Modeling the Energy Consumption of Multi-NIC Communication Mechanisms

Stefano Ferretti, Vittorio Ghini, Moreno Marzolla, Fabio Panzieri
Department of Computer Science, University of Bologna
Mura A. Zamboni 7, I-40127 Bologna, Italy
{sferrett, ghini, marzolla, panzieri}@cs.unibo.it

***Abstract*—This paper introduces a theoretical model of the Always Best Packet Switching (ABPS) communication mechanism, a novel mechanism for wireless communications we have developed in order to enable mobile nodes, equipped with multiple network interface cards (NICs), to determine dynamically the most appropriate NIC to use as single point of access to the Internet. Using ABPS, a mobile node can seamlessly switch to another NIC, based on the performance of the NIC currently in use, in order to get better performance. ABPS supports dynamic NIC switching without causing communication interruptions at the application level. We use Markov reward models to efficiently estimate the availability and power consumption of ABPS, SIP-based mechanisms and conventional connection-oriented protocols employing a single NIC.**

I. INTRODUCTION

Modern mobile telecommunication technologies allow users to employ different networks to access Internet services, while on the move [1]–[4]. Mobile Nodes (MNs) are commonly equipped with several Network Interface Cards (NICs) (e.g., 3G, UMTS, Wi-Fi). However, there is a lack of optimized software solutions that promote a smart use of these NICs. Indeed, a MN should be able to select dynamically the "best available" (see below) network depending on its current state and on the possible connections that can be open through some specific NIC at a given moment. This selection may depend on different criteria, e.g., the "best available" network may well be the one providing lowest communication latency, or that providing communications at the less expensive rate.

Basically, the idea is to optimize the communications and the data transmission by employing one of the available networks through reconfiguration of the communication when some handoff occurs (e.g., when the active connection fails, or its throughput degrades below a certain threshold). Thus, a MN can connect to another access point available to the NIC it is currently exploiting (horizontal handoff), or disable the current NIC and switch to a different one (vertical handoff). All this should be carried out transparently to the applications and without any user intervention. In other words, the applications exploiting the connections should not notice that a reconfiguration occurred.

We have recently proposed a novel mechanism, named Always Best Packet Switching (ABPS), which can determine the most appropriate NIC to use as single point of access to the Internet at any instant [4]. The idea is that a software component running in the MN elects dynamically a "preferred" NIC to use based on the available access points. If the performance of that preferred NIC degrades or the connection fails, the MN switch to a different NIC. Specifically, for each datagram, the MN selects the best NIC to use. NICs are configured in background. Thus, when properly configured they can be exploited concurrently; moreover, it is possible to switch from a network to another transparently to the application. A prototype implementation of our ABPS mechanism is available [4], and supports interactive multimedia services based on SIP/RTP/RTCP.

Quantitative evaluation of communication mechanisms which use multiple NICs is complex and time-consuming, since it involves setting up an experimental testbed to collect measures. In particular, given that MNs operate on battery power, understanding the energy consumption of these mechanisms is important, but may require complex hardware-based measurements.

In this paper we describe a family of performance models of multi-NIC communication mechanisms based on Continuous Time Markov Chains (CTMCs) with rewards. The models can be efficiently solved, and are therefore useful to perform "what-if" analysis over large spaces of configuration parameters. Also, the models are quite general and can be easily applied to other different single- and multi-NIC communication mechanisms (e.g., [5]). Finally, by defining suitable reward structures it is possible to estimate other performance parameters such as throughput and connection reliability (mean time to failure). We use the proposed models to estimate the power consumption of ABPS and compare that to conventional connection-oriented mechanisms such as Session Initiation Protocol (SIP). Results show that the ABPS mechanism exhibits consistently better connection availability, and may achieve lower power consumption in realistic scenarios.

This paper is organized as follows. Section II describes the ABPS mechanism. Section III briefly introduces Markov chain models with rewards. In Section IV we describe the Markov models for a "simple" communication approach using a single NIC, as well as for the SIP-based and ABPS mechanisms. The models are evaluated in Section V. Finally, concluding remarks are given in Section VI.

This work has been partially funded by the Italian Ministry of Education, University and Research (MIUR) through project STEM-Net (Prot. 2009S7RLXY_003 - Year 2009).

