On Efficient Topologies for Bluetooth Scatternets[*]

Daniele Miorandi, Arianna Trainito, and Andrea Zanella[1]

Department of Information Engineering
University of Padova, Via Gradenigo 6/B
35131 Padova, Italy
{daniele.miorandi, orsetto, andrea.zanella}@dei.unipd.it
phone [fax]: +39.049.827 7753 [7699]

Abstract. In this paper we investigate the relationship between network capacity and topology for Bluetooth scatternets. We start by considering the intrinsic capacity limits of a scatternet structure, and we show that limiting capacity may be achieved for very local traffic and under specific conditions on the scatternet structure. Then, we provide a description of the performance achievable with two basic scatternet configurations, namely star and closed–loop topologies, and show the role played by inter–piconet interference in the choice of efficient configurations. Finally, we present some efficient topologies, based on Platonic solids structure.

Keywords: Bluetooth, scatternet, efficient configuration, performance evaluation, topology.

1 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. The development of this technology is now focused on the area of the so–called Wireless Personal Area Networks (WPANs), where Bluetooth is expected to play a major role in the short and mid-term future. The commercial success of WPANs is intimately linked to their ability to support advanced digital services, like audio and video streaming, web browsing, etc. [2]. In such a scenario, the performance aspects of the radio technologies involved appear of primary importance.

Bluetooth has been designed to work in a scattered ad–hoc environment, where multiple independent overlapping networks, called piconets, may coexist and be inter–connected to form a multi-hop network, called *scatternet*. Recently, much attention has been devoted, by both academic and industrial world, to the issues concerning scatternets formation and management. In particular, scatternet formation algorithms [3],[4],[5],[6] have been widely investigated, and many solutions have been proposed to build up a scatternet starting from disconnected units. The focus is now moving to the characterization and design of efficient scatternet topologies [7], [8], [9], since it is

---

clear that the way piconets are interconnected to form scatternets may dramatically impact on network performance. To the best of authors knowledge, however, the literature still lacks in-depth investigation of the optimal scatternet topologies.

Clearly, the optimality of a scatternet topology does depend on the performance indexes considered. Typically, scatternet performance metrics are number of piconets, number of bridge nodes, number of roles per node etc. When traffic connections are considered, average traffic delay, throughput and maximum traffic delay are shown as metrics of interest. Nevertheless, packet delay may lead to misleading conclusions whether network capacity is not taken into account. In particular, results obtained by simulating an over–loaded network may be affected by the setting of simulator parameters, like buffers capacity, simulations run–time, traffic generation pattern, and so on. For this reason, we prefer to consider a different performance metric, namely the *network capacity*. This metric represents the maximum aggregated average traffic that nodes can inject into the network without overflowing.

In general, we say that a network configuration is *stable* if the total traffic offered to the network does not exceed the network capacity, otherwise the network is said to be *unstable*. It may be worth underlining that, for unstable configurations, some traffic connections will either experiment always increasing average packet delays or packet losses due to overflows. Conversely, in stable configurations the average packet delay is limited and, provided that the dimensions of the buffers are adequately chosen, packets are never dropped because of overflows.

In this paper, we investigate the relation between the scatternet topology, i.e., the way piconets are interconnected to form the scatternet, and the network capacity. We first discuss some conditions to achieve efficient piconets interconnection. Then, we investigate the performance achieved by some specific scatternet topologies, both "planar" and "solid", in case of uniform traffic matrix, that is, assuming that every node in the network generates an equal amount of traffic towards any other node.

This work is a first attempt to provide a mathematical insight into the relation between scatternet topology and performance. The analysis we provide allows us to determine some optimality criteria and to evaluate the performance achieved by some scatternet topologies. Such an analysis may be worth, for instance, for designing networks of static sensors or domestic appliances, where the end-to-end traffic matrix among the nodes may be known a priori. At this phase of the work, however, many aspects related to the network creation and management have still to be considered and the discussion presented in this paper is far to be complete. Nevertheless, it may provide a starting point for further, more practical–oriented, investigations in the field.

The paper is organized as follows. Section 2 outlines the primary characteristics of the Bluetooth technology. Section 3 describes the system model used for our analysis. Section 4 provides the analysis of the limiting performance, in terms of capacity, achievable by Bluetooth scatternets. In Section 5, system performance is evaluated under the assumption of a uniform traffic matrix. Section 6 concludes the paper with some remarks and directions for future work.

2   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 makes use of a frequency hopping (FH) spread spectrum technique, spanning 79 RF carriers, 1-MHz wide each. From 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, while the other units act as *slaves*. Time is divided into consecutive slots of 625μs each, which 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.

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. Piconets may communicate by sharing a device on a time division basis, forming what is commonly referred to as a *scatternet*. The shared device, usually named *gateway*, spends a given amount of time in each piconet it belongs to, forwarding the inter–piconet traffic. The scatternet configuration, in turn, is determined by:

- the assignment of the master and slave roles in each piconet composing the scatternet;
- the definition of the shared units.

