

A Non-Stationary Service Curve Model for Estimation of Cellular Sleep Scheduling

Nico Becker[✉] and Markus Fidler[✉], *Senior Member, IEEE*

Abstract—While steady-state solutions of backlog and delay have been derived for wireless systems, the analysis of transient phases still poses significant challenges. Considering the majority of short-lived and interactive flows, transient startup effects, as caused by sleep scheduling in cellular networks, have, however, a substantial impact on the performance. To facilitate reasoning about the transient behavior of systems, this paper contributes a notion of non-stationary service curves. Models of systems with sleep scheduling are derived and transient backlogs and delays are analyzed. Further, measurement methods that estimate the service of an unknown system from observations of selected probe traffic are developed. Fundamental limitations of existing measurement methods are explained by the non-convexity of the transient service and further difficulties are shown to be due to the super-additivity of network service processes. A novel two-phase probing technique is devised that first determines the shape of a minimal probe and subsequently obtains an accurate estimate of the unknown service. In a comprehensive measurement campaign, the method is used to evaluate the service of cellular networks with sleep scheduling (2G, 3G, and 4G), revealing considerable transient backlog and delay overshoots that persist for long relaxation times.

Index Terms—Cellular networks, sleep scheduling, DRX, network calculus, non-stationary service curves, transient delay

1 INTRODUCTION

THE majority of flows in today's computer networks are short-lived [2] and hence dominated by various transient effects, such as TCP slow start, that can have a significant impact on their performance. Another case in point is energy saving in wireless networks to increase the mobiles' battery lifetime by deactivating power-intensive parts, including the display and the radio interface. In cellular radio, mobiles enter different types of sleep states and wake up according to a defined schedule to receive information in the downlink or in case an uplink transmission is requested. The feature is referred to as discontinuous reception (DRX) [3].

An example of DRX is illustrated in Fig. 1. Initially, the mobile is in idle state and wake up is triggered by a channel request at P_0 . Activating the channel takes T_0 units of time before data can be transmitted. After the data transfer, the mobile enters sleep state at P_1 and wakes up again according to a defined sleep cycle after T_1 . The sleep, respectively, wake-up duration T_i can be deterministic or random. Similarly, the service during data transfer can have a constant or a variable service rate, e.g., due to wireless transmission outages. Generally while the mobile is asleep, data are buffered for later transmission, causing a non-negligible transient backlog overshoot and corresponding delays.

The effects of DRX on power consumption and battery lifetime of mobiles have been evaluated, e.g., for 2G and 3G in [4], and for 4G in [5], [6]. The re-activation of a dormant

connection requires, however, time [5], [7], e.g., to exchange signalling messages. This gives rise to an important trade-off between energy and delay [8], that motivated significant work on the optimization of DRX for certain traffic types in 3G [9], [10], 4G [6], [11], [12], [13], [14], [15], [16], and 5G [17] networks. An established method is modeling the DRX states by a semi-Markov chain to analyze the stationary mean wake-up delay [10], [11], [12], [13], [14], [15], [17]. Stationary queueing delays under DRX are also derived using an $M | G | 1$ model with vacations [6], [9].

The derivation of transient performance measures, such as the progression of backlog and delay during connection re-activation, causes fundamental difficulties. As an example consider the basic $M | M | 1$ queue, where the stationary state distribution follows readily from a set of linear balance equations, e.g., [18]. The transient state distribution, on the other hand, is expressed by a set of differential equations for which mainly approximate or numerical solutions are known [19]. As a consequence, transient solutions of queueing systems are sparse [20], [21], [22] or tailored to specific problems like TCP congestion control [2].

A framework that does without an assumption of stationarity is the deterministic network calculus [23], [24], [25] that is a theory of time-invariant linear systems under a min-plus algebra [24, pp. xiv-xviii]. Examples of time-invariant linear systems include, e.g., links with a constant capacity or leaky-bucket traffic shapers. Non-linear, time-variant, and possibly non-stationary systems are replaced by time-invariant linear bounds that consider the worst-case behavior. Hence, transient phases are not excluded, however, only maximal backlogs and delays can be expressed, as we will show in Section 2. The stochastic network calculus [25], [26], [27], [28], [29], [30], [31], [32], [33], on the other hand, takes time-variant systems into account. The service of a system is specified by a service process [25], [29], that is a random process, or a statistical service curve [26], [27], [28], that is a non-random function.

• The authors are with the Institute of Communications Technology, Leibniz Universität Hannover, Hanover 30167, Germany.
E-mail: {nico.becker, markus.fidler}@ikt.uni-hannover.de.

Manuscript received 12 Aug. 2016; revised 24 Jan. 2018; accepted 19 Mar. 2018. Date of publication 27 Apr. 2018; date of current version 3 Dec. 2018.

(Corresponding author: Markus Fidler.)

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TMC.2018.2821133




Fig. 1. Sleep scheduling as implemented in discontinuous reception (DRX).

Typically the stochastic network calculus either assumes stationary processes or uses stationary bounds.

In this work, we use the notion of time-variant systems [34], [35] to model non-stationary service characteristics. Time-variant systems are described by bivariate instead of univariate processes that enable considering changes over time. We contribute a non-stationary service curve model and derive solutions for systems with sleep scheduling that provide insights into their transient performance. Measures of interest include the transient overshoot and the relaxation time until the steady-state is approached. Second, we examine methods for estimation of a system's service curve from measurements of probe traffic. We refine known measurement methods to enable the estimation of non-stationary service curves, where we discover specific limitations that are explained by the non-convexity and super-additivity of the service. We devise a novel minimal probing technique that estimates a non-stationary service curve within a defined measure of accuracy. Using this technique, we conduct a comprehensive measurement campaign to evaluate the transient uplink service characteristics of cellular 2G (EDGE), 3G (HSPA), and 4G (LTE) sleep scheduling.

In comparison to our earlier work [1], major advancements of this paper are: improved numerical accuracy by using a bounding technique based on Doob's martingale; a comprehensive analysis of the effects of transient versus stationary delays; and a detailed measurement study of transient effects caused by DRX in different cellular networks.

The remainder of this work is structured as follows. In Section 2, we define non-stationary service curves, show a method for construction, and derive models of systems with sleep scheduling. In Section 3, we investigate the measurement-based estimation of non-stationary service curves. We reveal difficulties that arise and devise a new minimal probing method. In Section 4, we use our minimal probing method to estimate the service of productive cellular networks with sleep scheduling. We discuss further related works in the respective sections. Section 5 presents brief conclusions. A summary of our main notation is given in Table 1.

2 SERVICE MODELS AND PERFORMANCE BOUNDS

In this section, we define regenerative service processes, where regeneration points mark the start of new transient phases, as the basic model of non-stationary systems and show a first application that evaluates the performance of deterministic sleep scheduling (Section 2.1). Considering systems with random service and random sleep scheduling, we derive non-stationary service curves as a stochastic characterization of the regenerative service processes (Section 2.2). While we restrict the exposition to non-stationary systems, we note that non-stationary traffic can be dealt with in the same way.

2.1 Regenerative Service Processes

We consider systems with cumulative arrivals $A(t)$, where $A(t)$ denotes the number of bits that arrive in the time

TABLE 1
Summary of Main Notations

Notation	Description
τ, t	time index
$A(t), A(\tau, t)$	cumulative arrivals in $(0, t]$ and $(\tau, t]$
$D(t), D(\tau, t)$	cumulative departures in $(0, t]$ and $(\tau, t]$
$S(\tau, t)$	time-variant service in $(\tau, t]$
$S(t)$	time-invariant service in $(\tau, \tau + t]$ for $\tau \geq 0$
$B(t)$	backlog at t
$W(t)$	delay at t
P	probability
M	moment generating function (MGF)
ε	probability, quantile
θ	free parameter of the MGF
$\mathcal{A}^\varepsilon(t)$	statistical arrival envelope
$\mathcal{S}^\varepsilon(\tau, t)$	non-stationary service curve
$\Delta(\tau, t)$	deviation from additivity
$\delta(t)$	burst function
T	latency
R	rate
sl, slr	stationary latency (rate)
tl, tlr	transient latency (rate)
br	burst response
mp	minimal probe
rs	rate scanning

interval $(0, t]$. By convention, there are no arrivals for $t \leq 0$ so that we generally consider $t \geq 0$. Clearly, $A(t)$ is a non-negative, non-decreasing function, and $A(0) = 0$. Shorthand notation $A(\tau, t) = A(t) - A(\tau)$ is used to denote the arrivals in $(\tau, t]$ where $t \geq \tau \geq 0$. Trivially, $A(t, t) = 0$ for all $t \geq 0$. Similarly, $D(t)$ denotes the cumulative departures from the system.

The service that is provided by the system is characterized by a time-variant service process $S(\tau, t)$ that specifies the service as a bivariate function of the interval $(\tau, t]$. It establishes the departure guarantee [25], [34], [35], [36]

$$D(t) \geq \inf_{\tau \in [0, t]} \{A(\tau) + S(\tau, t)\} =: A \otimes S(t). \quad (1)$$

By convention, $S(\tau, t)$ is non-negative and $S(t, t) = 0$ for all $t \geq 0$. The operator \otimes that is defined by Eq. (1) is known as convolution under a min-plus algebra¹ [24], [25]. We note that \otimes is associative but not commutative in general. For linear systems,² such as a work-conserving link with a time-variant capacity, Eq. (1) holds with equality [25, Ex. 5.2.4]. Further examples of Eq. (1) include scheduling with cross-traffic [29] and networks of systems where a network service process $S^{net}(\tau, t)$ is computed from the service processes $S^i(\tau, t)$ of the individual systems $i = 1 \dots n$ by recursive insertion of Eq. (1) as $S^{net}(\tau, t) = S^1 \otimes S^2 \otimes \dots \otimes S^n(\tau, t)$ [25]. The notion of a network service process characterizes a network by a single equivalent system, implying an immediate extension of results for single systems to

1. We use the infimum \inf , i.e., the largest lower bound of a set, instead of the minimum that may not be defined in general. Similarly, the supremum \sup denotes the smallest upper bound of a set.

2. Given pairs of arrivals and corresponding departures $A_i(t), D_i(t)$ with index i . A system is min-plus linear if any min-plus linear combination of arrivals $\min\{A_i(t), A_j(t)\} + c$ results in the corresponding linear combination of the departures $\min\{D_i(t), D_j(t)\} + c$ where c is an arbitrary constant.

networks. Given this quality, we will use the terms system and network interchangeably.

The service guarantee of Eq. (1) enables the derivation of performance bounds. An upper bound of the backlog $B(t) = A(t) - D(t)$ at time t follows by insertion of Eq. (1) as

$$B(t) \leq \sup_{\tau \in [0, t]} \{A(\tau, t) - S(\tau, t)\}. \quad (2)$$

The definition of $B(t)$ considers all data that are in the system including data in buffers as well as data in transmission [24]. The First-Come First-Serve (FCFS) delay at time t defined as $W(t) = \inf\{w \geq 0 : A(t) \leq D(t+w)\}$ can be derived as $W(t) \leq \inf\{w \geq 0 : \sup_{\tau \in [0, t]} \{A(\tau, t) - S(\tau, t+w)\} \leq 0\}$.

As opposed to the state-of-the-art, that assumes stationary service processes, we consider the service as a non-stationary, regenerative process with regeneration points $\mathbb{P} = \{P_0, P_1, P_2, \dots\}$, where $P_0 = 0$ and $P_i < P_{i+1}$ for all $i \geq 0$. We divide $S(\tau, t)$ into segments

$$S_i(\tau, t) = S(\tau + P_i, t + P_i), \quad (3)$$

for all $0 \leq \tau \leq t \leq P_{i+1} - P_i$ and $i \geq 0$, so that $S_i(\tau, t)$ is the service process between the i th and the $(i+1)$ th regeneration point. The defining characteristic of a regenerative process [18] is that the $S_i(\tau, t)$ are statistical replicas of each other, i.e.,

$$\mathbb{P}[S_i(\tau, t) \leq x] = \mathbb{P}[S_j(\tau, t) \leq x], \quad (4)$$

for all $i, j, x \geq 0$, and $0 \leq \tau \leq t \leq \min\{P_{i+1} - P_i, P_{j+1} - P_j\}$. Owing to Eq. (4), we omit the index i in the sequel. Also, we will not explicitly mention the constraint $t \leq P_{i+1} - P_i$ and assume that the next regeneration point is generally spaced sufficiently apart.

