Performance Characterization of a Bluetooth Piconet with Multi–Slot Packets

Daniele Miorandi, Carlo Caimi, Andrea Zanella

Performance Characterization of a Bluetooth Piconet with Multi–Slot Packets

Daniele Miorandi, Carlo Caimi and Andrea Zanella

Department of Information Engineering
University of Padova, Via Gradenigo 6/B
35131 Padova, Italy
email: {daniele.miorandi, carlo.caimi, andrea.zanella}@dei.unipd.it

***Abstract*— In this paper we present a framework for performance evaluation of a Bluetooth piconet using multislot packets. In particular, under some classical assumptions, we develop a model of a Bluetooth network and derive the complete statistics of the one–hop delay, and other significant metrics such as the channel utilization parameter. Stability conditions are investigated and it is shown that the use of multi–slot packets enlarges the achievable rate region. Simulations results are then shown, which validate the proposed analysis.**

***Index Terms*— Bluetooth, S.A.R., multislot packets, performance evaluation**

I. Introduction

Originally born as a wireless replacement for cables connecting electronic devices, Bluetooth [1] has been gaining a lot of consideration and attention by the scientific community in the last few years [2], [3], [4]. The development of this technology is now focused on the so-called Wireless Personal Area Networks (WPANs), which are foreseen as the major application for Bluetooth devices in the short and mid-term future. While the diffusion of Bluetooth devices is believed to experience a rapid growth in the next years, thanks to its expected low cost (down to 5 USD, according to leading manufacturers), the success of such a technology will still be linked to its ability to track the demand for advanced applications. Indeed, performance improvement schemes have become an active research field, in order to let Bluetooth networks being able to support demanding services, like Internet, MP3 audio and low-quality video.

The basic brick to build up a Bluetooth–based network is the so–called piconet, a cluster of nodes (not more than 8) communicating with each–other by sharing a common FH channel. Access control to the radio channel is achieved by using a master–driven TDD scheme, based on a master–slave architecture. Piconets may communicate by sharing, on a time division basis, a device, called inter–piconet unit, which acts as gateway among the piconets it belongs to.

Even if much attention has been devoted to performance evaluation of Bluetooth networks, most of the results have been obtained through numerical simulations. Indeed, the literature still lacks an in–depth analytical investigation of Bluetooth network performance. Not even the analysis of a single piconet with multi–slot packets has been completed, even if some considerable effort has been devoted in the last years [5], [6]. In this paper, we aim at providing a mathematical framework which, under some simplifying but classical assumptions, allows us to evaluate the performance obtainable with the use of multi–slot packets. In particular, we provide the packet–delay statistics for one–hop transmissions. In this sense, our approach resembles that used in [6], from which it differs for some model assumptions, in particular on the role of the master and on the possibility of having different statistics for the packet length at different nodes. Furthermore, we provide a characterization of the channel utilization parameter and an in–depth investigation of the stability regions achievable under a given Segmentation And Reassemly (SAR) policy, showing that the use of multi–slot packets provides an enlargement of the achievable rate region.

The paper is organized as follows: Section II deals with the characteristics of the Bluetooth communications profile and describes the system model which will be used for the analysis. Section III deals with the performance analysis of a Bluetooth piconet, in terms of single–hop delay, together with an investigation of the achievable capacity regions. Finally, Section IV presents some concluding remarks and open issues for future work.

This work was supported by MIUR within the framework of the PRIMO project FIRB RBNE018RFY

II. Background: the Bluetooth technology

Bluetooth operates in the 2.4 GHz ISM unlicensed band, providing a raw bit rate of 1Mb/s by using a binary Gaussian-shaped FSK modulation. In order to reduce interference with other devices operating in the ISM band, Bluetooth adopts a frequency hopping (FH) spread spectrum technique, spanning 79 RF carriers, 1-MHz wide each. In order to communicate, two up to eight Bluetooth units may connect in a small network, called piconet. In each piconet, a unit acts as master, controlling the channel access by means of a simple polling scheme. Time is divided into consecutive slots of 625μs each, that are used for downlink (master-to-slave) and uplink (slave-to-master) transmissions, alternatively, in a time division duplex (TDD) fashion. Namely, each time-slot is associated to a hop in the hopping sequence, resulting in a nominal hop rate of 1600 hop/s. The master can transmit in even–numbered time slots, whereas odd–numbered slots are reserved for slaves' transmissions. The standard provides two types of service, Asynchronous ConnectionLess (ACL) and Synchronous Connection–Oriented (SCO). The latter provides a way of sending delay–sensitive services (tipically voice) over Bluetooth, and is based on a reservation scheme which estabilishes a virtual circuit between a slave and the master. The other type of link, ACL, provides packet–switching on