3   System Model

Before presenting our analysis, we need to introduce some simplifying assumptions for the network behavior.

At this phase of the work, we assume all nodes are in range, neglecting the spatial constraints that, however, may have a dramatic impact on the scatternet topology. For the sake of simplicity, we assume that only one–slot packets are used (namely, $DH1$ packets). Furthermore, we limit our analysis to the simplest polling scheme, namely Pure Round Robin (PRR). With PRR, all slaves get polled cyclically by the master, whether they have data to transmit or not. Although many efficient polling schemes have been proposed in the last few years [10–12], at this time available devices basically support only PRR. This is essentially due to the necessity of keeping the complexity of the firmware as low as possible, in order to reduce manufacturing costs and power consumption. PRR does not require complex logic to be embedded on the chip and, thus, it results the most attractive choice for low–cost, energy–conserving products.

Finally, we neglect the *switchover time*, i.e., the time a gateway spends to switch from one piconet to another. Such a time may determine a capacity loss to the system, since during the switchover time the gateway is not synchronized to any piconet and, thus, it cannot communicate. Nevertheless, the capacity loss may be reduced by increasing the *persistence time*, i.e., the period the gateway spends in each piconet before switching to the next one. Hence, we assume that the switchover time is a small fraction of the

persistence time, so that the capacity loss caused by the inter–piconets switchover time may be neglected.[1] Furthermore, we suppose gateways are equally shared among the piconets they belong to, so that a gateway connected to γ_i piconets spends $\frac{1}{\gamma_i}\%$ of the time in each piconet it belongs to.

Let us denote by N the total number of nodes (which will be considered static) in the scatternet and by M the number of resulting piconets.[2] For the generic piconet $i \in \{1, 2, \ldots, M\}$, let n_i be the total number of nodes (master, slaves and gateways) the piconet consists of. Moreover, let ϑ_i denote the number of units that act as gateways in the i-th piconet. Clearly, it has to be $\vartheta_i \leq n_i$ and $n_i \leq 8$, since a piconet cannot have more than 8 active units. It may be worth noting that each gateway is counted in the ϑ_i of all the piconets it belongs to. The total number of gateways in the scatternet is denoted by Θ. Thus, the following relation holds:

$$N = \sum_{i=1}^{M} (n_i - \vartheta_i) + \Theta \ . \tag{1}$$

We denote by Δ the *end–to–end traffic matrix*, so that the (i, j)–th entry, $\delta_{i,j}$, is the average user traffic (expressed in packet/slot) generated by node i to node j, with $i, j \in \{1, 2, \ldots, N\}$. Note that, the end–to–end traffic matrix Δ refers to the traffic produced at the transport layer, and, hence, it does not depend on the specific network topology.

Given the end–to–end traffic matrix, network topology and routing tables[3] we can univocally determine the *effective traffic matrix* Λ, which gives the actual data traffic flowing in the physical network links. Since, in Bluetooth, direct communication can occur between master and slave (or gateway) units only, for each pair (i, j) of slave units, the (i, j)–th entry, $\lambda_{i,j}$, of the effective matrix Λ is always equal to zero.

4   Network Capacity

We are interested in determining the ability of a given scatternet topology to carry information flows. Such an ability will be measured in terms of *network capacity*, defined as the maximum aggregate traffic that can be injected into the network preserving the network stability. In this context, a network is said to be stable if, for any of its nodes, the number of buffered packets does not grow indefinitely.

As a first step, network capacity will be evaluated without considering the inter–piconet interference [13, 14]. The effect of interference, however, will be considered in the successive sections.

---

[1] Note that, the persistence time cannot be increased indefinitely, since the longer the persistence time, the greater the forwarding delay introduced by the gateway and the buffer size required to hold the packets that have to be forwarded to the other piconets.

[2] We have as many piconets as nodes playing the role of master.

[3] In the following we will consider balanced routing algorithms, in the sense that, if node A wants to communicate with node B and there are P shortest paths, the traffic $A \longrightarrow B$ is equally divided among the P routes.

Given the scatternet topology, we can define the set Ω of all the end–to–end traffic matrices Δ that preserve the network stability.[4] Then, for a traffic matrix $\Delta = [\{\delta_{i,j}\}]_{i,j=1,2,\ldots,N}$, we define the *offered load*, $G(\Delta)$, as:

$$G(\Delta) = \sum_{i=1}^{N} \sum_{j=1}^{N} \delta_{i,j} \, . \tag{2}$$

It may be worth stressing that G gives the effective aggregate user traffic offered to the network, since the traffic carried over multiple hops is counted just once. Given these preliminary definitions, we can finally define the **network capacity** C as follows:

$$C = \max_{\Delta \in \Omega}\{G(\Delta)\} \, . \tag{3}$$

Analogously, for a given network topology, we can define the set Ω^* of all stable effective matrices Λ. Hence, we define the *aggregated throughput* $S(\Lambda)$ of the network as:

$$S(\Lambda) = \sum_{i=1}^{N} \sum_{j=1}^{N} \lambda_{i,j} \, ; \tag{4}$$

where also the traffic forwarded in a multi–hop fashion is taken into account. The parameter S gives the total traffic flowing in the physical links of the network. Finally, we define the **transport capacity** T as:

$$T = \max_{\Lambda \in \Omega^*}\{S(\Lambda)\}. \tag{5}$$

Trivially, under stability conditions, the offered load G cannot exceed the aggregated throughput S and the two quantities are equal if all the traffic goes through a single hop. Taking the maxima of G and S and using the notation defined above, we get the following lemma.