In the following, we present a first application to deterministic sleep scheduling, as introduced in Fig. 1. We define the respective time-variant service process (Section 2.1.1) and derive performance bounds thereof (Section 2.1.2). Further, we include a related time-invariant model of a link with vacations from the literature to show the general limitations of the time-invariant approach (Section 2.1.3). Finally, we include exact results for a special case to verify the tightness of the bounds (Section 2.1.4).

2.1.1 Deterministic Sleep Scheduling

We consider a system that if idle goes to sleep state according to a defined protocol, see Fig. 1. Wake up is scheduled deterministically T units of time after entering the sleep state. The transmission rate in sleep state is 0 and otherwise it is R . Taking advantage of the fact that each transition to sleep state is a regeneration point, we define the time-variant service process for $t \geq \tau \geq 0$ as

$$S^{tlr}(\tau, t) = R[t - \max\{\tau, T\}]_+, \quad (5)$$

where $[x]_+ = \max\{0, x\}$ is the non-negative part of x . Eq. (5) takes the form of a latency-rate function where the latency is transient, abbreviated by superscript tlr .

2.1.2 Performance Bounds

To evaluate the influence of the sleep cycle T on the performance of different types of traffic, we specify the traffic arrivals by a general statistical envelope function $\mathcal{A}^\varepsilon(t)$ that satisfies the sample path guarantee

$$\mathbb{P}[A(\tau, t) \leq \mathcal{A}^\varepsilon(t - \tau), \forall \tau \in [0, t]] \geq 1 - \varepsilon, \quad (6)$$

for all $t \geq 0$. Above, $\varepsilon \in (0, 1]$ is a probability of overflow. With Eq. (6), statistical performance bounds follow readily by substitution of $\mathcal{A}^\varepsilon(t - \tau)$ for $A(\tau, t)$, e.g., the backlog bound Eq. (2) yields

$$\mathbb{P}\left[B(t) \leq \sup_{\tau \in [0, t]} \{\mathcal{A}^\varepsilon(t - \tau) - S(\tau, t)\}\right] \geq 1 - \varepsilon. \quad (7)$$

In previous work [1], we used an established method [30] for construction of $\mathcal{A}^\varepsilon(t)$ that is based on Chernoff's theorem and the union bound. We improve this envelope by applying a recent technique [37], [38] that uses Doob's martingale inequality [39] instead. An envelope for processes with independent and identically distributed (iid) increments is

$$\mathcal{A}^\varepsilon(t) = \frac{1}{\theta} (\ln M_A(\theta, t) - \ln \varepsilon), \quad (8)$$

where $M_A(\theta, t) = \mathbb{E}[e^{\theta A(\tau, \tau+t)}] = (M_A(\theta, 1))^t$ for $\tau, t \geq 0$ is the moment generating function (MGF) of A and $\theta > 0$ is a free parameter.

Derivation of Eq. (8). We consider the parameterized envelope $\mathcal{A}^\varepsilon(t) = \rho_A t + \sigma_A$ with rate parameter $\rho_A > 0$ and burstiness parameter $\sigma_A \geq 0$. To obtain ρ_A and σ_A , we insert $\mathcal{A}^\varepsilon(t)$ into Eq. (6) and derive

$$\begin{aligned} & 1 - \mathbb{P}[A(\tau, t) \leq \rho_A(t - \tau) + \sigma_A, \forall \tau \in [0, t]] \\ &= \mathbb{P}\left[\max_{\tau \in [0, t-1]} \{A(\tau, t) - \rho_A(t - \tau)\} > \sigma_A\right] \\ &= \mathbb{P}\left[\max_{\tau \in [1, t]} \{e^{\theta(A(t-\tau, t) - \rho_A \tau)}\} > e^{\theta \sigma_A}\right], \end{aligned}$$

for $\theta > 0$, where we used that $A(t, t) = 0$. Next, we fix $t > 0$ and consider the process $U(\tau) = e^{\theta(A(t-\tau, t) - \rho_A \tau)}$, for $\tau \in [0, t]$. It follows that $U(\tau + 1) = U(\tau)e^{\theta(A(t-\tau-1, t-\tau) - \rho_A)}$ and using the independence the conditional expectation becomes

$$\begin{aligned} & \mathbb{E}[U(\tau + 1)|U(\tau), U(\tau - 1), \dots, U(1)] \\ &= U(\tau) \mathbb{E}[e^{\theta(A(t-\tau-1, t-\tau) - \rho_A)}]. \end{aligned}$$

Since for iid increments $\mathbb{E}[e^{\theta(A(t-\tau-1, t-\tau) - \rho_A)}] = M_A(\theta, 1)$, it follows that $\mathbb{E}[U(\tau + 1)|U(\tau), U(\tau - 1), \dots, U(1)] = U(\tau)$ is a martingale if $e^{\theta \rho_A} = M_A(\theta, 1)$. Hence, we obtain parameter $\rho_A = \ln M_A(\theta, 1)/\theta$ so that $\rho_A t = \ln M_A(\theta, t)/\theta$.

By application of Doob's martingale inequality [39, Theorem 3.2, p. 314] and the reformulation from [37] we have for non-negative martingales $U(\tau)$ for $\tau \geq 1$ that

$$x \mathbb{P}\left[\max_{\tau \in [1, t]} \{U(\tau) \geq x\}\right] \leq \mathbb{E}[U(1)].$$

With $\mathbb{E}[U(1)] = 1$ and $x = e^{\theta \sigma_A}$ it follows that

$$\mathbb{P}\left[\max_{\tau \in [0, t]} \{A(\tau, t) - \rho_A(\theta)(t - \tau)\} > \sigma_A\right] \leq e^{-\theta \sigma_A}.$$

Finally, we let $\varepsilon = e^{-\theta \sigma_A}$ to obtain $\sigma_A = -\ln \varepsilon/\theta$. \square

For a numerical study, we consider the special case of a discrete time equivalent of a stationary Poisson arrival process that enables an exact solution for comparison. In each

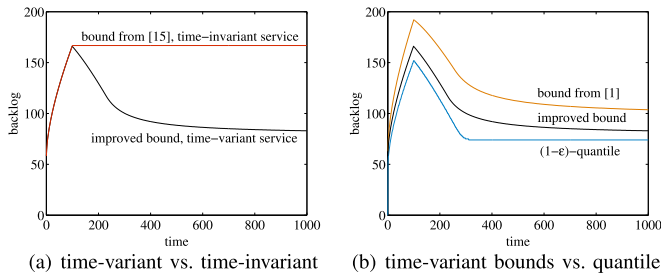


Fig. 2. Progression of the transient backlog over time. The time-variant service model correctly estimates the shape of the $(1 - \varepsilon)$ -quantile.

time-slot, the probability of a packet arrival is an independent Bernoulli trial with parameter $\alpha \in [0, 1]$ such that the number of packet arrivals $N(t)$ in an interval of t time-slots is binomial with the same parameter α . The individual packet sizes $Y(i)$ with index $i = 1, 2, \dots$ are iid geometric random variables with parameter $\beta \in (0, 1]$. Parameter α has the interpretation of an average arrival rate and $1/\beta$ is the average size of packets. For the special case of a system with constant service rate $R = 1$, α/β is the utilization and $\alpha < \beta$ is required for stability. The respective MGFs are $M_N(\vartheta, t) = (\alpha e^\vartheta + 1 - \alpha)^t$ and $M_Y(\theta) = \beta e^\theta / (1 - (1 - \beta)e^\theta)$ for $\theta \in [0, -\ln(1 - \beta)]$ [18]. The cumulative arrival process is the doubly stochastic process $A(\tau, t) = \sum_{i=N(\tau)+1}^{N(t)} Y(i)$. It has MGF $M_A(\theta, t - \tau) = M_N(\ln M_Y(\theta), t - \tau)$ [33] so that by insertion of $M_N(\vartheta, t)$ and $M_Y(\theta)$ we obtain

$$M_A(\theta, t) = \left(\frac{\alpha \beta e^\theta}{1 - (1 - \beta)e^\theta} + 1 - \alpha \right)^t. \quad (9)$$

Clearly, $M_A(\theta, t) = (M_A(\theta, 1))^t$ as $A(t)$ has iid increments.

In Fig. 2a, we illustrate the progression of the backlog bound Eq. (7) over time, labeled time-variant service. The curve shows an evident transient overshoot. The parameters of the service process given by Eq. (5) are $T = 100$ and $R = 1$. For the Poisson arrival process we use the MGF Eq. (9) with parameters $\alpha = 0.09$ and $\beta = 0.3$ corresponding to a utilization of 0.3. We choose $\varepsilon = 10^{-9}$ and optimize the free parameter $\theta > 0$ in Eq. (8) numerically. Since we normalized the rate $R = 1$, we do not provide units. If we define the duration of a time-slot as 1 ms, and the unit of data as 10 kbit, i.e., 1,250 Byte, the rate R equates to 10 Mbps.

2.1.3 Work-Conserving Link with Vacations [25]

A related model of a work-conserving link that takes deterministically bounded vacations whenever there is no data to transmit is proposed in [25, Th. 2.3.16]. The model is formulated in the deterministic network calculus where the service is expressed by the time-invariant, univariate function

$$S^{slr}(\tau, t) = R[t - \tau - T]_+ = S^{slr}(t - \tau), \quad (10)$$

that depends only on the width of the interval $(\tau, t]$ but not on its absolute position in time. Eq. (10) has the well-known type of a latency-rate function, however, the latency is stationary as it is not tied to any regeneration points as in Eq. (5).

Fig. 2a shows that both, the time-variant and the time-invariant service model, reveal the same growth of the backlog bound until service starts at $T = 100$. How the transient backlog is cleared after T and eventually converges to a stationary backlog bound is, however, only explained by the

time-variant model. To see why the time-invariant model cannot include this relaxation, note that the sup in Eq. (7) is generally non-decreasing in t if S is a univariate function.

2.1.4 Exact Markov Model

To evaluate the tightness of the performance bounds, we take advantage of the memorylessness of the arrivals and set up a Markov model that enables an exact (albeit numerical) transient solution. While an exact solution is feasible for the special case of the discrete-time Poisson process, we note that the envelope approach extends to a variety of non-trivial arrival processes including self-similar, long-range dependent [40], [41], and heavy-tailed processes [41]. Martingale bounds are also available for Markov modulated [38], and autoregressive processes [25], [38]. For details on the solution of the Markov model see [1].

For comparison with Fig. 2a, we depict in Fig. 2b the $(1 - \varepsilon)$ -quantile of $B(t)$ as derived from the Markov model as well as the previous bound [1, Eq. (7)]. We observe that the bounds that are derived from the time-variant service model provide good estimates that recover the shape of the quantile well. The deviations are due to inequalities that are invoked in the derivation of the envelope functions, where the use of Doob's martingale inequality shows a clear improvement.

2.2 Non-Stationary Service Curves

Based on the concept of time-variant, regenerative service process, we consider the service as a non-stationary random process and derive a basic stochastic characterization. We define a lower service envelope $\mathcal{S}^\varepsilon(\tau, t)$ that satisfies

$$\mathbb{P}[S(\tau, t) \geq \mathcal{S}^\varepsilon(\tau, t), \forall \tau \in [0, t]] \geq 1 - \varepsilon, \quad (11)$$

for all $t \geq 0$, where $\varepsilon \in (0, 1]$ is a probability of underflow. Compared to the state-of-the-art definition of univariate service and arrival envelope functions [26], such as Eq. (6), the essential difference is that Eq. (11) specifies a time-variant, bivariate envelope function. We emphasize that this generalization is crucial for modelling transient changes over time, as already shown for the deterministic case in Sections 2.1.1 and 2.1.3. It brings further basic differences about, e.g., the convolution is commutative for univariate but not for bivariate functions.