II. THE ABPS MECHANISM

The idea behind the ABPS mechanism is simple: during transmission (and reception) the MN can use all the available NICs it has on board simultaneously, differentiating the choice of the actual NIC to use from datagram to datagram. This is carried out transparently to both the application running in the MN producing (consuming) the data being transmitted (received) and its peer at the Correspondent Node (CN); thus, applications can continue to employ classic end-to-end application-layer protocols such as TCP regardless of ABPS.

In order to let the data flow through different NICs, and deliver the related contents as a single flow to the application, additional software components are needed, in charge of managing data coming from, and going out through, different NICs. Specifically, the ABPS requires the following two additional software components: the "ABPS client proxy" and the "ABPS server proxy" (see [4] for details).

ABPS Client Proxy: The ABPS Client Proxy runs in the MN. It is in charge of maintaining the seamless multipath communication channel between an application running in its home MN and that application peer at the CN. The Client Proxy handles two types of events: those related to the management of some NIC (corresponding to some state change in the model that we will present in the next Section), and those related to the communication of the MN with its CN.

As to the NICs' management, the Client Proxy handles the following three events: i) when an access point is detected for a given NIC, the Client Proxy starts a setup procedure to activate a connection and exploit that NIC; ii) upon successful connection setup, the Client Proxy includes that NIC in a list of NICs available to transmit data; iii) when a disconnection is detected through a NIC, that NIC is removed from the list of active NICs.

As to the communication, the Client Proxy manages data coming from/to the application interacting with the Server Proxy. The data flow is fragmented in a sequence of datagrams and sent to the Server Proxy via the preferred NIC. As multiple networks can be used during the communication, it is possible that the datagrams are received out of order. To this end, a sequence numbering scheme is used that enables the receiver Server Proxy to order the received datagrams and discard possible duplicates. The Client Proxy uses ACKs to identify lost datagrams which need to be re-transmitted. When a timeout for the reception of an ACK occurs, the Client Proxy tries to retransmit that datagram through a different NIC.

ABPS Server Proxy: The ABPS Server Proxy runs on a remote host, separate from the that of the Client Proxy. The Server Proxy operates basically as a relay between the MN and its CN. The Server collects all the datagrams coming from the MN via different NICs, and sends them to the CN (which is unaware of the presence of the ABPS Server Proxy).

III. MARKOV CHAINS

A stochastic process $\{X(t), t \geq 0\}$, defined over the discrete state space $\{1, \ldots, N\}$, is a CTMC if the probability that the system is in state $X(t_{n+1})$ at time t_{n+1} only depends on the previous state $X(t_n)$ at time t_n, for any $t_n < t_{n+1}$. A CTMC can therefore be fully defined in term of an *infinitesimal generator matrix* $\mathbf{Q} = [Q_{i,j}]$, where $Q_{i,j}$ is the transition rate from state i to state $j \neq i$. The elements $Q_{i,i}$ are defined such that $\sum_{j=1}^{N} Q_{i,j} = 0$.

We denote with $\boldsymbol{\pi}(t) = (\pi_1(t), \pi_2(t), \ldots, \pi_N(t))$ the state occupancy probability vector at time t, $\pi_i(t)$ being the probability that the system is in state i at time $t \geq 0$. Given the infinitesimal generator matrix $\mathbf{Q}$ and the initial state occupancy probabilities $\boldsymbol{\pi}(0) = (\pi_1(0), \pi_2(0), \ldots, \pi_N(0))$, $\boldsymbol{\pi}(t)$ can be computed as the solution of the differential equation $\frac{d\boldsymbol{\pi}(t)}{dt} = \boldsymbol{\pi}(t)\mathbf{Q}$ given initial condition $\boldsymbol{\pi}(0)$. $\boldsymbol{\pi}(t)$ can be expressed in closed form as:

$$\boldsymbol{\pi}(t) = \boldsymbol{\pi}(0) \exp(\mathbf{Q}t) \tag{1}$$

where $\exp(\mathbf{Q}t)$ is the matrix exponential of $\mathbf{Q}t$. Under certain conditions [6], there exists a stationary state occupancy probability $\boldsymbol{\pi} = \lim_{t \to +\infty} \boldsymbol{\pi}(t)$, which is independent from $\boldsymbol{\pi}(0)$. $\boldsymbol{\pi}$ can be computed as the solution of the linear system:

$$\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}, \quad \boldsymbol{\pi}\mathbf{1}^T = 1 \tag{2}$$