Lemma 1. *The network capacity C equals the transport capacity T: $C = T$.*

Proof. First we note that, for a given scatternet topology, the set of stable end–to–end matrices Ω includes the set of stable effective matrices Ω^*. Thus, if $\Lambda \in \Omega^*$ achieves the transport capacity T, then Λ belongs to $\Omega^* \subset \Omega$ and, consequently, by taking $\Delta = \Lambda$, we get an offered traffic $G(\Delta) = T$. Therefore, $C \geq T$.
Now, let $\Delta \in \Omega$ achieve the network capacity C. Thus, the corresponding effective matrix Λ, which is univocally associated to Δ (the mapping is surjective but not one–to–one), is stable: $\Lambda \in \Omega^*$. Since $S(\Lambda) \geq G(\Delta) = C$, then $T \geq C$ and the lemma is proved.

In short, Lemma 1 trivially says that the maximum aggregate traffic that can be injected into the network, maintaining stability, is equal to the aggregate link capacity of the network. Consequently, the optimal network utilization is reached when communication occurs between nodes directly connected.
In other words, the system capacity is reached when no forwarding over multiple hops is needed (local traffic) and, thus, all traffic gets through a single hop only.

---

[4] Note that, Ω is strictly related to the network topology.

4.1   Capacity of isolated piconets

As far as the capacity of a single piconet is concerned, it is clear that, under the collision–free MAC scheme provided by the Bluetooth system, the resulting capacity is $C = 1$ packet/slot. Hence, we have the following lemma:

Lemma 2. *Given M isolated piconets,[5] the network capacity is equal to $C = M$ packet/slot.*

This is clearly an upper bound on the capacity achievable by M piconets, since the effect of inter–piconet interference is not taken into account.

4.2   Capacity of connected piconets

In our lexicon, a piconet is said to be *connected* if it shares at least one unit with some other piconet.
Let us focus on a piconet with a single gateway, and let t be the gateway ID. Let us assume the gateway is shared among β piconets in such a way it is present in each piconet for a fraction $\alpha = 1/\beta$ of the time. We may distinguish two cases.

a) The shared unit is the master: $t = k$. Hence, the piconet is active only for a fraction α of the time and the piconet capacity is $C = \alpha$ packet/slot.
b) The shared unit is a slave: $t \neq k$. Therefore, the piconet capacity is $C = 1$ packet/slot when the piconet consists of more than two nodes, $C = \alpha$ packet/slot otherwise.

Two considerations arise from this simple analysis: to approach the bound given by Lemma 2 all interpiconet units ought to be slaves in all the piconets they belong to. Furthermore, to fully exploit the transport capacity of the network, each piconet has never to be empty, i.e., at least a slave must be present in each piconet at any time. Summing up, we get the following proposition:

Proposition 1. *A scatternet consisting of M piconets may achieve an aggregated throughput of at most M packet/slot. This happens if the following conditions hold:*

1. *the shared units are slaves in all the piconets they belong to;*
2. *in each piconet, there must be at least one slave present at any time.*

It is not difficult to find structures that satisfy the above conditions. An example is given by the simple chain configuration depicted in Fig. 1, where it suffices to synchronize the gateways in an odd–even fashion, in such a way that there is not more than one gateway in each piconet at any time.

 Under the conditions of Proposition 1, the capacity turns out to be linearly increasing with the number of piconets.[6] Therefore, it may be interesting to determine the maximum capacity a scatternet of N nodes may achieve. The answer is given by the following proposition:

---

[5] A piconet is said to be *isolated* when the nodes in the piconet do not belong to any other piconet.
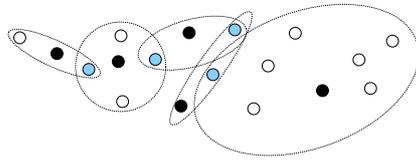[6] We recall that, in this section, interpiconet–interference is not considered.

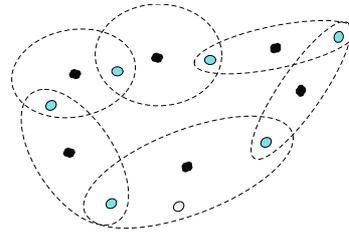**Fig. 1.** Chain configuration.

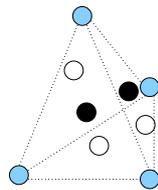**Fig. 2.** Closed loop configuration.