A system with service envelope $\mathcal{S}^\varepsilon(\tau, t)$ guarantees that

$$\mathbb{P}[D(t) \geq A \otimes \mathcal{S}^\varepsilon(t)] \geq 1 - \varepsilon, \quad (12)$$

for all $t \geq 0$. We refer to $\mathcal{S}^\varepsilon(\tau, t)$ as *non-stationary service curve*. To see that Eq. (12) follows from Eq. (11) for a system Eq. (1), add $A(\tau)$ to both sides of Eq. (11) to obtain

$$\mathbb{P}[A(\tau) + S(\tau, t) \geq A(\tau) + \mathcal{S}^\varepsilon(\tau, t), \forall \tau \in [0, t]] \geq 1 - \varepsilon. \quad (13)$$

Since Eq. (13) makes a sample path argument for all $\tau \in [0, t]$, it also holds that

$$\mathbb{P} \left[\inf_{\tau \in [0, t]} \{A(\tau) + S(\tau, t)\} \geq \inf_{\tau \in [0, t]} \{A(\tau) + \mathcal{S}^\varepsilon(\tau, t)\} \right] \geq 1 - \varepsilon,$$

where we finally substitute $D(t) = \inf_{\tau \in [0, t]} \{A(\tau) + S(\tau, t)\}$ using Eq. (1) to arrive at Eq. (12).

For a given service process $S(\tau, t)$, a non-stationary service curve $\mathcal{S}^\varepsilon(\tau, t)$ can be derived as

$$\mathcal{S}^\varepsilon(\tau, t) = -\frac{1}{\theta(\tau, t)} (\ln M_S(-\theta, \tau, t) + \rho(t - \tau) - \ln(\rho\varepsilon)), \quad (14)$$

where $M_S(-\theta, \tau, t) = \mathbf{E}[e^{-\theta S(\tau, t)}]$ is the negative MGF, respectively, Laplace transform and $\theta(\tau, t) > 0$ and $\rho \in (0, 1/\varepsilon]$ are free parameters.

Derivation of Eq. (14). The derivation of Eq. (14) extends [42] to non-stationary processes. It employs basic steps from the stochastic network calculus [27], [28]. We use the complementary formulation of Eq. (11) to define

$$\xi := \mathbf{P}[\exists \tau \in [0, t] : S(\tau, t) < \mathcal{S}^\varepsilon(\tau, t)] \leq \varepsilon.$$

We prove that $\mathcal{S}^\varepsilon(\tau, t)$ defined in Eq. (14) satisfies Eq. (11) by showing that $\xi \leq \varepsilon$. Using the union bound and Chernoff's lower bound $\mathbf{P}[X \leq x] \leq e^{\theta x} M_X(-\theta)$ for $\theta \geq 0$ it holds that

$$\xi \leq \sum_{\tau=0}^{t-1} \mathbf{P}[S(\tau, t) < \mathcal{S}^\varepsilon(\tau, t)] \leq \sum_{\tau=0}^{t-1} e^{\theta(\tau, t) \mathcal{S}^\varepsilon(\tau, t)} M_S(-\theta, \tau, t),$$

where $\theta(\tau, t) \geq 0$ is a set of free parameters. The case where $\tau = t$ is omitted since $S(t, t) = 0$ and $\mathcal{S}^\varepsilon(t, t) \leq 0$ by definition. By insertion of $\mathcal{S}^\varepsilon(\tau, t)$ from Eq. (14) we have

$$\xi \leq \rho\varepsilon \sum_{\tau=0}^{t-1} e^{-\rho(t-\tau)} = \rho\varepsilon \sum_{\nu=1}^t e^{-\rho\nu} \leq \rho\varepsilon \int_0^\infty e^{-\rho y} dy = \varepsilon,$$

where each summand is bounded by $e^{-\rho\nu} \leq \int_{\nu-1}^\nu e^{-\rho y} dy$ since $e^{-\rho y}$ is decreasing. Finally, letting $t \rightarrow \infty$ and solving the integral completes the proof that $\xi \leq \varepsilon$ for all $t \geq 0$. \square

2.2.1 Random Sleep Scheduling

The concept of non-stationary service curve enables the stochastic analysis of systems. We consider a work-conserving system with random sleep scheduling and random service increments $Z(t)$ for $t \geq 0$. When entering sleep state, the system regenerates and wakes up after a random time $T \geq 0$, i.e., $Z(t) = 0$ for $t \in [0, T]$. The service process is computed as $S(\tau, t) = \sum_{\nu=\tau+1}^t Z(\nu)$ for all $t > \tau \geq 0$ and $S(t, t) = 0$ for all $t \geq 0$. To derive the MGF of $S(\tau, t)$, we first consider the number of usable time-slots in $(\tau, t]$, i.e., after time T

$$U(\tau, t) = [t - \max\{\tau, T\}]_+.$$

The MGF of $U(\tau, t)$ is composed of three terms

$$\begin{aligned} M_U(\theta, \tau, t) &= e^{\theta(t-\tau)} \mathbf{P}[T \leq \tau] + \sum_{\nu=\tau+1}^t e^{\theta(t-\nu)} \mathbf{P}[T = \nu] + \mathbf{P}[T > t], \end{aligned} \quad (15)$$

that correspond to the cases where the start of the service T occurs before and including τ , within $(\tau, t]$, and after t , respectively. Given the service increments $Z(t)$ for $t > T$ are iid with MGF $M_Z(\theta)$, the MGF of the service process is

$$M_S(\theta, \tau, t) = \mathbf{E}[(M_Z(\theta))^{U(\tau, t)}] = M_U(\ln M_Z(\theta), \tau, t). \quad (16)$$

For an implementation, we use memoryless processes, as solutions for this specific case may also be derived, e.g., from a Markov model. This enables us to compute certain

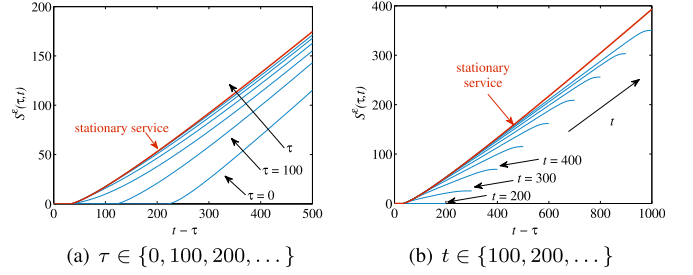


Fig. 3. Non-stationary service curves of random sleep scheduling.

reference results in Section 3. We note that the service curve in Eq. (14) is, however, not limited to memoryless processes. In detail, we model T as a geometric random variable with parameter p , where $\mathbf{P}[T = \nu] = p(1-p)^\nu$, and $Z(t)$ for $t > T$ following a basic wireless outage model [33] as iid Bernoulli trials with parameter q .

Regarding Eq. (15), we have $\mathbf{P}[T \leq \tau] = 1 - (1-p)^{\tau+1}$ and

$$\sum_{\nu=\tau+1}^t e^{\theta(t-\nu)} \mathbf{P}[T = \nu] = e^{\theta t} p \sum_{\nu=\tau+1}^t (e^{-\theta}(1-p))^\nu,$$

where we substitute $y = e^{-\theta}(1-p)$ and compute

$$\sum_{\nu=\tau+1}^t y^\nu = \frac{y^{\tau+1} - y^{t+1}}{1-y}.$$

Having obtained a solution of Eq. (15), the MGF of the service process follows from Eq. (16) by insertion of $M_Z(\theta) = qe^\theta + 1 - q$ for the Bernoulli service increment process. Finally, the non-stationary service curve is computed from Eq. (14).

Fig. 3 illustrates how $\mathcal{S}^\varepsilon(\tau, t)$ identifies the characteristic features of sleep scheduling. The parameters are $p = 0.1$ and $q = 0.5$, i.e., the mean transient latency is $\mathbf{E}[T] = (1-p)/p = 9$ and the mean stationary service rate is $\mathbf{E}[Z] = q = 0.5$. The service curves are computed for $\varepsilon = 10^{-6}$, $\rho = 10^{-4}$, and $\theta(\tau, t)$ is optimized numerically.

In Fig. 3a, we show how the service in an interval of width $t - \tau$ increases with increasing distance τ from the last regeneration point. For small τ a significant impact of the initial transient phase is noticed. For large τ we observe that $\mathcal{S}^\varepsilon(\tau, t)$ converges towards a stationary service curve that is computed for the same service increment process, however, without sleep scheduling. Fig. 3b displays the service curves for fixed t and variable τ . We emphasize that the transient phase is reflected in the non-convex shape of the curves, rightwards where τ approaches zero. In contrast, the initial delay at the origin, that also applies to the stationary service curve, is caused by the outages of the Bernoulli service increment process. For small intervals $t - \tau$ the process results in a service of zero with non-negligible probability. For $t = 100$ and $t = 200$ the effects bring about a service curve of zero.

2.2.2 Performance Bounds

The presentation of $\mathcal{S}^\varepsilon(\tau, t)$ in Fig. 3b conforms with the formulation of statistical performance bounds, where t is fixed and all $\tau \in [0, t]$ are evaluated. A statistical backlog bound follows with Eqs. (12) and (6) as

$$\mathbf{P}[B(t) \leq \sup_{\tau \in [0, t]} \{\mathcal{A}^\varepsilon(t - \tau) - \mathcal{S}^\varepsilon(\tau, t)\}] \geq 1 - 2\varepsilon.$$

A delay bound follows in the same way. Intuitively, the backlog and delay bound are the maximal vertical and horizontal deviation of $\mathcal{A}^\varepsilon(t - \tau)$ and $\mathcal{S}^\varepsilon(\tau, t)$, respectively. We present numerical results in [1].

3 MEASUREMENT-BASED ESTIMATION METHODS

Following the model-based approach in Section 2, we derived service curves of sleep scheduling. To identify a system's service curve without knowledge of its internals, we investigate measurement-based methods. We use these to evaluate full implementations of DRX in cellular networks, in Section 4.

The estimation of a system's service curve from measurements has received increasing attention in the past years [43], [44], [45], [46], [47], [48], [49]. The general approach can be viewed as system identification of a black box model, where the service curve of an unknown system is estimated from observations of its arrivals and departures. Usually, the black box model presumes only basic properties such as linearity and stationarity. The existing works can be classified accordingly to be based either on the assumption of a work-conserving system [43], [44], [45], [46] or a general min-plus linear system [47], [48], [49] and either use a deterministic model of a time-invariant system [45], [46], [47], [48] or a stochastic model of a stationary system [43], [44], [49]. For a more detailed comparison see also [31]. Related methods focus on estimating the long-term average capacity [50] or available bandwidth [51] of cellular networks.

While methods, such as rate scanning [49], are available for estimation of stationary service curves for the general class of min-plus linear systems with random service [49], the measurement-based estimation of non-stationary service curves that are the subject of this work is not elaborated in the current literature. Like [47], [48], [49], we consider min-plus linear systems, where Eq. (1) holds with equality, i.e.,

$$D(t) = \inf_{\tau \in [0, t]} \{A(\tau) + S(\tau, t)\}, \quad (17)$$

for all $t \geq 0$. The unknown $S(\tau, t)$ is a time-variant, random service process as in [49]. Compared to [49] we do, however, not assume stationarity of $S(\tau, t)$.

Based on Eq. (17), the goal of service curve estimation can be phrased as an inversion problem, i.e., given measurements of $A(t)$ and $D(t)$, solve Eq. (17) for $S(\tau, t)$. Due to the infimum, the min-plus convolution has, however, no inverse operation in general and a solution can only be obtained for certain functions $A(t)$ [48]. Finally, we seek to estimate a maximal non-stationary service curve $\mathcal{S}^\varepsilon(\tau, t)$, i.e., find a Pareto efficient function $\mathcal{S}^\varepsilon(\tau, t)$ that satisfies Eq. (12).

We perform measurements assuming a regenerative service process as defined in Section 2.1, where repeated network probes can observe samples of the service process $S_i(\tau, t)$ at regeneration point P_i as illustrated in Fig. 1.