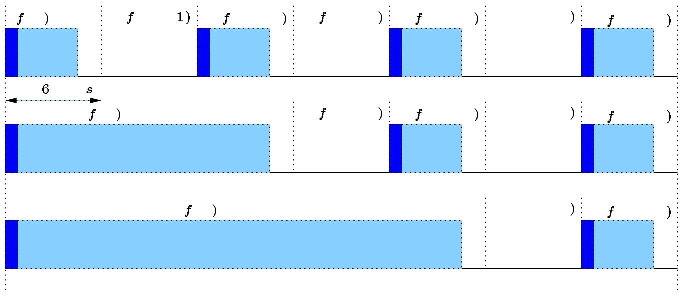

Fig. 1. Multi–slot packet transmission

Type	Slot occupancy	Max.payload length (bytes)	FEC rate
DM1	1	17	2/3
DM3	3	121	2/3
DM5	5	224	2/3
DH1	1	27	–
DH3	3	1 3	–
DH5	5	339	–

TABLE I
PACKET CHARACTERISTICS FOR ACL LINKS

the wireless channel; the standard presents six possible packet types, which differ for both coding scheme applied (either a $(15, 10)$ shortened Hamming code or none) and packet length ($1, 3$ and 5 time slots), as depicted in Fig. 1. A resume of the packet characteristics is reported in Tab.I.

Different piconets are associated to independent FH channels. This allows more piconets to share the same physical space and spectrum without increasing excessively the mutual interference. Piconet may communicate by sharing a device on a time division basis, forming what is commonly referred to as a scatternet. Both a piconet and a scatternet consisting of $M = 3$ piconets are depicted in Fig.2.

A. System Model

To carry on our analysis, we need some simplifying assumptions on the considered network. First of all, we limit our analysis to the simplest polling scheme, Pure Round Robin (PRR), also referred to as Limited-1 polling. Although many efficient polling schemes have been proposed in the last few years [7], [8], at this time available devices implement a basic PRR scheme. This is essentially due to the necessity of keeping the complexity of the firmware as low as possible, in order to reduce the manufacturing costs and lower the power consumption. PRR does not require complex logic to be

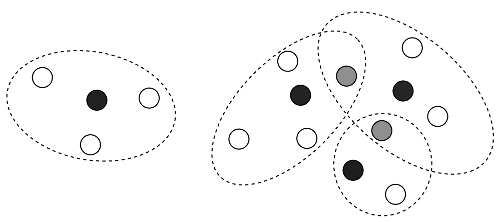

Fig. 2. Topologies of Bluetooth networks: piconet & scatternet

embedded on the chip and, thus, it results the most attractive choice for low–cost and power-aware solutions. Furthermore, in the following, we will consider ACL links only and limit our analysis to the use of unprotected packets (DHx, where $x = 1, 3, 5$). In fact, it has been shown [9] that, in most operating conditions (namely for not too low signal–to–noise ratios) the use of DHx packets results in higher goodput with respect to DMx.

We start by considering a piconet consisting of $N + 1$ nodes, $N \leq 7$; the network may, then, be represented by a set of $2N$ interacting queues. Since the solution for problems of interacting queues is still far from coming, we assume that the arrival processes at the various queues are independent. To treat also slave–to–slave communications, we will employ the classical tool of statistical routing [10], which enables us to get an approximate analysis of the network performance. In other words, the traffic from the master to a particular slave, say j, will be computed as the sum of the traffic generated in the whole piconet for that particular slave, and the various resulting flows will be considered independent.