Fig. 3. Tetrahedral configuration (only the nodes of two piconets are shown).
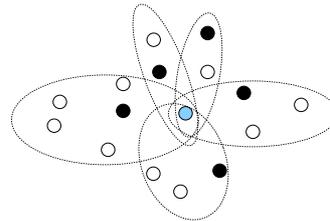

Fig. 4. Single gateway configuration.

Proposition 2. *A scatternet consisting of N interconnected nodes may achieve an aggregated throughput of at most $\lfloor \frac{N}{2} \rfloor$ packet/slot.*

Proof. The proof is constructive. We first find a configuration which achieves the bound. To this purpose, we organize the nodes in a closed–loop topology, as depicted in Fig. 2, so that each piconet consists of $n = 3$ nodes (except for at most one piconet, with $n_M = 4$ if N is odd) and $\vartheta = 2$ gateways. The total number of gateways is, then, equal to $\Theta = \lfloor \frac{N}{2} \rfloor$. The regularity of the structure enables us to find a scheduling of the gateways such that there is always at least one slave (or gateway) in each piconet. Thus, conditions of Proposition 1 are satisfied and, consequently, the limit capacity of the network is equal to M packet/slot.

To prove the proposition, it remains to verify that the capacity of a scatternet containing N nodes cannot exceed $\lfloor \frac{N}{2} \rfloor$ packet/slot. Let us assume, by contradiction, that a scatternet topology of N nodes achieves a network capacity C greater than $\lfloor \frac{N}{2} \rfloor$ packet/slot. This requires that, at each time, C master–slave pairs are actively exchanging data–packets. Since we cannot extract more than $\lfloor \frac{N}{2} \rfloor$ distinguished pairs of nodes out of a total of N nodes, the assert is clearly a contradiction. Hence, the scatternet capacity cannot exceed the limit $\lfloor \frac{N}{2} \rfloor$ packet/slot and the proposition is proved.

Note that, the above results may provide useful insight into efficient configuration design. For instance, we can observe that the closed–loop configuration is asymptotically optimal, since it achieves limiting capacity for any N. For a fixed N, however, it is easy to find different configurations which achieve the bound. For example, when $N = 8$, the limiting capacity may be achieved by a tetrahedral structure, as that shown

in Fig. 3. Another interesting observation concerns the ring architecture considered in [15], where all nodes act as master/slave relays. It is easy to prove that also this config-uration (which can be, in a certain sense, considered a degeneration of the closed–loop architecture) is asymptotically optimal. (In fact, it seems that these are the only two asymptotically optimal architectures). The ring topology, however, presents scalability problems and performance rapidly worsens when traffic has to be forwarded in a multi–hop fashion. Furthermore, it can be observed from the above propositions that also tree structures [4] are inefficient from the point of view of network capacity (and, hence, spectrum utilization).

In the above analysis, the inter–piconet interference has been neglected. However, the impact of inter–piconet interference will be considered within the following sec-tions, where some possible scatternet configurations are analyzed in case of homoge-neous end–to–end traffic matrices.

5   Efficient configurations for Bluetooth scatternets

From the analysis presented in the previous section, it turns out that network capacity may be achieved with very local traffic only. However, a scatternet topology ought to be efficient under a wide range of traffic patterns.
In this section we provide a sort of worst case analysis, considering uniform end–to–end traffic matrices, where each node generates, on average, the same amount of traffic to every other node in the network. It is clear that, in this situation, any reordering of the nodes would not increase the traffic locality. Consequently, network performance for uniform traffic matrices will dependent on the scatternet topology only, and not on single node position.
Therefore, we consider a uniform end–to–end traffic matrix Δ, having the following structure:

$$\Delta = \begin{pmatrix} 0 & \delta & \dots & \delta \\ \delta & 0 & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots \\ \delta & \delta & \ddots & 0 \end{pmatrix} ; \tag{6}$$

where the (i,j)-th entry, $i,j = 1,2,\dots,N$, represents the average traffic generated from node i to node j (in packet/slot). In this situation, which models a cooperative computing scenario, the scatternet configuration we look for should present three desir-able properties:

1. **efficiency**: the scatternet ought to support maximum traffic;
2. **fairness**: all the piconets should experience the same traffic load;
3. **scalability**: the resulting structure should scale well with respect to the number of nodes.

The simplest configuration that meets the above requirements is the single–gateway topology. In this case, the resulting scatternet is star–shaped, as depicted in Fig. 4. Next,

we will consider the closed–loop topology, to end up with a discussion on some efficient configurations based upon platonic solids.

We first introduce some stability issues, which are at the basis of the successive capacity evaluation. Then, we introduce a model for inter–piconet interference, derived from [13]. Finally, we show some numerical results.