Of particular interest for this paper are *Markov reward models*. We associate to each state $i \in \{1, \ldots, N\}$ a reward r_i; $\mathbf{r} = (r_1, \ldots, r_N)$ is the vector of rewards. Rewards have the following meaning: for each period of duration dt spent in state i, the total accumulated reward increases by $r_i dt$. Given the initial occupancy probability vector $\boldsymbol{\pi}(0)$, the mean total reward $R(t)$ accumulated in $[0, t]$ is

$$R(t) = \int_0^t \boldsymbol{\pi}(u)\mathbf{r}^T du \tag{3}$$

If there are no absorbing states (a state is absorbing if it has no outgoing transitions), the mean stationary reward rate can be defined as

$$R = \boldsymbol{\pi}\mathbf{r}^T \tag{4}$$

IV. MODELING NETWORK INTERFACES

We now present a model of ABPS and compare it with the following two different mechanisms: one based on a classic communication scheme employing a single network (termed "simple" in the rest of this paper), and another one based on the SIP scheme that allows the mobile terminal to use different network interfaces. We model the mechanisms above as CTMCs; then in Section V we will define appropriate reward structures to estimate various metrics related to power consumption. Our models go beyond existing results (e.g., [7]) as we consider the high-level behavior of the whole MN rather than the low level details of the network interfaces.

In the models below, we mainly focus on the local states of a MN, without going into the details of the creation and the management of a connection. In fact, this would entail to consider issues concerned with the communication with other distributed entities, and issues arising at the data link, transport and session layers. All the steps related to the configuration of a NIC, the connection of the MN to an access point, and the configuration of the NIC's IP address, are common to all

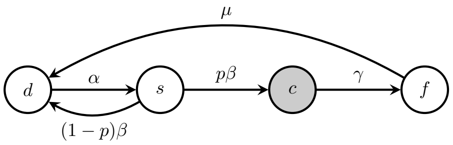

Fig. 1. Markov model of a single NIC

the considered models. Hence, we will model all the activities related to the setup of a communication with a single state. Similarly, the time interval during which the MN is properly configured to communicate with a CN is modeled with a single state, without discriminating if a real communication (e.g., a VoIP call with the CN) is active or not. In fact, this does not influence the network management at the MN.

A. Single interface

During a traditional connection-oriented communication, which employs a single network interface among those available, the behavior of a MN can be modeled using the four-state chain shown in Figure 1. At the beginning, the mobile terminal is disconnected (state d); at this state, the MN is scanning the selected network, looking for some access point to connect with. (State d includes also the time the system spends trying to connect unsuccessfully to unavailable access points.)

Once an access point is found, the MN sets up a connection with it; the setup is successful with probability p, and fails with probability $(1-p)$. After a successful setup, the NIC is configured and can be identified by an IP address (either dynamically or statically assigned). This setup phase corresponds to state s in Figure 1. This state includes also session layer setup operations, e.g., registration to some lobby server, such as the SIP registration server.

Then, the mobile terminal is connected to the service and it may communicate with other nodes (state c). If the connection is lost, for any reason, the mobile terminal passes to a state f. In this state, the communication failed and no datagrams can be delivered to the receivers, but the MN is not yet aware of such failure and continues to send data until it detects, eventually, that the connection was lost and enters in state d.

The transition from a state to another is modeled through transition rates (namely α, β, γ, μ) whose magnitude is inversely proportional to the average time a node passes in the departing state. These transitions depend on the particular NIC in use, since different network technologies usually offer different capabilities to connect to a given access point, depending on different characteristics such as its availability and signal strength while the MN is moving. In this CTMC, the MN can communicate only when it is in state c (the shaded states in Figure 1).