To get a more mathematically comprehensive framework, we need to introduce some preliminary notations. Let us enumerate units in a piconet from 0 to N. Let $k \in \ 0, \ldots, N\}$ be the master ID. For each $i, j \in \ 0, \ldots, N\}$, we use the suffix (i, j) to denote the link between node i and j. Note that, in case both i and j are slave units, the link is "virtual", since each packet exchange between two slaves has to be routed through the master, that will forward the packet to the recipient slave. The arrivals are modeled with a Poisson process with bulk arrivals and rate $\delta_{i,j}$ (expressed in packets/slot). The probabilities of packets generated at node i to node j being one, three and five slots long, are $p_{i,j}(1)$, $p_{i,j}(3)$ and $p_{i,j}(5)$, with $\sum_{l \ 1,3,5} p_{i,j}(l) = 1$. We assume that no segmentation–and–reassembly (SAR) procedures take place when a slave–to–slave communication passes through the master. A complete characterization of the piconet is thus given by:

$$\mathcal{P} = \ k, \Delta(s)\}, \tag{1}$$

where $\Delta(s)$ is the polynomial end–to–end traffic matrix, whose (i, j)-th entry is defined as:

$$\Delta(s)_{i,j} = \delta_{i,j} \sum_{l \ 1,3,5} p_{i,j}(l) \quad {}^{l}. \tag{2}$$

The end–to–end traffic matrix $\Delta(s)$ can be associated to the effective traffic matrix $\Lambda(s)$, which describes the actual traffic flowing between the master unit and each slave unit:

$$\Lambda(s)_{i,j} = \lambda_{i,j} \sum_{l \ 1,3,5} \pi_{i,j}(l) e^{-sl}. \tag{3}$$

Due to our assumptions of statistical routing and absence of SAR in slave–to–slave communications, we obtain:

$$\begin{cases} \lambda_{i,k} = \sum_{j=0}^{N} \delta_{i,j}, & i = 0, \ldots, N, \quad j = k; \\ \lambda_{k,j} = \sum_{i=0}^{N} \delta_{i,j} & j = 0, \ldots, N, \quad i = k; \\ \lambda_{i,j} = 0 & \text{t r i .} \end{cases} \tag{4}$$

and:

$$
\begin{cases}
\pi_{i,}\ (l) = \frac{}{\lambda_{i,}} \sum_j \quad p_{i,j}(l)\delta_{i,j}, & i = 0, \ldots, N, \quad l = 1, 3, 5; \\
\pi\ \ _j(l) = \frac{}{\lambda\ \ _{,i}} \sum_i \quad p_{i,j}(l)\delta_{i,j} & j = 0, \ldots, N, \quad l = 1, 3, 5; \\
\pi_{i,j}(l) = 0 & i, j \neq k, \quad l = 1, 3, 5.
\end{cases} \tag{5}
$$

There is, however, a more elegant way to pass directly from $(\ )$ to $\Lambda(\ )$. Let us define the matrix Q as:

$$
Q \ = \begin{pmatrix} 0 & 0 & \ldots & 0 \\ \vdots & \vdots & & \vdots \\ 1 & 1 & \ldots & 1 \\ 0 & 0 & \ldots & 0 \\ \vdots & \vdots & & \vdots \end{pmatrix}, \tag{6}
$$

where the k-th row is the only non–zero one, and U as the all–zero matrix presenting a single one in position (k, k). Then, we may write:

$$
\Lambda(\ ) = Q \ \cdot \ \ (\ ) \cdot (I - U\ ) \qquad (\ ) \cdot Q' \ \cdot (I - U\ ), \tag{7}
$$

where I is the identity matrix.

III. Performance Analysis

The delay on the link (i, j), $d_{i,j}$, consists of two components, the access delay $w_{A_{i,j}}$ (i.e. the time it takes for a packet to get to the head of the queue and start the service) and the transmission delay $w_{tx_{i,j}}$. By using the corresponding LSTs (Laplace-Stieltjes transforms) we have:

$$
D^*_{i,j}(\ ) = W^*_{A_{i,j}}(\ ) \cdot W^*_{tx_{i,j}}(\ ), \tag{8}
$$

where $W^*_{tx_{i,j}}(\ ) = \sum_{\ ,3,5} \pi_{i,j}(k) \qquad = \frac{\Lambda(\ )_{i,j}}{\lambda_{i,j}}$.
The access delay itself may be expressed as the sum of two components:

$$
w_{A_{i,j}} = w_V \quad w_{q_{i,j}}, \tag{9}
$$

where w_V is the time elapsed between the packet arrival and the first time the queue gets the token, whereas the second term takes into account the time spent for transmitting all the packets found waiting in queue. Passing to the corresponding LSTs, we obtain:

$$
W^*_{A_{i,j}}(\ ) = W^*_V(\ ) \cdot W^*_{q_{i,j}}(\ ). \tag{10}
$$

The first term may be thought as the residual life in a renewal process with renewal period equal to the cycle time. Thus, from the theory of random look [11] we have for the pmf:

$$
f\ \ _V(a) = \frac{1 - \ \ _T\ \ (a\ \ )}{E[T_C]}, \tag{11}
$$

where F_{T_C} is the probability distribution of the r.v. T_C. Passing to the LSTs, we get:

$$
W^*_V(\ ) = \frac{1 - X^*_{T_C}(\ )}{E[T_C]} \tag{12}
$$

The other term in (9) may be found by considering the equivalent $M|G|1$ model of the queue, with an equivalent service time equal to the cycle time, thus getting, in terms of LSTs,

$$
W^*_{q_{i,j}}(\ ) = \frac{(1 - \hat{\rho}_{i,j})}{\ - \lambda_{i,j} - \lambda_{i,j} X^*_{T_C}(\ )}, \tag{13}
$$

where $\hat{\rho}_{i,j}$ is the equivalent load factor for the (i, j)-th queue (see §III-A for more details).
Now we need to find the LST of the cycle time T_C. The cycle period may be expressed as the sum of the times spent for data exchange on the (i, j)-th link, $T_{d_{i,j}}$, which, due to our assumptions, are assumed independent, and thus:

$$
X^*_{T_C}(\ ) = \prod_{i,j} G^*_{i,j}(\ ). \tag{14}
$$

According to the notation of [6] we set $G^*_{i,j}(\ ) = E[\ \ \ ^{T_{d_{i,j}}}]$.
Finally, considering that in case of no data packets available a one–slot long dummy packet (POLL/NULL) is sent, we find:

$$
\begin{aligned}
G^*_{i,j}(\ ) = (\hat{\rho}_{i,j}\pi_{i,j}(1) \quad (1 - \hat{\rho}_{i,j})) \\
\hat{\rho}_{i,j}\pi_{i,j}(3) \quad ^3 \quad \hat{\rho}_{i,j}\pi_{i,j}(5) \quad ^5 .
\end{aligned} \tag{15}
$$

On the whole, we get:

$$
\begin{aligned}
D^*_{i,j}(\ ) = \frac{1 - X^*_{T_C}(\ )}{E[T_C]} \cdot \\
\cdot \frac{(1 - \hat{\rho}_{i,j})}{\ - \lambda_{i,j} - \lambda_{i,j} X^*_{T_C}(\ )} \cdot \sum_{\ ,3,5} \pi_{i,j}(k) \qquad .
\end{aligned} \tag{16}
$$

Equation (16) provides the complete statistics of the delay on the link (i, j). In particular, the average value, $D_{i,j} = E[d_{i,j}]$ may be found to be:

$$
\begin{aligned}
D_{i,j} = -\frac{\partial D^*_{i,j}(\ )}{\partial} \Big| \quad = \frac{E[T_C^2]}{2\ [T_C]} \\
\frac{\lambda_{i,j} E[T_C^2]}{2(1 - \lambda_{i,j} E[T_C])} \quad (1 \quad 2\pi_{i,j}(3) \qquad \pi_{i,j}(5)),
\end{aligned} \tag{17}
$$

where the moments of T_C may be found by deriving expression (14).
Finally, the average packet delay may be computed as:

$$
D = \frac{\sum_{i,j} \lambda_{i,j} D_{i,j}}{\sum_{i,j} \lambda_{i,j}}. \tag{18}
$$

A. Stability Conditions

Generally speaking, the load factor ρ of a queue is defined as [11]:

$$
\rho = E[b] \cdot E[r], \tag{19}
$$

where b is the service time and r the arrival rate. A polling system employing PRR is stable if and only if the following conditions are both verified [10]:

$$
\rho = \sum_{i,j} \rho_{i,j} < 1, \tag{20}
$$

$$
\hat{\rho}_{i,j} < 1, \qquad\qquad i, j = 0, \ldots, N. \tag{21}
$$

where the first condition applies to the system as a whole, whereas the second ensures that also every single queue is stable. According to the notation previously introduced, each load factor of the i-th queue may be described by: $\rho_{i,j} = \lambda_{i,j} E[w_{tx_{i,j}}] = \lambda_{i,j} \cdot [1 \quad 2\pi_{i,j}(3) \qquad \pi_{i,j}(5)]$. The equivalent load factor, $\hat{\rho}_{i,j}$ may be calculated by considering the equivalent $M|G|1$ system, which presents a service time equal to the service time T_C. Thus, $\hat{\rho}_{i,j} = \lambda_{i,j} E[T_C]$. Substituting, after a