5.1   Stability conditions

Let us consider again the effective traffic matrix Λ, introduced in Section 3. Each element $\lambda_{i,j}$ of this matrix represents the average number of packets per slot sent along the link from node i to node j. We remark that effective links are present only between master and slave or gateway units. Thus, if $\lambda_{i,j} > 0$ either i or j has to be a master unit, while the other has to be a slave or gateway unit. Applying basic queueing theory, it can be shown that for an admissible link (i,j), the stability condition, under the assumption of Pure Round Robin (PRR), is given by [16]:

$$\lambda_{i,j} < \frac{1}{2(n_{i,j} - 1)},$$

where $n_{i,j}$ is the number of nodes in the piconet that i and j belong to. The above inequality trivially means that for a piconet to be stable each link must not receive, on average, more than a packet per polling cycle.

When connected piconets are considered, however, we have to take into account the presence of gateways, which will be active only for a fraction of the time. Hence, adding dependence on the time t, we may write:

$$\lambda_{i,j}(t) < \frac{1}{2(n_{i,j}(t) - 1)} \cdot b_{i,j}(t), \qquad (7)$$

where $n_{i,j}(t)$ represents the number of nodes in the considered piconet at a given time instant t, and $b_{i,j}(t)$ is either 1 or 0 according to the active or idle state of link (i,j) at time t. Clearly, both $n_{i,j}(t)$ and $b_{i,j}(t)$ depend on the presence of gateways in the piconet of interest. For the regular configurations we are interested in, the network shows a periodical behavior, in the sense that there exist a cycle period of duration T_C,[7] such that $n_{i,j}(t) = n_{i,j}(t + kT_C)$ and $b_{i,j}(t) = b_{i,j}(t + kT_C)$ for each integer k. Furthermore, we may assume that the cycle time T_C may be partitioned in intervals in which the scatternet configuration does not change. Denoting with Q the cardinality of the partition, we have:

$$(0, T_C] = (0, T_1] \cup (T_1, T_2] \cup \ldots \cup (T_{Q-1}, T_Q]. \qquad (8)$$

Let n_m and b_m be the values assumed by $n_{i,j}(t)$ and $b_{i,j}$ in $(T_{m-1}, T_m]$, respectively. Furthermore, let $\overline{\lambda}_{i,j}$ be the average traffic offered over a period T_C, given by

$$\overline{\lambda}_{i,j} = \frac{1}{T_C} \int_0^{T_C} \lambda_{i,j}(t)dt \ .$$

---

[7] It is worth stressing that this cycle time has not to be confused with the cycle time of the polling algorithm.

Table 1. Relationship between N, n and M for different scatternet topologies

Topology	Number of piconets
single–gateway	$M = \frac{N-1}{n-1}$
closed–loop	$M = \frac{N}{n-1}$
tetrahedron	$M = 4$

Averaging both sides of inequality (7) over the cycle period, we get:

$$\overline{\lambda}_{i,j} < \frac{1}{2T_C} \sum_{m=1}^{Q} \frac{b_m}{n_m - 1}(T_m - T_{m-1});$$ (9)

where T_0 is assumed equal to zero. Thus, inequalities (9) for each pair (i,j) of network nodes gives the overall network stability conditions.

5.2   Performance evaluation for uniform traffic

Given the conditions on the offered traffic that assure network stability, we can derive the maximum throughput that can be achieved by the network for different scatternet topologies. Once more, we neglect the effect of inter–piconet interference, thus obtaining upper bounds for network performance. Such bounds will be, then, used in the following subsection to derive performance estimation in the presence of interference. Let N be the total number of units in the scatternet; M the number of piconets and n_i, $i = 1, \ldots, M$ the number of hosts in the i-th piconet. For the sake of simplicity, in the following we assume that all the piconets have the same number of units, i.e. $n_i = n, i = 1, \ldots, M$. Tab. 1 shows the relationship that ties N, n and M for some scatternet topologies that will be consider in the following.

The basis for the evaluation of network capacity is the stability condition given by (9). Applying equation (9) to the end–to–end traffic matrix Δ given by (6), we find a condition of the form $\delta < \zeta$, where ζ depends on the scatternet topology. Thus, the offered traffic $G(\Delta)$ will satisfy the relation $G(\Delta) < N(N-1)\zeta$. Considering the set $\mathcal{U}$ of all the uniform end–to–end traffic matrices Δ which preserve the network stability, we define the *uniform capacity* $\tilde{C}$ as the maximum offered traffic that can be reached within $\mathcal{U}$, i.e.,

$$\tilde{C} = \max_{\Delta \in \mathcal{U}} G(\Delta) = N(N-1)\zeta.$$ (10)

Due to our assumptions on routing and scatternet structure, the computation of $\tilde{C}$ for the different topologies we consider may be reduced to the evaluation of the maximum traffic rate achievable in a generic master–gateway link. Indeed, it is easy to realize that, with uniform end–to–end traffic connections, the stability conditions (9) are satisfied for each link (i,j) as long as they are verified by all the master–gateway links.

Planar topologies Let us, first, evaluate the uniform capacity for two simple planar structures, namely the star–shaped and the close–loop topologies.

The star topology relies on a single gateway to forward traffic among the M piconets.