The remainder of this section is structured as follows: We adapt two known methods, rate scanning (Section 3.1) and burst response (Section 3.2), for estimation of non-stationary service curves. Fundamental limitations of these methods are identified that are explained by the non-convexity and the super-additivity of the service. We devise a new two-phase probing method (Section 3.3) that first uses the burst response to determine the shape of a suitable probe, that is proven to be minimal under the conditions of Lemma 3.

The probe is then used to obtain a service curve estimate with a defined accuracy.

3.1 Rate Scanning

First, we consider the rate scanning method from [48], [49]. The method uses constant rate probes $A(t) = rt$ for a set of rates $r \in \mathbb{R}$ to excite the network for measurements. While [48] and [49] consider deterministic and stationary service curves, respectively, we adapt the method to provide estimates of non-stationary service curves for transient phases.

For constant rate probes $A(t) = rt$ the backlog of a system follows from Eq. (17) as $B(t) = \sup_{\tau \in [0, t]} \{r(t - \tau) - S(\tau, t)\}$. Consequently, it holds that $B(t) \geq r(t - \tau) - S(\tau, t)$, for all $\tau \in [0, t]$. Solving for $S(\tau, t)$ provides the lower bound

$$S(\tau, t) \geq r(t - \tau) - B(t), \quad (18)$$

for all $\tau \in [0, t]$.

Next, we use the definition of $(1 - \xi)$ -quantile

$$X^\xi = \inf\{x \geq 0 : \mathbb{P}\{X \leq x\} \geq 1 - \xi\}, \quad (19)$$

and denote $B^\xi(r, t)$ the backlog quantile at time t as a function of r . By insertion of $B^\xi(r, t)$ into Eq. (18), it holds that

$$\mathbb{P}\{S(\tau, t) \geq r(t - \tau) - B^\xi(r, t), \forall \tau \in [0, t]\} \geq 1 - \xi,$$

which takes the form of Eq. (11). By application of the union bound, it follows for $t \geq \tau \geq 0$ that

$$\mathcal{S}_{rs}^\varepsilon(\tau, t) = \max_{r \in \mathbb{R}} \{r(t - \tau) - B^\xi(r, t)\}, \quad (20)$$

is a non-stationary service curve as defined by Eq. (12) with probability $\varepsilon = \sum_{r \in \mathbb{R}} \xi$. Compared to [49], Eq. (20) uses the transient instead of the stationary backlog to estimate a non-stationary service curve.

For evaluation of the method, we consider a simulation of the random sleep scheduling model from Section 2.2.1 with identical parameters $p = 0.1$ and $q = 0.5$. We perform rate scanning with ten rates $r \in \{0.05, 0.1, \dots, 0.5\}$. For each rate we obtain 10^5 backlog samples from repeated experiments. From these, we extract an estimate of the backlog quantile $B^\xi(r, t)$ for $\xi = 10^{-4}$ so that $\varepsilon = \sum_{r \in \mathbb{R}} \xi = 10^{-3}$.

Fig. 4a shows the linear segments gathered by each of the ten probing rates $r \in \mathbb{R}$ and the resulting estimate $\mathcal{S}_{rs}^\varepsilon(\tau, t)$ that is obtained as the maximum of the linear segments by Eq. (20). We display bivariate functions for $t = 200$ and vary $\tau \in [0, t]$ to consider the influence of the width of the interval $(\tau, t]$. As a reference, we include an analytical upper bound.³ Any function that exceeds the upper bound for some $\tau \in [0, t]$ violates the definition of service envelope Eq. (11). Further, we show an analytical service curve⁴ that provides the lower guarantee specified by Eq. (12). Both, the analytical upper bound as well as the analytical service curve exhibit a flat area in the upper right for large $t - \tau$, respectively, small τ that is due to the initial transient phase starting at $\tau = 0$. In contrast, the service curve estimate of rate

3. For a Bernoulli service increment process with parameter q that starts after a geometrically distributed time T , an upper bound of the service provided in $(\tau, t]$ is derived as the $(1 - \varepsilon)$ -quantile of a binomial distribution with parameters $t - \max\{\tau, T\}$ and q . We note that the first parameter, i.e., the number of trials $t - \max\{\tau, T\}$, is a random variable.

4. The analytical service curve again applies the binomial distribution as in case of the upper bound but in addition makes a sample path argument using the union bound to satisfy Eq. (11).

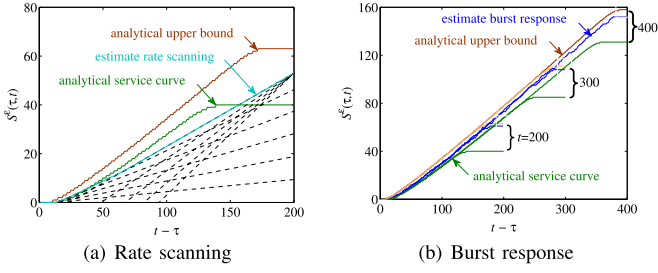


Fig. 4. Service curve estimates of random sleep scheduling. The estimate of rate scanning is the maximum of linear rate segments (dashed lines). By construction, it can only recover a convex hull. Burst probing can estimate non-convex service curves and performs close to the analytical upper bound.

scanning cannot recover the non-convex part of the analytical results.

This shortcoming of rate scanning is explained by the properties of Eq. (20) that has the form of a Legendre-Fenchel transform of the backlog [48]. The transform is also known as convex conjugate as the result is generally a convex function [52]. For convex functions, the Legendre-Fenchel transform is its own inverse, whereas the bi-conjugate of a non-convex function can only recover the convex hull [52], as seen for the estimate of rate scanning in Fig. 4a.

3.2 Burst Response

Motivated by the min-plus systems theory of the network calculus [24], [25], a canonical probe for system identification is the burst function that takes the role of the Dirac delta function in min-plus algebra. The burst function is defined as

$$\delta(t) = \begin{cases} 0 & \text{for } t = 0, \\ \infty & \text{for } t > 0. \end{cases} \quad (21)$$

The burst function is the neutral element of min-plus convolution, so that sending a burst probe $A(t) = \delta(t)$ reveals the service $S(0, t)$ for all $t \geq 0$ as the burst response of the system

$$D(t) = \inf_{\tau \in [0, t]} \{\delta(\tau) + S(\tau, t)\} = S(0, t). \quad (22)$$

For additive service processes as defined in [30, p. 6], $S(\tau, t)$ can be readily obtained from Eq. (22) for all $t \geq \tau \geq 0$ as

$$S(\tau, t) = S(0, t) - S(0, \tau). \quad (23)$$

For a stochastic analysis, we denote Ω the set of all causally possible sample paths. The individual sample paths are denoted $D_\omega(t)$ for $\omega \in \Omega$. For each $\omega \in \Omega$ we use the additivity as specified by Eq. (23) to obtain the service process from the burst response given by Eq. (22) as

$$S_\omega(\tau, t) = D_\omega(t) - D_\omega(\tau), \quad (24)$$

for all $\tau \in [0, t]$. We fix $t > 0$ and select a subset of the sample paths $\Psi_t \subseteq \Omega$ with probability $\mathbb{P}[\Psi_t] \geq 1 - \varepsilon$. Defining

$$S_{br}^e(\tau, t) = \inf_{\psi \in \Psi_t} \{S_\psi(\tau, t)\}, \quad (25)$$

it holds that $S_\psi(\tau, t) \geq S_{br}^e(\tau, t)$ for all $\tau \in [0, t]$ and all $\psi \in \Psi_t$. Further, we have $\mathbb{P}[\Psi_t] \geq 1 - \varepsilon$ so that $S_{br}^e(\tau, t)$ satisfies Eq. (11). Hence, $S_{br}^e(\tau, t)$ is a non-stationary service curve that conforms to Eq. (12).

In a practical probing scheme, we can only observe a finite set of sample paths Ω from repeated measurements of

the departures $D_\omega(t)$ for $t \geq 0$. We fix $t > 0$ and estimate the service from Eq. (24) for all $\tau \in [0, t]$ and $\omega \in \Omega$. To construct Ψ_t we remove a set of minimal sample paths from Ω . We define the minimal sample paths Φ to be the sample paths that attain the minimum

$$S_{\min}(\tau, t) = \min_{\omega \in \Omega} \{S_\omega(\tau, t)\},$$

for $\tau \in [0, t]$ most frequently. Formally, that is $\Phi = \arg \max_{\omega \in \Omega} \{X_\omega\}$, where for all $\omega \in \Omega$

$$X_\omega = \sum_{\tau=0}^{t-1} 1_{S_\omega(\tau, t) = S_{\min}(\tau, t)}.$$

The indicator function $1_{(\cdot)}$ is one if the argument is true and zero otherwise. We remove Φ to obtain $\Psi_t = \Omega \setminus \Phi$ and repeat the above steps for Ψ_t as long as $\mathbb{P}[\Psi_t] \geq 1 - \varepsilon$. Finally, we estimate the service curve from Eq. (25) for all $\tau \in [0, t]$.

Fig. 4b shows service curve estimates of burst probing for $t = 200, 300$, and 400 together with analytical upper bounds and analytical reference service curves for the same system as Fig. 4a. We observe that burst probing performs very well as it provides service curve estimates that are in between the analytical reference curves and closely track the analytical upper bound. Unlike rate scanning, burst probing recovers the flat, non-convex areas in the upper right of the analytical curves that are caused by the initial transient phase.

Aside from its good performance, burst probing has, however, downsides. Burst probes are considered intrusive as they cause non-linear behavior of certain systems. An example are FCFS multiplexers, where a burst probe can preempt other traffic resulting in a too optimistic service estimate [48].

Furthermore, we discover that the intuitive assumption of additive service processes [30], that is used as a basis of the method in Eq. (23), is not justifiable in general. While sub-additivity is less an issue, as (23) provides a conservative estimate of $S(\tau, t)$ that yields a valid lower service curve in this case, we find that there are relevant systems, where super-additivity applies and (23) overestimates $S(\tau, t)$.

Formally, a function $f(s, t)$ is super-additive if $f(s, u) \geq f(s, t) + f(t, u)$ for all $u \geq t \geq s \geq 0$. Trivially, an additive function $f(s, u) = f(s, t) + f(t, u)$ for all $u \geq t \geq s \geq 0$ is also super-additive. To quantify the magnitude of super-additivity, we define the maximal deviation of a super-additive function $f(s, t)$ from additivity as

$$\Delta(s, u) := f(s, u) - \inf_{t \in [s, u]} \{f(s, t) + f(t, u)\}. \quad (26)$$

With respect to Eq. (23), super-additivity implies that estimating $S(\tau, t)$ as $S(\tau, t) = S(0, t) - S(0, \tau)$ is too optimistic as in general only $S(\tau, t) \leq S(0, t) - S(0, \tau)$ holds.

3.2.1 Super-Additive Service Processes

To provide insights on the occurrence and the impact of non-additivity, we first consider the instructive example of deterministic latency-rate functions as specified by $S^{tlr}(\tau, t)$ in Eq. (5) for the case of a transient latency and $S^{slr}(\tau, t)$ in Eq. (10) for a stationary latency, respectively. In case of a transient latency, it can be verified that $S^{tlr}(\tau, t)$ is additive. For a stationary latency, however, we find that $S^{slr}(\tau, t)$ is super-additive but not generally additive. An example is




Fig. 5. Service curve estimates of deterministic sleep scheduling. Latency-rate service curves with a transient latency, with a stationary latency, and with both are compared. The transient latency equals 20 and the stationary latency 10. In case of the transient latency solely, the latency-rate service curve is additive and burst probing recovers the exact result. In contrast, in case of a stationary latency, the service curve is super-additive and burst probing overestimates the service curve. Minimal probing (see Section 3.3) provides a corresponding lower estimate that matches the service curve exactly in case of additivity.

$S^{slr}(0, \tau) + S^{slr}(\tau, t) = R(t - 2T)$ for $t \geq \tau + T$ and $\tau \geq T$, whereas $S^{slr}(0, t) = R(t - T)$. The example also achieves the maximal deviation from additivity $\Delta(\tau, t) = RT$ for $t \geq \tau + 2T$.