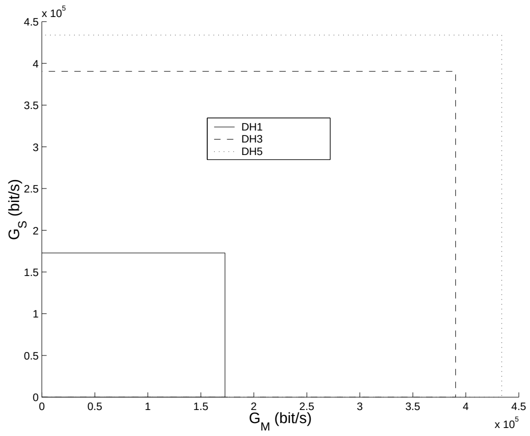

Fig. 3. Stability region, $N = 1$, in terms of offered traffic (DHx packets only)

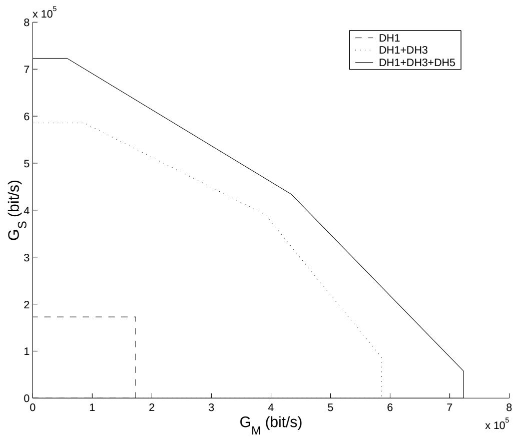

Fig. 4. Stability region, $N = 1$, in terms of offered traffic

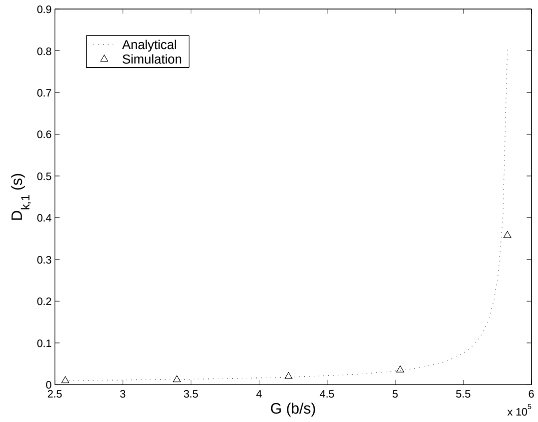

Fig. 5. Average packet delay versus offered traffic, $N = 7$

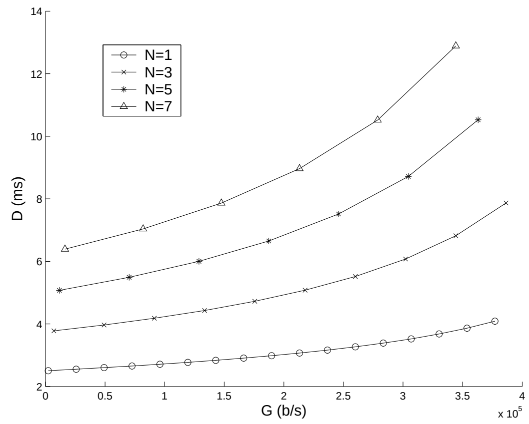

Fig. 6. Average packet delay versus offered traffic for various N, $\pi(1) = \pi(3) = \pi(5)$

few algebra we get that condition (21) implies (20) and is thus not only necessary but also sufficient to get a stable network. In Fig.3 we reported the results for the case when only DHx packets are used on both forward and reverse links for a piconet with $N = 1$ slave. The achievable rate region for $N = 1$, is shown in Fig.4. Note that, as expected, due to the higher efficiency (in terms of payload/packet length ratio), enabling the use of multi–slot packets effectively enlarges the achievable rate region. The same reasoning applies to the general case of a piconet with N slaves, where the achievable rate region may be depicted as a polyhedron in a $2N$–dimensional space.

