottleneck content and thus the additional delay grows. Then, one can use the MSVA method at the receiver to estimate the delay values.

The MSVA method exhibited a low performance, *i.e.*, overestimation of the additional delay, especially for cases when the sender’s current behavior deviated significantly from its average behavior, or in case of

traces with data loss. The estimations provided by this method can be treated as a worst case scenario; the other methods can be used for more precise estimations of additional delay.

6.3 Missing Measurement Data

COIA is relying upon comparison of performance data. In this subsection, we consider cases in which exchange of the data does not work as expected. We follow the classification presented in Section 3.

In cases (a) and (b), lacking throughput data should not be interpreted as zero values in the first hand due to the recursive nature of Equation (5). R_i^{in} values set to zero would drive the system towards incausality and underestimations of T^{add} , while R_i^{out} values set to zero would suggest persistent bottleneck content, loss and overestimations of T^{add} . A rather simple but straightforward mitigation strategy in case (a), which is also supported by the examples shown in Figures 3 and 4, is to increase the estimation of T^{add} by the time that R_i^{in} is overdue, as the latter might be strongly correlated with the fact that the data path is blocked. In case (b), such a practice might yield overestimations when the delay occurs on the return path. Indeed, complete loss of monitoring data puts COIA as such at stake. Eventually, data loss discovery on some higher level and the causality condition (7) can help to reconstruct the buffer content estimation. It is important to note that Equations (5) and (6) require data from the same interval to be compared. For this reason, it is advisable to include sequence numbers i in the messages containing the R_i values.

In case (c), the content X_i can be estimated locally *e.g.* from drain from a jitter buffer, which should be unproblematic in terms of availability of data. In case (d), the sender’s average rate may be estimated in an incorrect way. However, given fact that the MSVA method works best for streaming-type services in view of vanishing loss, the risk for crude receiver-side estimations appears limited.

6.4 Desynchronization

Incorrect synchronization of sender and receiver can imply residual buffer contents or (non-causal) buffer under runs and thus biased (over- and under-) estimations of T^{add} . Eventual *time synchronization* problems are typically seen through a bias of the content of the equivalent bottleneck X_i , and thus of T^{add} , that is constant over time and that might be corrected by modifying X_0 as described in Section 2.3. *Frequency synchronization* problems manifest themselves in $\Delta T^{\text{in}} \neq \Delta T^{\text{out}}$, and eventually in values of ΔT that vary over time. Assume a relative deviation between ΔT values of $\varepsilon (\ll 1)$, then a drift of approximately $\varepsilon T^{\text{add}} \ll T^{\text{add}}$ would be observed per time interval. A mitigation strategy consists in estimating the ΔT value used by the sender of the measurements, and to adapt the ΔT of the receiver of

the measurements. To help with the latter, the messages containing the R_i values might contain time stamps that allow for estimating the duration of interval i on the sender side.

A detailed investigation of the issues described in sections 6.3 and 6.4 would motivate a quantitative study of its own.

6.5 Mobility

The mobility of a mobile user may result in horizontal or vertical network handovers. Such handovers may cause significant impacts on the performance for data communication by increasing the additional delay. Naturally, the latter is revealed by all five methods, which means that COIA allows to quantify the performance impact of mobility, both regarding coverage and control. Furthermore, during handovers, both data and measurement traffic might get lost; the latter situation needs to be dealt with as described in Section 6.3. Finally, a failed handover might imply a connection loss and thus naturally a detectable termination of the COIA measurements.