If we enumerate the states as $d=1$, $s=2$, $c=3$ and $f=4$, the infinitesimal generator matrix $\mathbf{Q}$ for the chain in Figure 1 is:

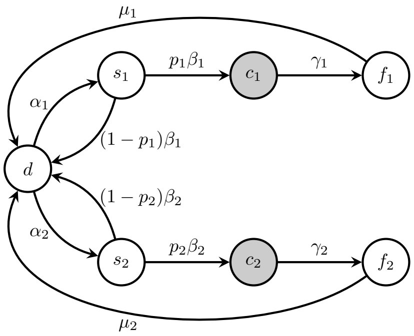

Fig. 2. SIP-based communication with two NICs

$$\mathbf{Q} = \begin{pmatrix} -\alpha & \alpha & 0 & 0 \\ (1-p)\beta & -\beta & p\beta & 0 \\ 0 & 0 & -\gamma & \gamma \\ \mu & 0 & 0 & -\mu \end{pmatrix} \tag{5}$$

The chain above is *irreducible*, meaning that it represents a strongly connected graph (each state can be reached from any other state). Therefore, there exists a stationary state occupancy probability vector $\boldsymbol{\pi}$ which can be computed using Eq. 2.

B. Multiple Interfaces: SIP

The Session Initiation Protocol (SIP) is a signaling protocol widely used for controlling communication sessions [8]. When dealing with multiple wireless NICs, it is quite usual to employ a SIP-based approach [1], [2], [4], [9]–[11]. When modeling a SIP-based approach, we assume that the MN has the possibility of accessing different networks through different NICs, among those it has on board. Once the MN selects a network, it sets up the connection through a specific NIC and then can communicate through that NIC until that connection is released or a failure occurs.

Figure 2 shows such a model where two networks are considered; the model can be easily extended to include an arbitrary number of networks (if there are K interfaces, the model will have $1+3K$ states). From the state d, the MN has multiple choices, one for each NIC (hence, two in the example depicted in Figure 2). Once a selection is made, the operation mode is the same as that described the previous Subsection. Depending on the networks involved, the transition rates may be different. Again, the states in the shaded nodes in Figure 2 represent those states where a communication is available. Note that the complexity of a SIP-based approach, with respect to a simple communication based on a single network, is hidden behind the message exchange between the MN and the SIP server, during the setup phase (state si).

C. Multiple Interfaces: ABPS

The ABPS mechanism can be modeled as the CTMC depicted in Figure 3. For the sake of simplicity, in the model

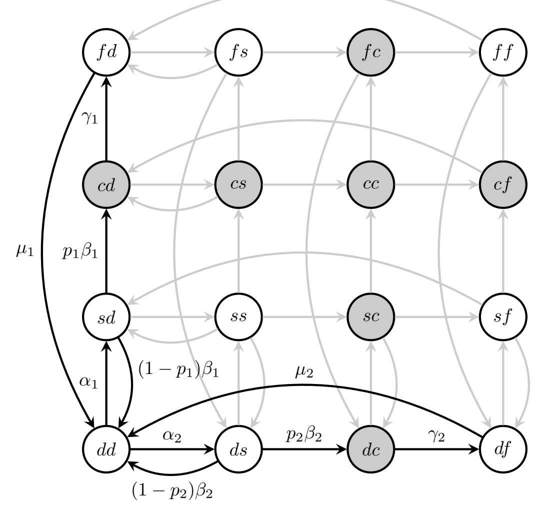

Fig. 3. ABPS communication approach, 2 NICs employed.

we assume that a MN embodies two active NICs. Needless to say, as in the case of the SIP model, the ABPS model can be easily extended to deal with MNs incorporating an arbitrary number of NICs (if there are K NICs, the ABPS model will have 4^K states). In Figure 3, states are identified by two letters; the first letter indicates the state of the first NIC, while the second indicates the state of the second NIC. The two NICs behave independently of each other; thus, the set of possible states is given by the Cartesian product of the states that each single NIC can assume. The system starts from state dd, where both NICs are disconnected; hence, the MN has no active connections, nor any network is in a set up phase. Once an access point is found for a given NIC (e.g., NIC 1), the MN sets up a connection (entering state sd) and eventually moves to the state where that NIC is connected (cd, in our example). In general, owing to the above mentioned independent behavior of the NICs, it is possible that, as a NIC goes through a state transition, the other NIC (namely, NIC 2 in Figure 3) changes its state as well, regardless of the state of the first one. As in the previous models, there is the possibility that a connection fails. In this case, the state of the considered NIC passes from state c to a state f, before going back to d. It is worth pointing out that the d, s, c states for a NIC correspond to events handled by the Client Proxy during the NIC's management. As already mentioned, state f is not handled by the Client Proxy as it corresponds to the time period during which the MN is unaware of a disconnection.