B. Channel Utilization

In wireless systems, bandwidth is a scarce resource and, hence, it is of vital importance to fully exploit the available bandwidth. The time division duplex (TDD) architecture of Bluetooth systems makes impossible to fulfil this requirement in the case of strongly asymmetric traffic. (An example may be a multicasting in a conference room, where downlink traffic only is present). Hence, a metric of great interest is the so–called channel utilization parameter, defined as the average percentage of slots occupied by data packets. The channel utilization parameter, which will be denoted by η, may be obtained by the ratio between the total number of slots used to send data during a service cycle and the cycle duration expressed in slots. Hence, using average values we get:

$$\eta = \frac{1}{E[T_C]} \cdot \sum_{i,j} \hat{\rho}_{i,j} E[w_{tx_{i,j}}] = \sum_{i,j} \lambda_{i,j} E[w_{tx_{i,j}}]. \qquad (22)$$

C. Numerical Results

In this section we report some results obtained through numerical simulations. To validate our model, we implemented a simulator using OPNET; we considered a scenario of a full piconet ($N = 7$), with download traffic only and various SAR policies. The results, in terms of average packet delay, are plotted in Fig.5 for $\pi(l) = \frac{1}{3},\ l = 1, 3, 5$.

In Fig.6, we reported the impact of the number of slaves on the packet delay for various N, in a balanced scenario with $\pi_{i,j}(l) = \frac{1}{3}$ for any (i, j). In case of slave–to–slave communications, the end–to–end packet delay clearly depends on the order the master polls the slaves. To avoid such a dependency (which results in fairness loss), we should consider a PRR in which the polling order is randomized at each cycle; thus, an approximate expression of the packet delay for a (i, j) communication may be given by:

$$D_{i,j} = D_{i,k} + D_{k,j}. \qquad (23)$$

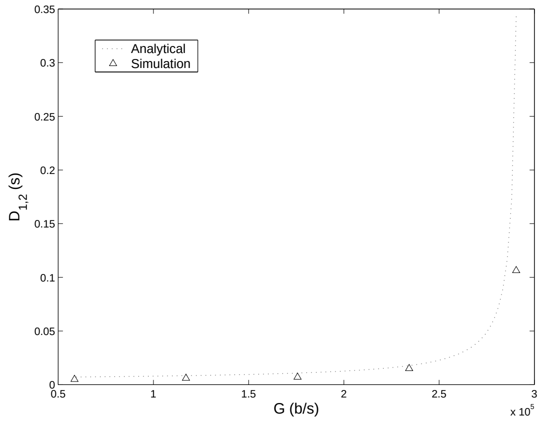

Fig. 7. Average packet delay for a slave–to–slave communication, $N = 2$

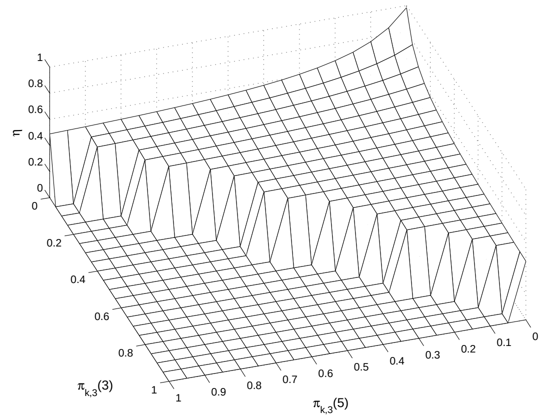

Fig. 8. Channel utilization parameteras a function of $\pi(3)$ and $\pi(5)$

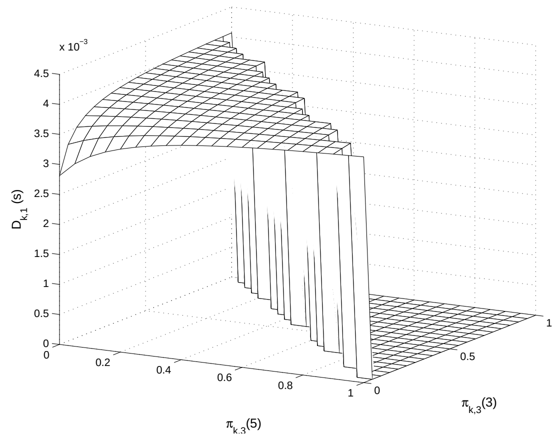

Fig. 9. Average packet delay for the delay–sensitive application as a function of $\pi(3)$ and $\pi(5)$