For illustration, Fig. 5 shows different types of latency-rate functions, labelled analytical service curve: with a transient latency of $T = 20$ according to Eq. (5) in Fig. 5a, as it applies for deterministic sleep scheduling; a stationary latency of $T = 10$ according to Eq. (10) in Fig. 5b; and for a combined system with both a transient and a stationary latency in Fig. 5c. The rate is $R = 1$. In all cases, we fix $t = 100$ and show the service curves for an increasing interval $t - \tau$.

As can be seen from the flat end of the curve in the region marked with tl in Fig. 5a, the transient latency has an effect only for large intervals $t - \tau > 80$, respectively, $\tau < 20$. The stationary latency, on the other hand, applies for all τ resulting in the right shift of the curve in Fig. 5b marked with sl . Finally, Fig. 5c shows both effects.

We also present estimates of the service curves as obtained from the burst response. Clearly, the burst response recovers the transient latency-rate function in Fig. 5a exactly. It overestimates, however, the stationary latency-rate function in Fig. 5b and similarly in Fig. 5c. The effect is due to the super-additivity induced by the stationary latency. In this case, the assumption of additivity used by Eq. (23) falsely assigns the stationary latency to large intervals $t - \tau$ only, as seen in the region marked with Δ . The figures further include estimates obtained by a method that we introduce as minimal probing in Section 3.3, where we also provide details on the curves.

While we designed the numerical experiments in Fig. 5 to investigate stationary and transient latencies in isolation, we note that in real networks stationary latencies, such as propagation delays, can not be avoided. In our experiments in cellular networks in Section 4, we generally have to deal with both types of latencies. As a consequence burst probing is not suitable as it overestimates the service.

3.2.2 Super-Additivity of \min and \otimes

A second notable exception of additive service processes are networks of systems. We consider n systems with service processes $S^i(\tau, t)$ for $i = 1, 2, \dots, n$ in tandem. The end-to-end service process of the network is denoted by $S^{net}(\tau, t)$. It is derived as the min-plus convolution of the service processes of the individual systems $S^i(\tau, t)$ for $i = 1, 2, \dots, n$ [25]. For additive S^i , the following Lemma 1 proves that S^{net} is super-additive in general and additive only in certain special cases. The result is significant as it refutes the assumption of

additive service processes for the large class of tandem systems. As before, it implies that the burst probing method is too optimistic. Our numerical evaluation in [1] shows significant deviations from additivity for tandem systems.

The following lemmas formalize the results and show that the \otimes and \min operators in general maintain only super-additivity but not additivity.

Lemma 1 (Super-additivity of \otimes). *Given two bivariate functions $f(s, t)$ and $g(s, t)$ for $t \geq s \geq 0$ where $f(t, t), g(t, t) = 0$ for all $t \geq 0$. Define $h(s, t) = f \otimes g(s, t)$.*

- 1) *If f and g are super-additive, then h is super-additive.*
- 2) *If f and g are additive and univariate, then h is additive.*

Proof of Lemma 1. By definition of h we have

$$\begin{aligned} h(s, t) + h(t, u) &= f \otimes g(s, t) + f \otimes g(t, u) \\ &= \inf_{\tau \in [s, t]} \inf_{v \in [t, u]} \{f(s, \tau) + f(t, v) + g(\tau, t) + g(v, u)\}. \end{aligned} \quad (27)$$

i. Given f and g are super-additive. From Eq. (27) we have

$$\begin{aligned} h(s, t) + h(t, u) &\leq \inf_{\tau \in [s, t]} \inf_{v \in [t, u]} \{f(s, v) - f(\tau, t) + g(\tau, t) + g(v, u)\} \\ &= \inf_{v \in [t, u]} \{f(s, v) + g(v, u)\} + \inf_{\tau \in [s, t]} \{g(\tau, t) - f(\tau, t)\} \\ &\leq \inf_{v \in [t, u]} \{f(s, v) + g(v, u)\}. \end{aligned} \quad (28)$$

In the first line, we estimated $f(s, \tau) + f(t, t) + f(t, v) \leq f(s, v)$ due to the super-additivity of f . In the second line, we rearranged the infima, and in the third line, we estimated $\inf_{\tau \in [s, t]} \{g(\tau, t) - f(\tau, t)\} \leq g(t, t) - f(t, t) = 0$ since $f(t, t), g(t, t) = 0$ for all $t \geq 0$. Similarly, using the super-additivity of g we derive from Eq. (27) that

$$h(s, t) + h(t, u) \leq \inf_{\tau \in [s, t]} \{f(s, \tau) + g(\tau, u)\}. \quad (29)$$

Combining Eqs. (28) and (29) we obtain

$$h(s, t) + h(t, u) \leq \inf_{\tau \in [s, u]} \{f(s, \tau) + g(\tau, u)\} = h(s, u),$$

which proves the super-additivity of h .

ii. For the special case of additive univariate functions $f(s, t) = f(t - s)$ and $g(s, t) = g(t - s)$, that depend only on the difference $t - s$ and not on the absolute values of s and t , it follows that $h(s, t) = f \otimes g(t - s) = h(t - s)$ is

also univariate. Using the additivity of f and g , Eq. (27) yields that

$$\begin{aligned} & h(t-s) + h(u-t) \\ &= \inf_{\tau \in [s,t]} \inf_{\nu \in [t,u]} \{f(\tau-s+\nu-t) + g(t-\tau+u-\nu)\} \\ &= \inf_{\zeta \in [s+t, t+u]} \{f(\zeta-s-t) + g(t+u-\zeta)\} \\ &= \inf_{\zeta \in [0, u-s]} \{f(\zeta) + g(u-s-\zeta)\} = h(u-s), \end{aligned}$$

which proves the additivity of h . \square

From Lemma 1 the convolution of two additive bivariate functions $f(s, t)$ and $g(s, t)$ for $t \geq s \geq 0$ is super-additive, however, in general it is not additive. As a counterexample consider the additive functions $f(s, t) = t - s$ and $g(s, t) = 2(\lfloor t/2 \rfloor - \lfloor s/2 \rfloor)$. It follows that $h(s, t) = f \otimes g(s, t) = 2\lfloor t/2 \rfloor - s$ is not additive. Additivity is, however, achieved in case of univariate functions, where Lemma 1 ii. extends a known result for min-plus convolution of sub-additive univariate functions in [24, p. 142].

Lemma 2 (Super-additivity of min). *Given two super-additive bivariate functions $f(s, t)$ and $g(s, t)$ for $t \geq s \geq 0$. The minimum $h(s, t) = \min\{f(s, t), g(s, t)\}$ is super-additive.*

Proof of Lemma 2. By definition of h we have

$$\begin{aligned} & h(s, t) + h(t, u) \\ &= \min\{f(s, t), g(s, t)\} + \min\{f(t, u), g(t, u)\} \\ &\leq \min\{f(s, u), g(s, u), f(s, t) + g(t, u), g(s, t) + f(t, u)\} \\ &\leq h(s, u). \end{aligned}$$

In the second line we used the super-additivity of f and g and in the third line $\min\{f(s, u), g(s, u)\} = h(s, u)$ and $\min\{h(s, u), x\} \leq h(s, u)$ for any x . \square

The minimum of two additive bivariate functions $f(s, t)$ and $g(s, t)$ for $t \geq s \geq 0$ is super-additive, as a special case of Lemma 2, but in general not additive. A counterexample is $f(s, t) = t - s$ and $g(s, t) = 2(\lfloor t/2 \rfloor - \lfloor s/2 \rfloor)$. Clearly f and g are additive, however, $h = \min\{f, g\} = 2\lfloor t/2 \rfloor - s$ is not.

3.3 Minimal Probing

Given the limitations identified above, we devise a new probing method to characterize transient service processes that are neither convex nor additive. The method comprises two phases. In the first step, a minimal probe as well as an upper bound of the service are estimated using the burst response as in Section 3.2. In the second step, the minimal probe is used to estimate a non-stationary service curve with a defined accuracy. We show that the minimal probe reveals the service of the system, whereas any smaller or larger probe estimates only a lower bound. The importance of the probe traffic intensity is also discussed in [48].

3.3.1 Estimation Using Arbitrary Probes

First, we consider how to obtain a service curve estimate from an arbitrary probe $A(\tau)$. The estimate will then be used to derive conditions for the shape of a minimal probe.

From Eq. (17) it follows that $D(t) \leq A(\tau) + S(\tau, t)$ for all $\tau \in [0, t]$ so that

$$S(\tau, t) \geq D(t) - A(\tau), \quad (30)$$

for all $\tau \in [0, t]$. An equivalent expression using the backlog is $S(\tau, t) \geq A(\tau, t) - B(t)$ for all $\tau \in [0, t]$. By insertion of the backlog quantile it follows that

$$S^\varepsilon(\tau, t) = A(\tau, t) - B^\varepsilon(t), \quad (31)$$

satisfies Eq. (11), i.e., $S^\varepsilon(\tau, t)$ is a non-stationary service curve as defined by Eq. (12). To estimate $S^\varepsilon(\tau, t)$, the probing procedure uses a defined probe $A(\tau, t)$ to obtain $B^\varepsilon(t)$ from repeated backlog measurements. In the special case $A(\tau, t) = r(t - \tau)$ the rate scanning method can be recovered.

3.3.2 Definition of Minimal Probe

So far, we did not constrain the shape of the probe and defining a suitable probe is non-trivial. Intuitively, a probe that is too small will provide little information about the service as the observed departures are mostly limited by the arrivals. A too large probe, on the other hand, will deteriorate the estimate, e.g., in the extreme case of a burst probe $A(\tau) = \delta(\tau)$, where δ is defined in Eq. (21), the lower bound in Eq. (30) will only be useful for $\tau = 0$ but not for $\tau > 0$. The same restriction applied to burst probing earlier. The following lemma defines a necessary and sufficient condition for a minimal probe to obtain the true service from Eq. (30).

Lemma 3 (Minimal Probe). *Fix $t > 0$ and define the probe*

$$A_{mp}(\tau) = S(0, t) - S(\tau, t), \quad (32)$$

for $\tau \in [0, t]$. Eq. (30) holds with equality if and only if $A(\tau) = A_{mp}(\tau)$ for all $\tau \in [0, t]$.

Proof of Lemma 3. By substitution of Eqs. (32) into (17) we have $D(t) = S(0, t)$ and finally by insertion into Eq. (30) $D(t) - A_{mp}(\tau) = S(\tau, t)$ for all $\tau \in [0, t]$, which proves that Eq. (32) is sufficient.

To see that Eq. (32) is necessary, consider any other (smaller or larger) probe $A(\tau)$ expressed as $A(\tau) = A_{mp}(\tau) + f(\tau)$ where $f(\tau)$ is a real function with $f(0) = 0$ and $\exists \nu \in (0, t] : f(\nu) \neq 0$. It follows by the same steps that $D(t) - A(\tau) = S(\tau, t) + \inf_{\nu \in [0, t]} \{f(\nu)\} - f(\tau)$ which implies that $\exists \tau \in [0, t] : D(t) - A(\tau) < S(\tau, t)$. \square

While Lemma 3 proves the optimality of $A_{mp}(\tau)$, it depends on the unknown service and cannot be constructed a priori. To gather information on $A_{mp}(\tau)$, we initially consider the system's burst response $D(\tau) = \delta \otimes S(\tau) = S(0, \tau)$, where δ is defined in Eq. (21), and estimate Eq. (32) by $\tilde{A}_{mp}(\tau) = S(0, \tau)$ for $\tau \in [0, t]$. We use tilde to indicate that $\tilde{A}_{mp}(\tau)$ is an approximation of $A_{mp}(\tau)$ that is exact only if $S(\tau, t)$ is additive. In the stochastic case, we employ $S_{br}^\varepsilon(\tau, t)$ from Eq. (25) to estimate

$$\tilde{A}_{mp}(\tau) = S_{br}^\varepsilon(0, t) - S_{br}^\varepsilon(\tau, t), \quad (33)$$

for $\tau \in [0, t]$. We note that additivity of $S_{br}^\varepsilon(\tau, t)$ cannot be assumed by construction of Eq. (25), see Lemma 2.

3.3.3 Accuracy of the Estimates

We show that the backlog at the end of the probe $\tilde{A}_{mp}(\tau)$ is a measure of the estimation accuracy. Further, we verify that

the backlog is bounded by the deviation of a super-additive service process from additivity.