The shaded nodes in Figure 3 show that the MN is able to continue its communication when at least one of the available NICs is in the connected state c. The Markov model of ABPS assumes that the probability that both NICs change state at the same time is negligible [6].

V. ANALYSIS

The Markov models described in the previous section can be used to compute many useful performance measures. In this paper we will focus on metrics related to power consumption, which will be described shortly. We consider the scenario in which the MN has two NICs (a UMTS and Wi-Fi card, respectively). The ABPS approach employs both networks on a per-packet basis, provided that the required connections be active. The SIP approach configures one of the available connections, while the "simple" approach uses the UMTS connections only. We used the queueing toolbox for GNU Octave [12] to analyze the CTMCs.

TABLE I
PARAMETER VALUES

	α	β	γ	μ	p
UMTS (NIC 1)	$1/6.024$	$1/1.5$	$1/500$ – $1/2500$	1	0.99
Wi-Fi (NIC 2)	$1/7.5$	$1/1.5$	$1/5$ – $1/50$	1	0.9

Parameter settings: Table I shows the parameter values we consider in this paper. The values have been selected as the inverse of estimates we measured experimentally. We considered a mobile user walking at an average speed of 6 Km/h in an urban area. We assumed that, during his/her walk, the user is able to find a Wi-Fi access point every 20–50 m, on average (this is true in our university district in Bologna–Italy). Therefore, a new connection is attempted on average every $20s$, which results in $\alpha_{\text{WiFi}} = 1/20$. For the parameters γ_{WiFi} and γ_{UMTS}, which control the duration of an active connection, we consider a range of possible values and study how the performance measures change. We keep the values of other parameters fixed (β_{WiFi}, β_{UMTS}, μ_{WiFi}, μ_{UMTS}). The reason is that the duration of an active connection (controlled by the rates γ) are more likely to change due to different mobility patterns or availability of access points, while the time needed to set up a connection, or detect that a connection has failed, are likely to remain unaffected.

A. Power Consumption

We now estimate the power consumption of the three communication models; to do so, we define a reward model in which the rewards r_i represent the average power absorbed by the NICs in state i. We can estimate the total power consumption of the whole MN by using higher reward rates, representing the total device power consumption in the different states. However, in this paper we choose to focus on the NICs alone in order to reduce the number of model parameters.

While the energy profile of wireless network interfaces have been analyzed in the literature [13], [14], such measurements are extremely device-dependent. Furthermore, it is difficult to find precise measurements of the energy consumption for all the device states we consider in our model. For each NIC we define the rewards based on data from [13]; we empirically scaled down the peak power ($7.6J$ for Wi-Fi and $12.5J$ for UMTS, over a $20s$ interval). We get the following values:

	d	s	c	f
$\mathbf{r}_{\text{UMTS}}$	0.12	0.31	0.62	0.25
$\mathbf{r}_{\text{Wi-Fi}}$	0.08	0.19	0.38	0.15

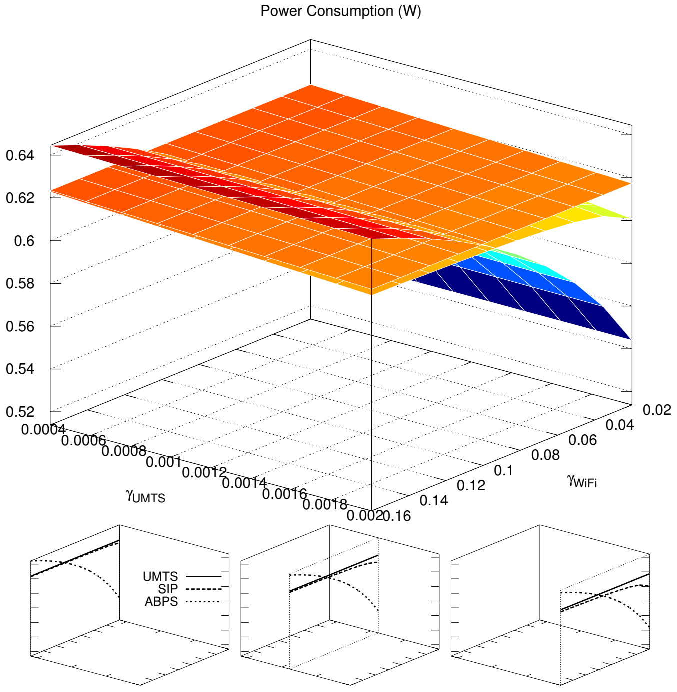

Fig. 4. Combined NICs Power Consumption (Lower is Better)