7 CONCLUSIONS AND OUTLOOK

In this paper, we propose five simple, yet powerful methods for the estimation of additional delay for an application-level data induced by heterogeneous data communication network behaviour; *e.g.*, radio link problems, protocol stack processing or queuing phenomena. The methods are based on both sender and receiver observed application-level data rates and rather simple calculations of additional delay. Both sender and receiver implement the Comparative Output-Input Analysis (COIA) methods. The methods are based on lightweight exchange and comparison of throughput time series between sender and receiver. Using traces of four different network providers under a variety of traffic conditions for two different mobile applications, in particular transfer of health-care and multimedia data, each using different transport protocol for uploading or downloading of the data, we show that the most precise methods for additional delay estimation are based on the sender view on the content of the equivalent bottleneck and the current throughput of the sender. Furthermore, the pragmatic method that is assuming the maximum value of three sender- or receiver-view-based methods helps to yield conservative estimations of additional delays, especially for the TCP-based application. For the UDP-based applications that exhibits occasional packet loss, the receiver-view based estimation relying on the expected throughput of the sender yields conservative estimation of additional delay. The conservative estimations for additional delay may be of particular interest for selected safety- or mission-critical mobile systems.

The estimation error of all methods depends on the size of the observation interval for the sender and receiver throughput time series. This is determined by the granularity (*i.e.*, time resolution) of the events that need to be discovered, application instrumentation possibilities and capabilities of the mobile device to support

both the application and its COIA instrumentation. Our results show that “oversampling”, in terms of using an observation time interval smaller than the major additional delays to be observed, does not necessarily improve the quality of the estimation. Indeed, significant additional delay values can be observed and approximately quantified on time scales that are significantly larger than those of the additional delay phenomena themselves. This speaks in favour of the methods to be implemented in a constrained mobile environment, where the time interval choice and hence the frequency of such observations are limited by the resources of the mobile device. Here, the methods can help to quickly discover extraordinary additional delays and inform the application how late the data is. Thus, the methods can facilitate an implementation of lightweight monitoring and control loops for adapting an application’s data stream to volatile network conditions, *e.g.* by compressing or even suppressing low-priority data.

Furthermore, as we have shown in this paper, all methods have different but complementary profiles with regards to their sensitivity and failure or success regarding the estimation of additional delays at different time scales. To maximize the accuracy of additional delay estimation, all methods could be run in parallel and be combined in a hybrid method. This hybrid method could continuously track the estimations provided by all methods and implement a type of voting mechanism to decide upon the value of additional delay and decide whether there is a peak or not, based on knowledge of the strengths and weaknesses of the methods with respect to their tendency to indicate peaks (*i.e.*, their false positive or negative ratios).

Future work will address detailed investigations of the capability of the methods to detect arbitrary peaks of additional delay in other application domains. In this context, we will further analyze the impact of the observation interval size in relation to time scales used by the application. We will also address issues of systematic under- or over-estimations of additional delays, undefined estimations, sender jitter, data loss time synchronization, and clock drifts between sender and receiver.

ACKNOWLEDGMENTS

Support of the EU-COST Action IC0703, the Swedish KKS project QoEMoS (Dnr 2008/0520), and Swiss NSF (SER-C08.0025, PBGEP2-125917) is acknowledged.

REFERENCES

- [1] U. Hansmann, L. Merk, M. Nicklous, and T. Stober. *Pervasive Computing: The Mobile World*. Springer, 2003.
- [2] ITU-T. Rec. E.800, definitions of terms related to QoS and network performance including dependability, 2008.
- [3] ITU-T. Rec. P.10/G.100, vocabulary for performance and QoS.
- [4] D. Chalmers and M. Sloman. A survey of QoS in mobile computing environments. *IEEE Com. Surveys and Tutorials*, 2(2):2–10, 1999.