The average aggregate traffic offered to the gateway by the nodes of each piconet is equal to $(n-1)(N-n+1)\delta$ packets per slot. Since the gateway spends a fraction $1/M$ of its time in each piconet, getting a fraction $1/(2(n-1))$ of piconet capacity during its permanence time, the limiting capacity (expressed in packet/slot) turns out to be

$$\tilde{C} = \frac{N(N-1)}{2M(n-1)^2(N-n+1)} = \frac{N}{2(n-1)(N-n+1)}\,. \tag{11}$$

It is clear that, in such a situation, the gateway represents a bottleneck that limits the performance of the network. In order to overcome this weakness, a closed–loop structure, as depicted in Fig. 2, can be considered. In this case, routing should be taken into account. We assume that the routing algorithm preserves somehow the symmetry of the structure (i.e., each gateway handles equal load), and that gateways are synchronized, so that a single gateway at a time is active in each piconet. Under these assumptions, each gateway node gets a traffic rate of approximately $\delta N^2/8$ packet/slot from each one of the two piconets it belongs to. Since the gateway gets a fraction $1/(2(n-2))$ of the piconet capacity, the uniform capacity (in packet/slot) is approximately equal to:[8]

$$\tilde{C} = \frac{8N(N-1)}{4(n-2)N^2} = \frac{2(N-1)}{N(n-2)}\,. \tag{12}$$

Solid topologies   In order to achieve better performance, we might add another dimension to the network, so that the resulting structure may be represented by a *solid* structure. In this case, each piconet can be represented as a regular polygon having the gateways as vertices. To have a "fair" structure, the resulting polyhedron has to be regular. To avoid problems with the routing schemes employed, the resulting polyhedron has also to be convex. However, it is a well–known result (which goes back to the ancient Greeks) that there are just five regular convex polyhedra, namely tetrahedron, cube, octahedron, dodecahedron and icosahedron. These are usually referred to as the Platonic solids.

Denoting by F and V the number of facets and vertices of the polyhedron, respectively, and by E the number of edges of the building polygon, we have that the minimal and maximal number of Bluetooth nodes such a structure may handle are equal to:

$$N_{min} = F + V; \tag{13}$$

$$N_{max} = (8 - E) \cdot F + V. \tag{14}$$

We focus on the tetrahedron structure, the analysis for the other cases being similar. We assume perfect synchronization among gateways, so that at each time instant there will be only one gateway present in each piconet. Such a scheduling of the gateways is made possible by the regular structure of the tetrahedron topology. Thus, a gateway will spend one third of its time in each piconet it belongs to. During the permanence in a given piconet, the gateway gets $1/(n-3)$ of the piconet capacity. Let us, now, introduce some notation. Let GW_1, GW_2, GW_3, GW_4 denote the four gateways at the

---

[8] For the sake of conciseness, we omit the derivation of the exact expression of the uniform capacity that, however, can be easily obtained by distinguishing the cases of M even and odd.

vertices of the tetrahedron. Furthermore, let P_i denote the only piconet that GW_i does not belong to, i.e., the piconet on the facet opposite to vertex i. Finally, we will use the notation $\mathcal{X}_i \to \mathcal{Y}_j$ to indicate the aggregate traffic produced by the node(s) in $\mathcal{X}_i$ to the node(s) in $\mathcal{Y}_j$, where $\mathcal{X}_i, \mathcal{Y}_j$ can be piconets or gateways.

Given the symmetry of the network, we can focus on a specific link, for instance the link between the master of P_4 and the gateway GW_1. The traffic along this link is given by the sum of the following terms:

a) $P_4 \to GW_1$, equal to $(n-3)\zeta$;
b) one half of $\{P_2, P_3\} \to P_4$, equal to $(n-3)^2\zeta$;
c) one third of $P_4 \to GW_4$, equal to $(n-3)\zeta/3$;
d) one third of $P_1 \to GW_1$, equal to $(n-3)\zeta/3$;
e) one sixth of $GW_2 \to P_2$ and $GW_3 \to P_3$, equal to $2(n-3)\zeta/6$;
f) one half of $\{GW_2, GW_3\} \to GW_1$, equal to ζ.

Thus, the uniform capacity turns out to be:

$$\tilde{C} = \frac{N(N-1)}{6(n-3)(n-2)^2} \; . \tag{15}$$

5.3   Performance with inter–piconet interference

The mutual interference among overlapping piconets may have a strong effect on scatternet performance. Let $P_s(M)$ be the probability that two stations in piconet A successfully exchange a packet when $M-1$ other piconets are collocated in the scatternet. Therefore, the effective uniform network capacity can be expressed as:

$$C' = \tilde{C} \cdot P_s(M)\,; \tag{16}$$

where $\tilde{C}$ is given by (11), (12) and (15). We evaluate inter–piconet interference in the worst case, i.e., with always–transmitting unsynchronized piconets. Thus, according to [13], we have $P_s(M) = a^{M-1}$, where the parameter a is equal to $a = 0.9852$.