The rewards for the simple model with a single UMTS interface are set to $\mathbf{r}_{\text{UMTS}}$. For SIP, we set the rewards $\mathbf{r}_{\text{SIP}}$ as follows:

$$\begin{array}{lll} r_{\text{SIP};d} = r_{\text{UMTS};d} + r_{\text{WiFi};d} & & \\ r_{\text{SIP};s_1} = r_{\text{UMTS};s} & r_{\text{SIP};c_1} = r_{\text{UMTS};c} & r_{\text{SIP};f_1} = r_{\text{UMTS};f} \\ r_{\text{SIP};s_2} = r_{\text{WiFi};s} & r_{\text{SIP};c_2} = r_{\text{WiFi};c} & r_{\text{SIP};f_2} = r_{\text{WiFi};f} \end{array}$$

Since SIP can use a single interface at a time, we can safely assume that the unused one is either shut down or put in some deep energy saving state, resulting in almost no energy consumption. The only SIP state in which both interfaces are active is d: the reward for this state is the sum of state d rewards for the UMTS and Wi-Fi NIC. Finally, the rewards $\mathbf{r}_{\text{ABPS};i;j}$ for ABPS state (i, j) are computed as the sum of the rewards of the first interface (UMTS) in state i and the second interface (Wi-Fi) in state j. The only exception is state (c, c) in which both NICs are connected, but only the fastest one (Wi-Fi) is used for actual data transmission. We assume that the idle interface consumes 20% of its peak power. Therefore, the rewards for the ABPS model are defined as:

$$r_{\text{ABPS};i,j} = \begin{cases} r_{\text{UMTS};i} + r_{\text{WiFi};j} & \text{if } (i,j) \neq (c,c) \\ 0.2\, r_{\text{UMTS};c} + r_{\text{WiFi};c} & \text{otherwise} \end{cases}$$

Figure 4 shows the steady-state power consumption of the three communication mechanisms for all combinations of γ_{WiFi} and γ_{UMTS}. We observe that the SIP-based approach exhibits a slightly lower energy consumption with respect to using the UMTS interface only; this is expected, since SIP can use also the Wi-Fi interface, which requires less power. ABPS has higher power consumption for values of γ_{WiFi} up to around 0.05, corresponding to an average duration of the Wi-Fi connection of $1/0.05 = 20s$. If the Wi-Fi connection lasts more than that, then the power consumption of ABPS drops below that of the other two alternatives.

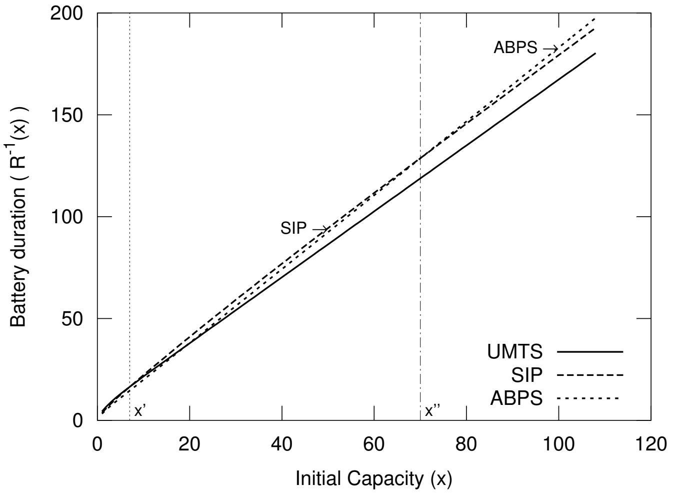

Fig. 5. Battery Duration

B. *Battery duration*

The accumulated reward of the model just introduced can be used to estimate the battery duration of a MN. Using the reward vectors defined in the previous section, the accumulated reward $R(t)$ from Eq. (3) represents the total energy consumed over the time interval $[0, t]$. If $R(t)$ is invertible, the inverse $R^{-1}(x)$ is of particular interest, since it represents the device operational period assuming that it is equipped with an "ideal" battery of initial capacity x.