- [5] L. Zhang, Z. Liu, and C. Xia. Clock synchronization algorithms for network measurements. In *Proc. INFOCOM*, 2002.
- [6] J. Prokkola, P.H.J. Perala, M. Hanski, and E. Piri. 3G/HSPA performance in live networks from the end user perspective. In *Proc. IEEE ICC*, pages 1–6, 2009.
- [7] L. Kosten. Stochastic theory of a multi-entry buffer. *Delft Progress Report*, 1:10–18, 1974.
- [8] D. Anick et al. Stochastic theory of a data handling system with multiple sources. *Bell Syst. Tech. J.*, 61:1871–1894, 1982.
- [9] A. van Halteren et al. Mobile patient monitoring: the Mobihealth system. *J. on IT in Healthcare*, 2(5):365–373, 2004.
- [10] MobiHealth B.V. (www.mobihealth.com). "Putting care in motion", 2007–2012, visited on 25/02/12.
- [11] S. Tachakra, X. Wang, R. Istepanian, and Y. Song. Mobile e-health: the unried evolution of telemedicine. *Telemedicine Journal and e-Health*, 9(3):247–257, 2003.
- [12] P. Carlsson, D. Constantinescu, A. Popescu, M. Fiedler, and A. Nilsson. Delay performance in IP routers. In *Proc. HET-NETs*, 2004.
- [13] J. Nielsen. *Usability Engineering*. Morgan Kaufman, 1994.
- [14] K. Wac, P. Arlos, M. Fiedler, S. Chevul, L. Isaksson, and R. Bulst. Accuracy evaluation of application-level performance measurements. In *Proc. NGI*, Trondheim, Norway, 2007.
- [15] P. Arlos. *On the Quality of Computer Network Measurements*. Ph.D. thesis, Blekinge Institute of Technology, School of Engineering, SE, 2005.
- [16] K. Wac and R. Bulst. *Performance Evaluation of a Transport System supporting the MobiHealth BANip: Methodology and Assessment*. M.Sc. thesis, University of Twente, the Netherlands, 2004.
- [17] M. Claypool. The effect of latency on user performance in real-time strategy games. *Elsevier Computer Networks - Networking Issues in Entertainment Computing*, 49(1):52–70, 2005.
- [18] M. Fiedler, K. Tutschku, P. Carlsson, and A. Nilsson. Identification of performance degradation in IP networks using throughput statistics. In *Proc. ITC-18*. Elsevier, 2003.
- [19] M. Fiedler, L. Isaksson, S. Chevul, J. Karlsson, and P. Lindberg. Measurement and analysis of application-perceived throughput via mobile links. In *Proc. HET-NETs*, 2005.
- [20] D. V. Lindley. The theory of queues with a single server. *Proc. Cambridge Philos. Soc.*, 48:277–289, 1952.
- [21] G. Hasslinger and et al. Variability of broadband ADSL traffic: Measurement and time slotted modelling. In *Proc. Euro-NGI*, 2006.
- [22] S. Aalto. Output from an A-M-S type fluid queue. In *Proc. ITC*, pages 421–443. Elsevier, 1994.
- [23] M. Fiedler and K. Tutschku. Application of the stochastic fluid flow model for bottleneck identification and classification. In *Proc. DASD*, pages 35–42, 2003.
- [24] J.D.C. Little. A proof for the queuing formula: $L=\lambda W$. *Operations Research*, 9(3):383–387, 1961.
- [25] M. N. Ismail and A. M. Zin. Comparing the accuracy of end-to-end network performance measurement testbed and simulation model for data transfers in heterogeneous environment. In *Proc. ICMS*, pages 124–131. IEEE Computer Society, 2008.
- [26] C. Sarr and I. Guerin. Estimating average end-to-end delays in IEEE 802.11 multihop wireless networks. Technical report, CITI Insa Lyon / INRIA Grenoble Rhône-Alpes – ARES, 2007.
- [27] S. Borst. Flow-level performance and user mobility in wireless data networks. *Philosophical Trans. of the Royal Society: Mathematical, Physical and Engineering*, 366(1872):2047–2058, 2008.
- [28] M. M. Barry, A. Campbell, and A. Veres. Distributed control algorithms for service differentiation in wireless packet networks. In *Proc. INFOCOM*, 2001.
- [29] H. Graja and et al. A statistical estimation of average IP packet delay in cellular data networks. In *Proc. IEEE WCNC*, 2005.
- [30] M. Garetto and D. Towsley. Modeling, simulation and measurements of queuing delay under long-tail internet traffic. In *ACM SIGMETRICS Perform. Eval. Rev.*, volume 31, pages 47–57, 2003.
- [31] G. Aniba and S. Aissa. A general traffic and queuing delay model for 3G wireless packet networks. In *Telecommunications and Networking - ICT 2004 (LNCS)*, volume 3124, pages 125–137. Springer Berlin / Heidelberg, 2004.
- [32] B. Ngamwongwattana and R. Thompson. Measuring one-way delay of VoIP packets without clock synchronization. In *Proc. I2MTC*, 2009.
- [33] M. Cola, , and et al. Covert channel for one-way delay measurements. In *Proc. ICCCN*, pages 1–6, 2009.
- [34] J. Kogel. One-way delay measurement based on flow data: Quantification and compensation of errors by exporter profiling. In *Proc. ICOIN*, pages 25–30, 2011.
- [35] K. Wac, M. Fiedler, R. Bulst, and H. Hermens. Estimations of additional delays for mobile application data from comparative output-input throughput analysis. In *Proc. IEEE NOMS*, pages 171–178. IEEE, 2010.
- [36] N. Dokovsky, A. van Halteren, and I. Widya. BANip: Enabling remote healthcare monitoring with body area networks. In N. Guelfi, E. Astesiano, and G. Reggio, editors, *Proc. FIJ103*, volume LNCS 2952, pages 62–72. Springer Verlag, 2004.
- [37] SUN. The Jini™ technology surrogate architecture overview. Technical report, SUN Microsystems Inc., 2001.
- [38] R. Bulst, K. Wac, A. van Halteren, D. Konstantas, and V. Nicola. Goodput analysis of 3G wireless networks supporting m-health services. In *Proc. IEEE ConTEL*, 2005.
- [39] P. Arlos and M. Fiedler. Influence of the packet size on the one-way delay on the down-link in 3G networks. In *Proc. ISWPC'10*, pages 573–578, Piscataway, NJ, USA, 2010. IEEE Press.