First we investigate a time-variant, deterministic system, where Eq. (30) provides a lower estimate of the service. The estimate of the minimal probe $\tilde{A}_{mp}(\tau) = S(0, \tau)$ is obtained as the system's burst response from Eq. (22). By insertion into Eq. (30), a lower service estimate is

$$S(\tau, t) \geq \inf_{v \in [0, t]} \{S(0, v) + S(v, t)\} - S(0, \tau),$$

where we used Eq. (17) to compute $D(t)$ for $\tilde{A}_{mp}(\tau)$. On the other hand, the burst response provides the upper estimate $S(\tau, t) \leq S(0, t) - S(0, \tau)$ if $S(\tau, t)$ is super-additive. Taking the difference of the upper and the lower estimate, the service process is bounded in an interval of width

$$\Delta(0, t) = S(0, t) - \inf_{v \in [0, t]} \{S(0, v) + S(v, t)\},$$

that is the maximum deviation of $S(0, t)$ from additivity as defined in Eq. (26). Further, since $S(0, t) = \tilde{A}_{mp}(t)$ we have

$$\Delta(0, t) = \tilde{A}_{mp}(t) - D(t) = B(t), \quad (34)$$

i.e., the estimation accuracy is determined by the backlog at the end of the probe. If $S(\tau, t)$ is additive, $\Delta(0, t) = 0$, i.e., both the lower and the upper estimate recover $S(\tau, t)$ exactly. Further, the backlog $B(\tau)$ that is caused by $\tilde{A}_{mp}(\tau)$ is zero during the entire probe if $S(\tau, t)$ is additive.

In the stochastic case, we obtain a non-stationary service curve estimate from Eq. (31) as $\mathcal{S}_{mp}^e(\tau, t) = \tilde{A}_{mp}(\tau, t) - B^e(t)$ using the probe defined in Eq. (33). Further, we have by use of Eq. (33) that $\tilde{A}_{mp}(\tau, t) = \tilde{A}_{mp}(t) - \tilde{A}_{mp}(\tau) = \mathcal{S}_{br}^e(\tau, t)$ since $\mathcal{S}_{br}^e(t, t) = 0$ by definition of Eqs. (24) and (25). By insertion into Eq. (31) it holds for all $\tau \in [0, t]$ that

$$\mathcal{S}_{mp}^e(\tau, t) = \mathcal{S}_{br}^e(\tau, t) - B^e(t). \quad (35)$$

We conclude that $B^e(t)$ observed by minimal probing at the end of the probe is a measure of accuracy that separates the generally conservative estimate of minimal probing from the possibly too optimistic estimate of burst probing.

To investigate $B^e(t)$, we consider the backlog expression $B(t) = \sup_{\tau \in [0, t]} \{A(\tau, t) - S(\tau, t)\}$ that follows from Eq. (17). By insertion of the probe defined in Eq. (33) we have $B(t) = \sup_{\tau \in [0, t]} \{\mathcal{S}_{br}^e(\tau, t) - S(\tau, t)\}$. Next, we substitute $\mathcal{S}_{br}^e(\tau, t) = \inf_{\psi \in \Psi_t} \{S_\psi(0, t) - S_\psi(0, \tau)\}$ from Eq. (25) so that for any sample path $\varphi \in \Psi_t$ it holds that

$$\begin{aligned} B_\varphi(t) &= \sup_{\tau \in [0, t]} \left\{ \inf_{\psi \in \Psi_t} \{S_\psi(0, t) - S_\psi(0, \tau)\} - S_\varphi(\tau, t) \right\} \\ &\leq \sup_{\tau \in [0, t]} \{S_\varphi(0, t) - S_\varphi(0, \tau) - S_\varphi(\tau, t)\} = \Delta_\varphi(0, t), \end{aligned}$$

i.e., $B_\varphi(t)$ is bounded by the maximal deviation of $S_\varphi(0, t)$ from additivity, defined as $\Delta_\varphi(0, t)$ in Eq. (26). If $S(\tau, t)$ is additive, it follows that $B_\varphi(t) = 0$ for all $\varphi \in \Psi_t$. Since $\mathbb{P}[\Psi_t] \geq 1 - \varepsilon$, it holds that $B^e(t) = 0$ and $\mathcal{S}_{mp}^e(\tau, t)$ recovers $\mathcal{S}_{br}^e(\tau, t)$ exactly. The same applies if $S(\tau, t)$ is sub-additive.

For a first example, we consider the latency-rate service curves of deterministic sleep scheduling in Fig. 5, where we also include the results of minimal probing. For the case of a transient latency depicted in Fig. 5a, minimal probing recovers the estimate of burst probing,

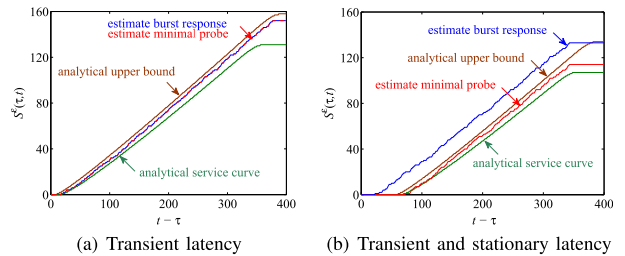


Fig. 6. Service curve estimates of random sleep scheduling plus a stationary latency. The estimate of minimal probing stays between the analytical curves, whereas burst probing exceeds the upper bound in case of a stationary latency.

which confirms that the estimate is exact. This is due to the additivity of the service.

In case of a stationary latency, as in Figs. 5b and 5c, the estimates of minimal probing and burst probing differ by a constant offset that is equal to the backlog at the end of the minimal probe $B(t) = 10$, as derived by Eq. (34). The deviation indicates that the estimate from burst probing is overly optimistic. This is a result of the super-additivity of the service that has a maximal deviation from additivity defined by Eq. (26) of $\Delta(0, t) = RT = B(t)$.

Unlike burst probing, minimal probing correctly identifies the stationary latency in the region marked with sl in Figs. 5b and 5c. It matches the analytical service curve up to the region marked with Δ , where it provides a conservative estimate with an accuracy of $\Delta(0, t)$. Finally, it reproduces the flat area in the region marked with tl in Figs. 5a and 5c that is due to the transient latency.

Next, we regard random sleep scheduling as in Fig. 4 where we find that the estimates of minimal probing and burst probing in Fig. 6a match. As before, we include as a reference an analytical upper bound as well as an analytical service curve that provides the lower guarantee specified by Eq. (12). Further, in Fig. 6b we include a stationary latency of 50. In this case, the estimates of burst probing and minimal probing differ: Both identify the correct rate, i.e., the slope, as well as the transient latency, expressed by the flat area in the upper right. In the lower left, a service of zero is caused by outages of the Bernoulli increment process as well as by the stationary latency that is recovered, however, only by minimal probing. Due to the super-additivity of the service process, the estimate of burst probing is too optimistic and exceeds the analytical upper bound. Minimal probing, on the other hand, provides a valid service curve that resides between the analytical reference curves.

4 ESTIMATION OF CELLULAR SLEEP SCHEDULING

We use the minimal probing method from Section 3.3 to estimate service curves of cellular networks with sleep scheduling, specifically DRX. Compared to state-of-the-art steady-state solutions of semi-Markov DRX models [10], [11], [12], [13], [14], our non-stationary service curve approach provides transient performance measures, such as the transient overshoot and corresponding relaxation time. Further, the service curve applies for arbitrary traffic arrivals, including TCP traffic, and is not limited to Markovian systems. While we investigate sleep scheduling of the mobile that is woken up for transmission of pending uplink data, we expect that similar studies can be performed, e.g., for the downlink, in case of mobility [53], or for green cellular networks with base station sleeping [54].




Fig. 7. The measurement setup comprises a cellular data connection from client (A) to server (D) for estimation, and a separated local control network.

In the following, we show our cellular measurement setup (Section 4.1). We present estimates of service curves that explain characteristics of the cellular data service, including transient delays caused by sleep scheduling, capacity limits, and service outages, e.g., due to the radio channel (Section 4.2). We conclude this section with a comparison of service curves obtained for different technologies, i.e., 2G (EDGE), 3G (HSPA), and 4G (LTE) (Section 4.3), as well as for different times of day showing typical diurnal behavior (Section 4.4).

4.1 Measurement Setup

The main components of our measurement setup are displayed in Fig. 7. We operate in our lab a cellular client (A) and a wired server (D), that are connected to the Internet by a major German commercial cellular provider and the German National Research and Education Network, respectively. For the cellular connection, the user equipment (UE) is a stationary category 3 Teldat RS232j-4G modem for EDGE and LTE and a Teltonika HSPA+ RUT500 modem for HSPA. The nominal uplink rates as stated by the network provider are 220 kbps for EDGE, 5,76 Mbps for HSPA, and 50 Mbps for LTE. The wired connections are 1 Gbps Ethernet links.

In addition to the cellular data connection, we maintain a separated local control network for the client and the server that permits time synchronization in the order of a few ms using the Network Time Protocol (NTP). Further, it enables operating the client and the server remotely. To automate the measurements we use the tool *sshlauncher*,⁵ that facilitates repeated execution of distributed network experiments. We repeat each measurement 100 times to obtain a statistical basis.

We use the UDP traffic generator *rude&crude*⁶ for transmission of probe traffic by the client. To measure the probe arrivals $A(t)$ and departures $D(t)$, packet traces are captured by *libpcap* at the client and the server, respectively. We choose packets of 1,400 Bytes size for HSPA and LTE and 500 Bytes for EDGE to accommodate the low uplink data rate.

4.2 Transient Service of LTE

We evaluate the performance impact of the transient phase that occurs if an uplink transmission is requested while the LTE UE is dormant, i.e., the radio resource control protocol is in idle state, see Fig. 1. After each measurement, a sufficiently long pause ensures that the UE enters idle state again. For the network under observation, this amounts to 10.5 sec [7].

5. <https://github.com/bozakov/sshlauncher>

6. <http://rude.sourceforge.net/>




Fig. 8. Transient backlog of LTE for CBR traffic.

4.2.1 Transient Overshoot and Relaxation Time

First, we consider the backlog process that is induced by constant bit rate (CBR) traffic with rates from 10 to 40 Mbps. Corresponding mean backlogs and 0.95-quantiles are depicted in Fig. 8. All curves show the typical transient overshoot and approach stationarity after a certain relaxation time. Next, we investigate the different phases in detail.

At startup, connection establishment is triggered and data are queued in a buffer at the UE. For CBR traffic the backlog increases linearly with a slope that corresponds to the traffic rate, e.g., after 120 ms, marked by a vertical line in Fig. 8, the backlog amounts to 1.2 Mb for a rate of 10 Mbps.

While we do not observe packet loss at 10 Mbps, packet loss is measured at rates of 20 Mbps and above and occurs deterministically once the backlog reaches 2.2 Mb, respectively, 200 packets. To see the deterministic behavior, note that both the mean and the 0.95-quantile show an identical plateau of 2.2 Mb that extends until 120 ms. Experiments with different packet sizes substantiate a buffer limit in the UE of about 200 packets, as packet loss occurs regardless of the packet size once the backlog reaches 200 packets.

Interestingly, the backlog starts to grow again after 120 ms, that is the time to establish the connection [7]. The increase of the backlog, computed as $B(t) = A(t) - D(t)$, is caused by packets that have been served from the buffer but are still in transmission while new packets enter the buffer at the offered traffic rate. For the mean backlog, the effect lasts for about 25 ms, that corresponds to the one-way delay (OWD) observed in [7], until packets depart from the network. Regarding the 0.95-quantile, the effect extends to up to 50 ms.

Afterwards, the backlog is depleted at a rate that is determined as the difference of the service rate and the traffic arrival rate. For traffic rates that are close to the capacity limit of about 45 Mbps, this causes significantly prolonged relaxation times until the backlog overshoot is cleared. Further, we observe an increasing volatility for higher traffic rates, as can be seen from the 0.95-quantile that grows more strongly than the mean.