Fig. 5 depicts the uniform capacity C' versus the number of nodes. We have plotted two sets of curves, obtained for $n = 8$ and $n = 3$, respectively. (Clearly, the parameter n has been varied only for the single–gateway and closed–loop configurations, since for the tetrahedral configuration, the number of nodes per piconet is univocally determined by N.) As expected, the single–gateway topology always performs worse than the others, due to the bottleneck introduced by the gateway. The tetrahedron topology achieves the highest capacity, but only for a very small number of nodes N. On the contrary, the closed–loop topology achieves good performance for a wide range of N values, provided that the piconets are "slim", i.e., the number n of nodes per piconet is small.

In Fig. 6, the uniform capacity for single–gateway and closed–loop topologies is depicted for a high number of nodes N. As expected, the overall capacity tends towards zero as N grows, because of the increasing interference.

Given the number of nodes N and the scatternet topology, it may be interesting to evaluate the uniform capacity for different number of piconets M. The left–hand side of Fig. 7 shows the number of piconets M that maximizes the uniform network capacity at
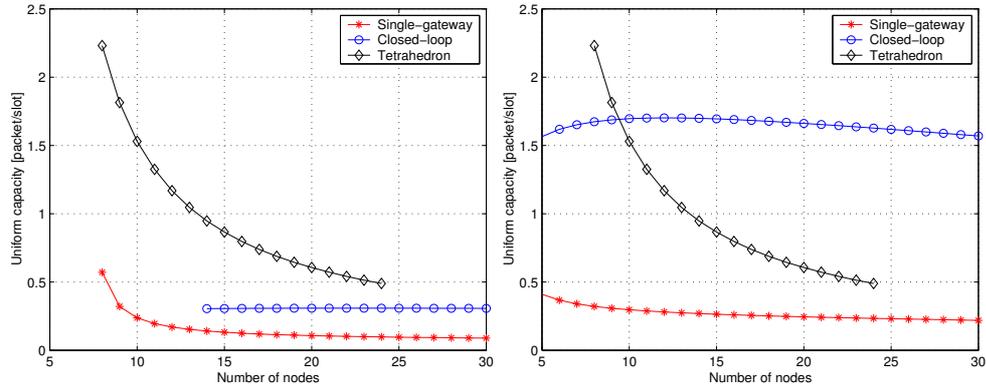
Fig. 5. Uniform capacity C' versus number of nodes N for $n = 8$ (left–hand side) and $n = 3$ (right–hand side).
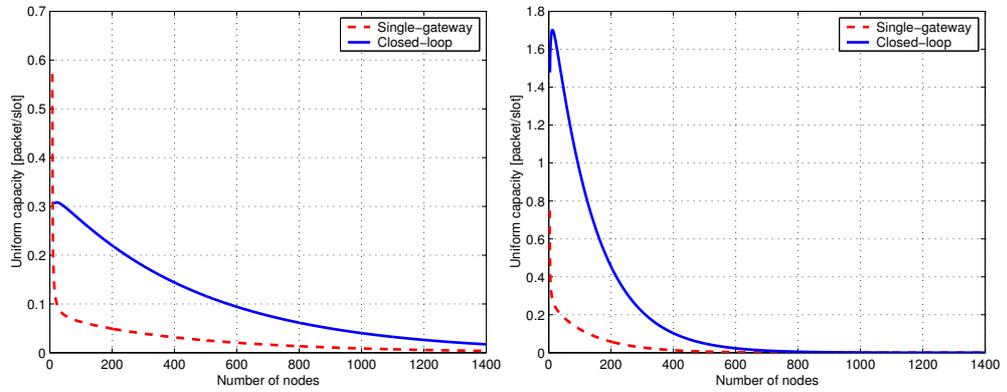

Fig. 6. Uniform capacity C' for high number of nodes N, for $n = 8$ (left–hand side) and $n = 3$ (right–hand side).
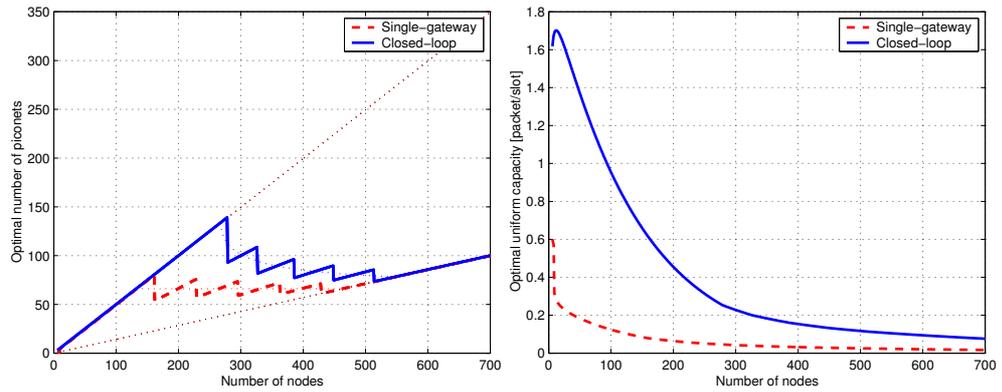