Markus Fiedler has received his PhD in electrical engineering/ICT from Universität des Saarlandes, Saarbrücken, Germany, in 1998. Since then, he has been with Blekinge Institute of Technology, Karlskrona, Sweden. As Professor within the School of Computing (COM) and Head of the Communications and Computer Systems Research Lab (CCS) at BTH, he performs and supervises research on Quality of Experience; seamless communications; network virtualization; service chains; and networks of the future. He is leading and participating in several national and European projects. He is serving on the Steering Board of the European Network of Excellence Euro-NF and coordinating its research activities. Furthermore, he is co-chair of the Future Internet Cluster of the EC.

Katarzyna Wac is since 2010 leading the Quality of Life group at the Institute of Services Science of University of Geneva (Switzerland). She holds a BSc and MSc in Computer Science, MSc in Telematics, and PhD in Information Systems. In 2003–2004 she was a researcher at University of Twente, the Netherlands and in 2004–2009 at University of Geneva. In 2009–2010 she was on leave at Carnegie Mellon University (USA). She researches pervasive mobile computing and communication in mHealth.

Richard Bulst is the co-founder and CTO of MobiHealth B.V., a company specialized in development of (mobile) telemonitoring services for the healthcare market. He holds an MSc (cum laude) in Telematics and a BSc in Technical Computer Science. He has more than 20 years of experience in performance analysis of (mobile) data communication networks. He researches methods and techniques for auto-adaptation of mobile health applications behaviour based on performance measurements of data communication networks.

Patrik Arlos is Assistant Professor at the School of Computing, Blekinge Institute of Technology. He holds a MSc degree in Electrical Engineering from Hogskolan i Karlskrona/Ronneby, and in 2005 he defended his PhD in Telecommunications Systems at Blekinge Institute of Technology. His research interests lies in network measurements, network performance, Quality of Service and Experience.