Eventually, the backlog process approaches stationarity, where it is mostly caused by packets in transmission. Hence, the stationary mean backlog can be approximately determined as the product of the OWD of 25 ms and the traffic rate, e.g., for a rate of 10 Mbps a mean backlog of 0.25 Mb applies. The 0.95-quantile depicts larger backlogs that correspond to delays of up to 50 ms.

4.2.2 Non-Stationary Service Curves

Next, we identify the transient service of the LTE network by a non-stationary service curve, that is obtained by minimal probing as defined in Section 3.3. Compared to the transient backlogs shown for individual CBR traffic arrivals in

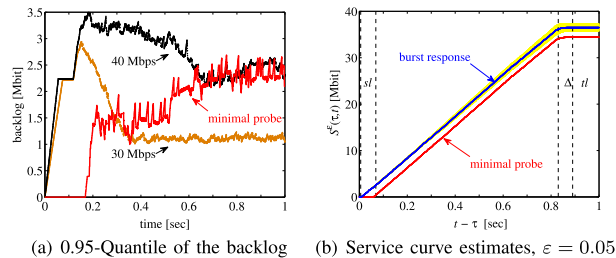


Fig. 9. Backlog of minimal probing and LTE service curve estimates.

Fig. 8, the service curve is a single characteristic function of the system that provides transient performance measures for any type of traffic arrivals.

The measurement method operates in two phases: first, an estimate of the minimal probe $\hat{A}_{mp}(\tau)$ is obtained from the burst response; second, $\hat{A}_{mp}(\tau)$ is used to estimate a service curve $\mathcal{S}_{mp}^\varepsilon(\tau, t)$. The accuracy of $\mathcal{S}_{mp}^\varepsilon(\tau, t)$ is determined as the backlog $B^\varepsilon(t)$ that is induced by $\hat{A}_{mp}(\tau)$. The accuracy can also be visualized using a service curve estimate $\mathcal{S}_{br}^\varepsilon(\tau, t)$ from the burst response, where $\mathcal{S}_{br}^\varepsilon(\tau, t) = \mathcal{S}_{mp}^\varepsilon(\tau, t) + B^\varepsilon(t)$ from Eq. (35) serves as the upper bound.

In practice, we emulate the burst function $\delta(t)$ by sending packets at a high rate that exceeds the nominal capacity of 50 Mbps for a duration of t . We choose $t = 1$ sec and measure 100 sample paths of the burst response to obtain $\mathcal{S}_{br}^\varepsilon(\tau, t)$ for $\varepsilon = 0.05$ from Eq. (25). An estimate of the minimal probe $\hat{A}_{mp}(\tau)$ follows from $\mathcal{S}_{br}^\varepsilon(\tau, t)$ by Eq. (33). We add a single packet to the start of the probe that acts as a trigger to initiate the wake-up procedure of the UE. We make use of the minimal probe to measure the backlog $B(t)$ at the end of the probe. After collecting 100 backlog samples, we select the 0.95-quantile $B^\varepsilon(t)$. A service curve $\mathcal{S}_{mp}^\varepsilon(\tau, t)$ for $\varepsilon = 0.05$ is estimated from Eq. (31) by insertion of $\hat{A}_{mp}(\tau)$ and $B^\varepsilon(t)$.

The advantage of the minimal probe is that it is adapted to the system's transient service characteristics. As a consequence, minimal probing does not sent further packets during connection establishment, such that the transient overshoot and the following relaxation time are eliminated. The effect is shown in Fig. 9a, where a representative sample path of the backlog of minimal probing is compared to the 0.95-backlog quantile of 30 and 40 Mbps CBR traffic, respectively. After the initial waiting time, the minimal probe has an average rate of 44 Mbps that results in a 0.95-backlog quantile of about 2.1 Mb at the end of the probe.

In Fig. 9b, we present the mean of ten estimates of $\mathcal{S}_{mp}^\varepsilon(\tau, t)$ and $\mathcal{S}_{br}^\varepsilon(\tau, t)$ obtained by minimal probing and the burst response, respectively. We also include the 0.95-confidence interval of $\mathcal{S}_{br}^\varepsilon(\tau, t)$, depicted as a yellow area, that confirms stable estimates. For clarity, we omit the confidence interval

of $\mathcal{S}_{mp}^\varepsilon(\tau, t)$, as it provides little additional information. Further, we note that the estimate $\mathcal{S}_{mp}^\varepsilon(\tau, t)$ shows a good accuracy, determined as the backlog $B^\varepsilon(t)$ of the minimal probe, that separates the upper bound $\mathcal{S}_{br}^\varepsilon(\tau, t)$ from the lower service estimate $\mathcal{S}_{mp}^\varepsilon(\tau, t)$. The deviation is a consequence of the super-additivity of the service, see Section 3.3.3.

Notably, the service curve estimates show the same distinct features as illustrated in Figs. 5c and 6, previously:

Service outages: For intervals $t - \tau \leq 8$ ms, both service curve estimates $\mathcal{S}_{mp}^\varepsilon(\tau, t)$ and $\mathcal{S}_{br}^\varepsilon(\tau, t)$ are equal to zero, indicating service outages on short time-scales, e.g., due to the characteristics of the radio channel.

Stationary latency: The region marked with sl expresses a stationary latency of about 50 ms. As in Fig. 5c, $\mathcal{S}_{mp}^\varepsilon(\tau, t)$ identifies this region correctly, whereas $\mathcal{S}_{br}^\varepsilon(\tau, t)$ overestimates the service and attributes the stationary latency to the region marked with Δ . The effect is due to the super-additivity of the service that is caused by the stationary latency, see Section 3.2.1.

Transient latency: The region marked with tl shows a transient latency of about 120 ms that is due to sleep scheduling.

Capacity limit: The upward segment at the center has a slope of 44 Mbps. It denotes the effective capacity limit with respect to ε . The almost constant slope evidences a stable transmission rate for intervals of $t - \tau \geq 58$ ms.

4.3 Comparison with HSPA and EDGE

We compare the performance of the LTE DRX mode with the preceding technologies HSPA and EDGE. For evaluation, we estimate non-stationary service curves using the measurement method as in Section 4.2.2. Taking into account smaller capacities and higher latencies, we reduce the probe traffic rate and extend the measurement duration accordingly.

Fig. 10 evaluates service curve estimates of LTE, HSPA, and EDGE that are obtained for each technology during the day and during the night respectively. First, we consider only the solid curves that apply for the night.

Compared to LTE, the service curves of HSPA and EDGE exhibit the same shape, however, with less pronounced segments, particularly for EDGE. Further, the estimates $\mathcal{S}_{mp}^\varepsilon(\tau, t)$ of minimal probing and $\mathcal{S}_{br}^\varepsilon(\tau, t)$ of burst probing show a larger relative difference for EDGE, indicating a lower accuracy. For HSPA the accuracy is good, like in case of LTE.

Next, we pay closer attention to the service curve estimates of HSPA in Fig. 10b. Following the same reasoning as in Section 4.2.2, we find service outages on short time-scales that result in a service of zero for intervals $t - \tau \leq 15$ ms, a stationary latency of 130 ms, a transient latency of approximately 1 sec, and a capacity limit of close to 4 Mbps.

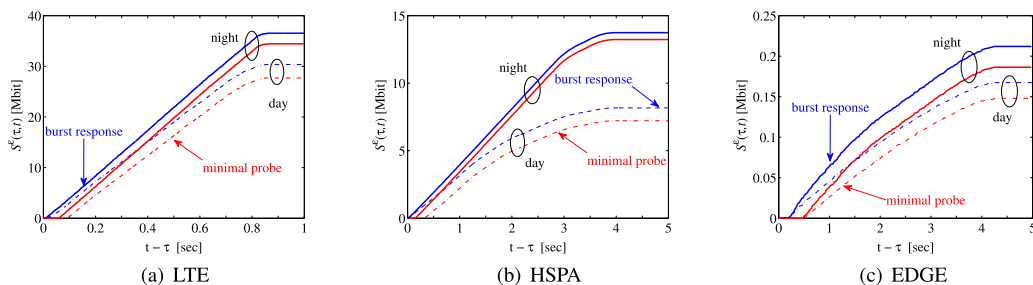


Fig. 10. Service curve estimates of LTE, HSPA, and EDGE. Solid lines show estimates obtained during the night, and dashed lines during the day, respectively.




Fig. 11. Transient 0.95-backlog quantiles of HSPA and EDGE for CBR traffic.

Further, we notice that after the transient latency, the capacity limit is approached more slowly than in case of LTE, expressed by the bent segment in the region $3 \leq t - \tau \leq 3.9$ sec. The segment emphasizes once more the advantage of minimal probing that features a correspondingly reduced rate during this time.

In case of EDGE, the service curve estimates show an even stronger bend than in case of HSPA that reveals less explicit parameters. The estimates indicate service outages of up to 200 ms, a stationary delay of 300 ms, a transient delay of 500 ms, and an effective capacity limit of about 70 kbps.

Corresponding 0.95-quantiles of the backlog for CBR traffic with different rates are shown in Fig. 11. Clearly, the long transient delay causes a backlog overshoot that can be an order of magnitude larger than the stationary backlog. Compared to LTE, the backlog of HSPA shows a less sharp peak that has a rounded top in the region around 1.5 sec. The effect is due to the slow ramp-up of the transmission rate after connection establishment. The backlog is cleared after a relaxation time of, e.g., about 5 sec for a traffic rate of 3 Mbps.

4.4 Diurnal Characteristics

We also include service curve estimates that are obtained during the day in Fig. 10, plotted as dashed lines. Compared to the night, all daytime measurements show the same general trend, i.e., a reduction of the service, that is most pronounced for HSPA. The effect may be attributed to the utilization of the network by other users that is known to follow a diurnal pattern. Further, we notice a greater volatility of the estimates at daytime, except for EDGE, that is reflected by a reduced accuracy, i.e., a larger deviation $B^e(t)$ of the lower estimate $S_{mp}^e(\tau, t)$ from the upper bound $S_{br}^e(\tau, t)$. Notably, the transient plus stationary latency, expressed by the flat area in the upper right of the curves, is, however, only marginally affected by the time of measurement.

For an example, we consider the case of LTE where the effective capacity limit is reduced from 44 Mbps at night to 38 Mbps during the day while $B^e(t)$ increases from 2.1 to 2.7 Mb. An analysis of the backlog for selected traffic rates, see Fig. 12, confirms the observations. While there is little difference between day and night for a traffic rate of 30 Mbps, we notice a significant increase of the backlog for 40 Mbps.

5 CONCLUSIONS

This work contributed a notion of non-stationary service curves that enables the analysis of transient phases. We contributed models of systems with sleep scheduling and provided insights into their transient behavior. Further, we considered measurement-based methods for identification of systems, using a black-box model. We discovered that existing probing methods cannot accurately estimate non-




Fig. 12. Transient 0.95-backlog quantiles of LTE during day and night.

stationary service curves. The difficulties were related to the non-convexity and super-additivity of the service. We devised a novel two-phase probing method. First a minimal probe is estimated that is adapted to the network. In a second step, the minimal probe is used to obtain a service curve estimate with a defined accuracy. Taking advantage of the estimation method, we reported results from a comprehensive measurement study of cellular networks with sleep scheduling, including EDGE, HSPA, and LTE. The service curve estimates showed characteristic features of the cellular data service that explained the observation of significant transient overshoots and long relaxation times. While delays in the range of seconds have been measured for EDGE and HSPA, it has been found that LTE achieves an improvement by an order of magnitude. We believe that the general non-stationary service curve approach and the modelling and measurement-based identification methods lay a foundation for transient analysis that has the potential to provide insights into a variety of relevant other cases, such as TCP slow start, transient network failures, and reactive routing.

ACKNOWLEDGMENTS

This work was supported by an ERC Starting Grant (UniQue, StG 306644). Parts of this work were presented at ITC 27 [1].