To show how the assumption of independent flows affects the results, we simulated a piconet with 2 slaves communicating each other: theoretical and simulation results are shown in Fig.7. It may be noted that the assumption of independent flows, although providing good results at low traffic load, leads to substantial mismatch with the simulation results as the system gets close to the stability limit. Finally, we investigated the impact of the SAR policy employed: namely, we analyzed a piconet with $N = 3$ slaves; two of them communicate with the master at a bit rate of 8 Kb/s, using $DH1$ packets only (which can be thought as a model for voice over ACL), whereas the third slave is downloading from the master at a rate of 300 Kb/s. We varied the $\pi(l)$ for the last connection, and plotted the channel utilization parameter and the average packet delay on the voice links, in Fig.8 and Fig.9 respectively. The results are worth some comments: the use of multi–slot packets, indeed, while enlarging the capacity region, achieves lower efficiency in terms of bandwidth utilization. This result is, in practice, due to the higher efficiency of multi–slot packets, and, thus, the terms in (22) decreases when longer packets are used. Furthermore, longer packets may result in detrimental impact on links carrying delay–sensitive flows. In a QoS–oriented scenario, it is clear that the choice of using multi–slot packets should be somehow negotiated at the master side, taking into account the possibly negative effects on link performances.

IV. Conclusion

In this paper we presented a mathematical framework, based on queueing theory tools, for performance evaluation in a Bluetooth piconet using multi–slot packets. We discussed stability conditions and showed that the use of multi–slot packets, due to their higher efficiency, effectively enlarges the boundaries of the achievable capacity region. We showed how the channel utilization parameter, a significant metric in wireless networks performance, may be calculated, and discussed the tradeoffs involved in the choice of a given SAR policy. Three subjects seem to be of great interest for future work. The first is the study of how channel failures impact on piconet performances, and their integration into a queueing–theoretic framework. The second is the effective modeling of more realistic SAR policies, in which the packet length is chosen in function of the buffer length, and, as a consequence, the analysis of the packet delay at L2CAP level. The last interesting issue would be the extension of such a framework to accomodate a scatternet structure, where other problems (the modeling of the gateways' behavior and inter–piconet interference) arise.

References

[1] Bluetooth special interest group. [Online]. Available: http://www.bluetooth.com

[2] J. C. Haartsen, "The Bluetooth radio system," *IEEE Pers. Comm.*, vol. 7, no. 1, pp. 28–36, Feb 2000.

[3] J. Haartsen, M. Naghshineh, J. Inouye, O. Joeressen, and W. Allen, "Bluetooth: Vision, goals, and architecture," *ACM Mobile Computing and Communications Review*, vol. 2, no. 4, pp. 38–45, Oct 1998.

[4] P. Johansson, M. Kazantzidis, R. Kapoor, and M. Gerla, "Bluetooth: an enabler for personal area networking," *ACM/IEEE Trans. Net.*, vol. 15, no. 5, pp. 28–37, Sep.-Oct. 2001.

[5] R. Bruno, M. Conti, and E. Gregori, "Bluetooth: architecture, protocols and scheduling algorithms," *ACM Cluster Computing*, vol. 5, pp. 117–131, Apr 2002.

[6] V. Misic and J. Misic, "On Bluetooth piconet traffic performance," in *Proc. of IEEE PIMRC 2002*, Lisbon, Portugal, 2002.

[7] A. Capone, M. Gerla, and R. Kapoor, "Efficient polling schemes for Bluetooth picocells," in *Proc. IEEE ICC*, Helsinki, 2001.

[8] M. Kalia, D. Bansal, and R. Shorey, "Data scheduling and sar for Bluetooth mac," in *Proc. IEEE VTC (Spring)*, Tokyo, 2000.

[9] M. C. Valenti, M. Robert, and J. H. Reed, "On the throughput of Bluetooth data transmissions," in *Proc. of IEEE WCNC 2002*, Orlando, 2002.

[10] H. Levy and M. Sidi, "Polling systems: applications, modeling and optimization," *IEEE Transactions on communications*, no. 10, pp. 1750–1760, Oct 1990.

[11] L. Kleinrock, *Queueing Systems*. New York: John Wiley & Sons, 1975.