While this model is not accurate–it does not take into account the discharge process of real batteries, which is not linear–it can be useful to quickly estimate the battery duration under different parameters. As a specific example, we consider the parameters in Table I, with $\gamma_{\text{WiFI}} = 1/24$ and $\gamma_{\text{UMTS}} = 1/602.4$. With these parameters, $R(t)$ is strictly monotonic increasing and hence invertible.

Figure 5 shows the battery duration $t = R^{-1}(x)$ as a function of the initial battery capacity x. We already know from the previous section that, with the parameters we considered, ABPS has the lower stationary power consumption rate. This means that, given a battery with sufficiently large initial capacity, a MN node using the ABPS communication model has the longer operational time before the battery is fully discharged. However, if the initial capacity is not large enough, interesting things happen. We observe in Figure 5 that there exist two values x', x'' such that a MN equipped with a single UMTS NIC has longer battery duration if the initial capacity is less than x''. The SIP based multi NIC approach provides longer operational lifetime is the initial capacity is in (x', x''), and then ABPS offers the better lifetime when $x > x''$.

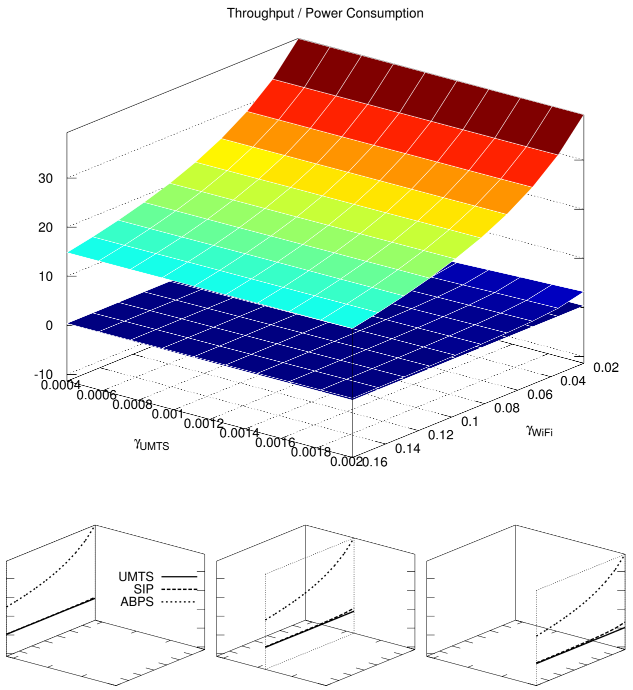

Fig. 6. Throughput / Power consumption ratio (Higher is Better)

C. Throughput / power consumption ratio

We now turn our attention to another important metric: the throughput / power consumption ratio, which indicates how many bytes can be sent with each unit of energy consumed. This metric can be computed as the ratio of the steady-state throughput and power consumption rate. The latter has already been computed in Section V-A. The stationary throughput is estimated with reward models, as follows. Let $T_{\text{UMTS}} = 0.2$ and $T_{\text{WiFi}} = 26$ be the maximum throughput (in Mbps) of UMTS and WiFi, respectively. We assign reward T_{UMTS} to state c of the single NIC model of Figure 1; all other states are assigned reward zero. For the SIP model of Figure 2 we assign reward T_{UMTS} to state c_1, and reward T_{WiFi} to state c_2. Finally, for the ABPS model of Figure 3, the reward $r_{i,j}$ assigned to state (i, j) is:

$$r_{\text{ABPS};i,j} = \begin{cases} T_{\text{UMTS}} & \text{if } j \neq c, i = c \\ T_{\text{WiFi}} & \text{if } j = c \\ 0 & \text{otherwise} \end{cases}$$

since when both interfaces are connected, ABPS uses the fastest one (WiFi) for actual communication. This reward structure can be used to compute the average steady-state throughput, which we can divide by the mean power consumption rate to obtain the throughput / energy ratio shown in Figure 6.