REFERENCES

- [1] N. Becker and M. Fidler, "A non-stationary service curve model for performance analysis of transient phases," in *Proc. 27th Int. Teletraffic Congr.*, Sep. 2015, pp. 116–124.
- [2] M. Mellia, I. Stoica, and H. Zhang, "TCP model for short lived flows," *IEEE Commun. Lett.*, vol. 6, no. 2, pp. 85–87, Feb. 2002.
- [3] 3GPP specification TS 36.321, "Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification," Mar. 2012, release 8, version 8.12.
- [4] G. P. Perrucci, F. H. Fitzek, G. Sasso, W. Kellerer, and J. Widmer, "On the impact of 2G and 3G network usage for mobile phones' battery life," in *Prof. IEEE Eur. Wireless Conf.*, 2009, pp. 255–259.
- [5] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4G LTE networks," in *Proc. ACM Int. Conf. Mobile Syst. Appl. Services*, 2012, pp. 225–238.
- [6] E. Chu, N.-O. Song, and D. K. Sung, "A batch packet server (BPS) based discontinuous reception mechanism for energy saving," *IEEE Trans. Veh. Technol.*, vol. 66, no. 11, pp. 10 557–10 561, Nov. 2017.
- [7] N. Becker, A. Rizk, and M. Fidler, "A measurement study on the application-level performance of LTE," in *Proc. IFIP Netw. Conf.*, 2014, pp. 1–9.
- [8] M.-R. Ra, J. Paek, A. B. Sharma, R. Govindan, M. H. Krieger, and M. J. Neely, "Energy-delay tradeoffs in smartphone applications," in *Prof. ACM Int. Conf. Mobile Syst. Appl. Services*, 2010, pp. 255–270.
- [9] S.-R. Yang, "Dynamic power saving mechanism for 3G UMTS system," *Mobile Netw. Appl.*, vol. 12, no. 1, pp. 5–14, 2007.
- [10] S.-R. Yang, S.-Y. Yan, and H.-N. Hung, "Modeling UMTS power saving with bursty packet data traffic," *IEEE Trans. Mobile Comput.*, vol. 6, no. 12, pp. 1398–1409, Dec. 2007.

- [11] L. Zhou, H. Xu, H. Tian, Y. Gao, L. Du, and L. Chen, "Performance analysis of power saving mechanism with adjustable DRX cycles in 3GPP LTE," in *Proc. IEEE 68th Veh. Technol. Conf.*, 2008, pp. 1–5.
- [12] R. S. Bhamber, S. Fowler, C. Braimiotis, and A. Mellouk, "Analytic analysis of LTE/LTE-advanced power saving and delay with bursty traffic," in *Proc. IEEE Int. Conf. Commun.*, 2013, pp. 2964–2968.
- [13] J. Wu, T. Zhang, and Z. Zeng, "Performance analysis of discontinuous reception mechanism with web traffic in LTE networks," in *Proc. IEEE 24th Annu. Int. Symp. Personal Indoor Mobile Radio Commun.*, 2013, pp. 1676–1681.
- [14] K. Zhou, N. Nikaiein, and T. Spyropoulos, "LTE/LTE-A discontinuous reception modeling for machine type communications," *IEEE Wireless Commun. Lett.*, vol. 2, no. 1, pp. 102–105, Feb. 2013.
- [15] H.-W. Ferng and T.-H. Wang, "Exploring flexibility of DRX in LTE/LTE-A: Design of dynamic and adjustable DRX," *IEEE Trans. Mobile Comput.*, vol. 17, no. 1, pp. 99–112, Jan. 2018.
- [16] K.-T. Feng, W.-C. Su, and Y.-P. Yu, "Design and analysis of traffic-based discontinuous reception operations for LTE systems," *IEEE Trans. Wireless Commun.*, vol. 16, no. 12, pp. 8235–8249, Dec. 2017.
- [17] M. K. Maheshwari, M. Agiwal, N. Saxena, and A. Roy, "Hybrid directional discontinuous reception (HD-DRX) for 5G communication," *IEEE Commun. Lett.*, vol. 21, no. 6, pp. 1421–1424, Jun. 2017.
- [18] S. M. Ross, *Introduction to Probability Models*, 9th ed. New York, NY, USA: Academic, 2007.
- [19] J. Zhang and E. J. Coyle, "The transient solution of time-dependent M/M/1 queues," *IEEE Trans. Inf. Theory*, vol. 37, no. 6, pp. 1690–1696, Nov. 1991.
- [20] C.-Y. Wang, D. Logothetis, K. S. Trivedi, and I. Viniotis, "Transient behavior of ATM networks under overloads," in *Proc. IEEE INFOCOM*, Mar. 1996, pp. 978–985.
- [21] E. de Souza eSilva and H. R. Gail, "An algorithm to calculate transient distributions of cumulative rate and impulse based reward," *Stochastic Models*, vol. 14, no. 3, pp. 509–536, 1998.
- [22] A. Horváth, M. Paolieri, L. Ridi, and E. Vicario, "Transient analysis of non-Markovian models using stochastic state classes," *Perform. Eval.*, vol. 69, no. 7, pp. 315–335, 2012.
- [23] R. L. Cruz, "A calculus for network delay, part I and II: Network elements in isolation and network analysis," *IEEE Trans. Inf. Theory*, vol. 37, no. 1, pp. 114–141, Jan. 1991.
- [24] J.-Y. Le Boudec and P. Thiran, *Network Calculus A Theory of Deterministic Queuing Systems for the Internet*. Berlin, Germany: Springer-Verlag, 2001.
- [25] C.-S. Chang, *Performance Guarantees in Communication Networks*. Berlin, Germany: Springer-Verlag, 2000.
- [26] A. Burchard, J. Liebeherr, and S. Patek, "A min-plus calculus for end-to-end statistical service guarantees," *IEEE Trans. Inf. Theory*, vol. 52, no. 9, pp. 4105–4114, Aug. 2006.
- [27] C. Li, A. Burchard, and J. Liebeherr, "A network calculus with effective bandwidth," *IEEE/ACM Trans. Netw.*, vol. 15, no. 6, pp. 1442–1453, Dec. 2007.
- [28] F. Ciucu, A. Burchard, and J. Liebeherr, "Scaling properties of statistical end-to-end bounds in the network calculus," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, pp. 2300–2312, Jun. 2006.
- [29] M. Fidler, "An end-to-end probabilistic network calculus with moment generating functions," in *Proc. 14th IEEE Int. Workshop Quality Service*, Jun. 2006, pp. 261–270.
- [30] Y. Jiang and Y. Liu, *Stochastic Network Calculus*. Berlin, Germany: Springer-Verlag, Sep. 2008.
- [31] M. Fidler, "A survey of deterministic and stochastic service curve models in the network calculus," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 1, pp. 59–86, Jan.–Mar. 2010.
- [32] F. Ciucu and J. Schmitt, "Perspectives on network calculus: No free lunch but still good value," in *Proc. ACM SIGCOMM Conf. Data Commun.*, Aug. 2012, pp. 311–322.
- [33] M. Fidler and A. Rizk, "A guide to the stochastic network calculus," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 1, pp. 92–105, Mar. 2015.
- [34] C.-S. Chang, R. L. Cruz, J.-Y. Le Boudec, and P. Thiran, "A min, + system theory for constrained traffic regulation and dynamic service guarantees," *IEEE/ACM Trans. Netw.*, vol. 10, no. 6, pp. 805–817, Dec. 2002.
- [35] R. Agrawal, F. Baccelli, and R. Rajan, "An Algebra for queueing networks with time-varying service and its application to the analysis of Integrated Service networks," *Math. Operation Res.*, vol. 29, no. 3, pp. 559–591, Aug. 2004.
- [36] C.-S. Chang and R. L. Cruz, "A time varying filtering theory for constrained traffic regulation and dynamic service guarantees," in *Proc. IEEE INFOCOM*, Mar. 1999, pp. 63–70.
- [37] Y. Jiang, "A note on applying stochastic network calculus," Tech. Rep., Norwegian University of Science and Technology (NTNU), 2010.
- [38] F. Poloczek and F. Ciucu, "Scheduling analysis with martingales," *Perform. Eval.*, vol. 79, pp. 56–72, Sep. 2014.
- [39] J. L. Doob, *Stochastic Processes*. Hoboken, NJ, USA: Wiley, 1953.
- [40] A. Rizk and M. Fidler, "Non-asymptotic end-to-end performance bounds for networks with long range dependent FBM cross traffic," *Comput. Netw.*, vol. 56, no. 1, pp. 127–141, Jan. 2012.
- [41] J. Liebeherr, A. Burchard, and F. Ciucu, "Delay bounds in communication networks with heavy-tailed and self-similar traffic," *IEEE Trans. Inf. Theory*, vol. 58, no. 2, pp. 1010–1024, Feb. 2012.
- [42] M. Fidler, R. Lübben, and N. Becker, "Capacity-delay-error-boundaries: A composable model of sources and systems," *IEEE Trans. Wireless Commun.*, vol. 14, no. 3, pp. 1280–1294, Mar. 2015.
- [43] C. Cetinkaya, V. Kanodia, and E. W. Knightly, "Scalable services via egress admission control," *IEEE Trans. Multimedia*, vol. 3, no. 1, pp. 69–81, Mar. 2001.
- [44] S. Valaee and B. Li, "Distributed call admission control for ad hoc networks," in *Proc. IEEE 56th Veh. Technol. Conf.*, Sep. 2002, pp. 1244–1248.
- [45] L. Alcuri, G. Barbera, and G. D'Acquisto, "Service curve estimation by measurement: An input output analysis of a softswitch model," in *Proc. Int. Workshop Quality Service Multiservice IP Netw.*, Feb. 2005, pp. 49–60.
- [46] M. Bredel, Z. Bozakov, and Y. Jiang, "Analyzing router performance using network calculus with external measurements," in *Proc. IEEE 18th Int. Workshop Quality Service*, Jun. 2010, pp. 1–9.
- [47] T. Hisakado, K. Okumura, V. Vukadinovic, and L. Trajkovic, "Characterization of a simple communication network using Legendre transform," in *Proc. Int. Symp. Circuits Syst.*, May 2003, pp. 738–741.
- [48] J. Liebeherr, M. Fidler, and S. Valaee, "A system theoretic approach to bandwidth estimation," *IEEE/ACM Trans. Netw.*, vol. 18, no. 4, pp. 1040–1053, Aug. 2010.
- [49] R. Lübben, M. Fidler, and J. Liebeherr, "Stochastic bandwidth estimation in networks with random service," *IEEE/ACM Trans. Netw.*, vol. 22, no. 2, pp. 484–497, Apr. 2014.
- [50] F. Michelinakis, N. Bui, G. Fioravanti, J. Widmer, F. Kaup, and D. Hausheer, "Lightweight capacity measurements for mobile networks," *Comput. Commun.*, vol. 84, pp. 73–83, 2016.
- [51] S. Shiobara and T. Okamawari, "A novel available bandwidth estimation method for mobile networks using a train of packet groups," in *Proc. ACM Int. Conf. Ubiquitous Inf. Manage. Commun.*, 2017, Art. no. 59.
- [52] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ, USA: Princeton Univ. Press, 1972.
- [53] M. Akselrod, N. Becker, and M. Fidler, "4G LTE on the road – what impacts download speeds most?" in *Proc. IEEE 86th Veh. Technol. Conf.*, 2017, pp. 1–6.
- [54] H. Tabassum, U. Siddique, E. Hossain, and M. J. Hossain, "Downlink performance of cellular systems with base station sleeping, user association, and scheduling," *IEEE Trans. Wireless Commun.*, vol. 13, no. 10, pp. 5752–5767, Oct. 2014.



Nico Becker received the Dipl-Math degree in mathematics from the Leibniz Universität Hannover, Germany, in 2012. Since 2012, he has been working toward the PhD degree in the Institute of Communications Technology, Leibniz Universität Hannover, Hannover, Germany.



Markus Fidler (M'04–SM'08) received the doctoral degree in computer engineering from RWTH Aachen University, Germany, in 2004. He was a post-doctoral fellow of NTNU Trondheim, Norway, in 2005 and the University of Toronto, ON, Canada, in 2006. During 2007 and 2008, he was an Emmy Noether Research Group leader with Technische Universität Darmstadt, Germany. Since 2009, he has been a professor of communications networks with Leibniz Universität Hannover, Germany. He is a senior member of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.