Fig. 7. Optimal number of piconets M and maximum capacity versus the number of nodes N, for the closed–loop and single–gateway topologies.

the varying of the total number N of nodes in the scatternet, for the single–gateway and closed–loop topologies. The right–hand side of Fig. 7 shows the values of the maximum capacity that can be achieved by choosing the optimal value of M for each N. We can note that building up "slim" piconets seems to be preferable, with respect to the overall throughput, for N smaller than ≈ 250. Then, the behavior is rather variable until N is bigger than ≈ 550, after which "fat" piconets, i.e., piconets consisting of $n = 8$ nodes, represents the best choice because of the strong influence of interference. Note that, the "comb"–like trend that can be observed in the left–hand side plot of Fig. 7 for values of N between 250 and 550 is due to the assumption that all the piconets contain the same integer number n of nodes.

6 Conclusions

In this paper we presented an analysis on the impact of the nodes configuration in a Bluetooth scatternet. We discussed how Bluetooth nodes should be organized to build up a scatternet, where the efficiency of the resulting configuration is measured in terms of network capacity. We presented a theoretical study of intrinsic capacity limits of a scatternet, where the maximum throughput may be achieved under local traffic. Hence, we performed a sort of worst–case analysis by considering uniform end–to–end matrix. We showed the impact of mutual interference term on network capacity, and studied how it affects the optimal configuration choice. We thus introduced some regular configurations based upon Platonic solids, which achieve high capacity, even if just for a low number of nodes. As an overall results, closed–loop structure seems to represent a valuable choice, since it scales well with respect to the number of nodes. Some promising directions for future work seem to be, on the one hand, the definition of scatternet formation algorithms which connect nodes building up an efficient network and, on the other hand, the incorporation into such a framework of the constraints due to the spatial dimension of the network, which has not been treated in the present paper.

References

1. Bluetooth special interest group. [Online]. Available: http://www.bluetooth.com
2. J. Karaoguz, "High-rate wireless personal area networks," *IEEE Communications Magazine*, pp. 96–102, Dec. 2001.
3. C. Law, A. K. Mehta, and K.-Y. Siu, "Performance of a new bluetooth scatternet formation protocol," in *Proc. ACM MobiHoc*, Long Beach, CA, 2001.
4. G. V. Zaruba, S. Basagni, and I. Chlamtac, "Bluetrees - scatternet formation to enable bluetooth-based ad hoc networks," in *Proc. IEEE ICC*, Helsinki, 2001.
5. S. Basagni and C. Petrioli, "Multihop scatternet formation for bluetooth networks," in *Proc. IEEE VTC (Spring)*, Birmingham, AL, 2002.
6. G. Tan, A. Miu, J. Guttag, and H. Balakrishnan, "Forming scatternets from bluetooth personal area networks," MIT Laboratory for Computer Science, Tech. Rep. TR-826, 2001.
7. M. Kalia, S. Garg, and R. Shorey, "Scatternet structure and inter-piconet communication in the bluetooth system," in *Proc. IEEE National Conference on Communications*, New Delhi, 2000.
8. S. Baatz, C. Bieschke, M. Frank, P. Martini, C. Scholz, and C. Kuehl, "Building efficient bluetooth scatternet topologies from 1- factor," in *Proc. IASTED WOC*, Banff, Alberta, 2002.
9. P. Bhagwat and S. P. Rao, "On the characterization of bluetooth scatternet topologies," 2001, submitted for publication.
10. A. Capone, M. Gerla, and R. Kapoor, "Efficient polling schemes for bluetooth picocells," in *Proc. IEEE ICC*, Helsinki, 2001.
11. M. Kalia, D. Bansal, and R. Shorey, "Data scheduling and sar for bluetooth mac," in *Proc. IEEE VTC (Spring)*, Tokyo, 2000.
12. R. Kapoor, A. Zanella, and M. Gerla, "A fair and traffic dependent polling scheme for bluetooth," in *Proc. of the joint IEEE conferences ICWLHN 2002 and ICN 2002*, Atlanta, USA, 26-29 Aug. 2002.
13. A. El-Hoiydi, "Interference between bluetooth networks—upper bound on the packet error rate," *IEEE Comm. Lett.*, vol. 5, no. 6, pp. 245–247, Jun 2001.
14. A. Zanella, A. M. Tonello, and S. Pupolin, "On the impact of fading and inter-piconet interference on bluetooth performance," in *Proc. of IEEE WPMC 2002*, Honolulu, Hawaii, Oct. 27-30, 2002.
15. C. C. Foo and K. Chua, "Bluerings - bluetooth scatternets with ring structures," in *Proc. IASTED WOC*, Banff, Canada, 2002.
16. L. Kleinrock, *Queueing Systems*. New York: John Wiley & Sons, 1975.