Unsurprisingly, the ABPS mechanism achieves higher throughput / energy ratio, since it can make better use of the Wi-Fi interface which has higher throughput and lower power consumption.

VI. CONCLUSION

In this paper we proposed a family of analytical models for the quantitative evaluation of multi-NIC communication mechanisms. The models have been used to analyze the power consumption of ABPS and compare them to results obtained for SIP-based and traditional single NIC systems. The model is based on CTMCs with rewards and can be efficiently solved. The results of this comparison show that ABPS outperforms the other two mechanisms in terms of availability in the scenarios we have considered. We found that, under some realistic conditions, ABPS also requires less energy than the other approaches. We are currently using the models described in this paper to compute other performance-related measures such as throughput and reliability; our goal is to use the results to get useful insights on how to improve ABPS.

REFERENCES

[1] P. Bellavista, A. Corradi, and L. Foschini, “Ims-compliant management of vertical handoffs for mobile multimedia session continuity,” *Comm. Mag.*, vol. 48, pp. 114–121, April 2010.
[2] G. Camarillo and M.-A. Garca-Martn, *The 3G IP Multimedia Subsystem: Merging the Internet and the Cellular Worlds*, 3rd ed. Wiley Publishing, 2008.
[3] S. Ferretti and V. Ghini, “A web 2.0, location-based architecture for a seamless discovery of points of interests,” in *The Fifth Advanced Int. Conf. on Telecommunications, AICT 2009*, 2009, pp. 226–231.
[4] V. Ghini, S. Ferretti, and F. Panzieri, “The ”always best packet switching” architecture for SIP-based mobile multimedia services,” *Journal of Systems and Software*, vol. 84, no. 11, pp. 1827–1851, 2011.
[5] T. Pering, Y. Agarwal, R. Gupta, and R. Want, “Coolspots: reducing the power consumption of wireless mobile devices with multiple radio interfaces,” in *Proc. 4th int. conf. on Mobile systems, applications and services*, ser. MobiSys ’06. New York, NY, USA: ACM, 2006, pp. 220–232.
[6] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi, *Queueing networks and Markov chains - modeling and performance evaluation with computer science applications; 2nd Edition*. Wiley, 2006.
[7] M. K. Chibesakunda and P. S. Kritzinger, “A methodology for analyzing power consumption in wireless communication systems,” Dept. of Computer Science, University of Cape Town, Tech. Rep. CS03-28-00, Dec. 15 2003.
[8] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, “SIP: Session initiation protocol,” Internet Engineering Task Force, United States, 2002.
[9] A. Devlic, “SIP-based context distribution: does aggregation pay off?” *SIGCOMM Comput. Commun. Rev.*, vol. 40, pp. 35–46.
[10] M. Femminella, R. Francescangeli, F. Giacinti, E. Maccherani, A. Parisi, and G. Reali, “Design, implementation, and performance evaluation of an advanced SIP-based call control for voip services,” in *Proceedings of the 2009 IEEE international conference on Communications*, ser. ICC’09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 1465–1469.
[11] A. Peddemors, H. Zandbelt, and M. Bargh, “A mechanism for host mobility management supporting application awareness,” in *Proc. 2nd int. conf. on Mobile systems, applications, and services (MobiSys ’04)*. New York, NY, USA: ACM, 2004, pp. 231–244.
[12] M. Marzolla, “The qnetworks toolbox: A software package for queueing networks analysis,” in *Proc. 17th int. conf on Analytical and Stochastic Modeling Techniques and Applications, ASMTA 2010, Cardiff, UK*, ser. LNCS, K. Al-Begain, D. Fiems, and W. J. Knottenbelt, Eds., vol. 6148. Springer, Jun.14–16 2010, pp. 102–116.
[13] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, “Energy consumption in mobile phones: a measurement study and implications for network applications,” in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, ser. IMC ’09. New York, NY, USA: ACM, 2009, pp. 280–293.
[14] M. Marcu, D. Tudor, and S. Fuicu, “Power efficiency profile evaluation for wireless communication applications,” in *Advances in Computer and Information Sciences and Engineering*, T. Sobh, Ed. Springer Netherlands, 2008, pp. 290–294.