On the Connectivity of Bluetooth-Based Ad Hoc Networks *

Pilu Crescenzi[1], Carlo Nocentini[1],
Andrea Pietracaprina[2], Geppino Pucci[2], and Carlo Sandri[2]

[1] Dipartimento di Sistemi e Informatica, Università di Firenze, Viale Morgagni 65, 50134, Firenze, Italy. E-mails: {piluc,nocentini}@dsi.unifi.it
[2] Dipartimento di Ingegneria dell'Informazione, Università di Padova, Via Gradenigo 6/B, 35131, Padova, Italy. E-mails: {capri,geppo,sandrica}@dei.unipd.it

Abstract. We study the connectivity properties of a family of random graphs which closely model the Bluetooth's device discovery process, where each device tries to connect to other devices within its visibility range in order to establish reliable communication channels yielding a connected topology. Specifically, we provide both analytical and experimental evidence that when the visibility range of each node (i.e., device) is limited to a vanishing function of n, the total number of nodes in the system, full connectivity can still be achieved with high probability by letting each node connect only to a "small" number of visible neighbors. Our results extend previous studies, where connectivity properties were analyzed only for the case of a constant visibility range, and provide evidence that Bluetooth can indeed be used for establishing large ad hoc networks.

1 Introduction

A critical problem in setting up mobile multi-hop radio networks, also known as *ad hoc networks*, is guaranteeing connectivity while minimizing power consumption and, in some cases, the number of active connections per node. Among others, *Bluetooth* [1] is a popular enabling technology for ad hoc networks, which was originally introduced in 1999 by a Special Interest Group formed by more than 1800 manufacturers for the deployment of Personal Area Networks (PANs), typically consisting of cellular phones, laptops, wireless peripherals and PDAs. Several arguments have been raised to foster the use of Bluetooth for the establishment of large ad hoc networks, due to its low cost, availability, suitability for small devices, and low power consumption (see, for example, [2]). However, a number of challenges arise in this context, particularly for what concerns network formation [3, 4].

Specifically, Bluetooth features a hierarchical organization where the nodes are grouped into *piconets*, with each piconet containing one master and multiple slaves. Piconets are then interconnected through bridge nodes to form a *scatternet*. Scatternet formation can be decomposed into three main steps, namely, *device discovery*, *piconet formation*, and *piconet interconnection*. Each of these steps poses interesting algorithmic challenges for which several solutions have been proposed [1]. In particular, during

* This work was supported in part by MIUR of Italy under project *MAINSTREAM*, and by the EU under the EU/IST Project 15964 *AEOLUS*.

the first step each device attempts at discovering other devices contained within its visibility range and at establishing reliable communication channels with them, in order to form a connected topology, called the *Bluetooth topology*, which underlies the subsequent piconet formation and piconet interconnection steps. Since requiring each device to discover *all* of its neighbors is too time consuming [3], a crucial problem consists of deciding how many neighbors have to be selected in order to guarantee that the resulting Bluetooth topology is connected. Indeed, obtaining connectivity under degree limitations has been indicated in [2] as a major challenge for the adoption of the Bluetooth technology for large ad hoc networks.

In [5] the device discovery step has been effectively modeled as follows. The devices are regarded as a set of n nodes randomly and uniformly distributed in a square of unit side. Each node has a visibility range of $r(n)$, i.e., it can "see" all other nodes within Euclidean distance $r(n)$. Given a function $c(n)$, each node selects as neighbors $c(n)$ visible nodes at random, picking all visible nodes if their number is less than $c(n)$. Observe that the process is unidirectional in nature, however, each link established in this way becomes bidirectional. As a consequence, the final degree of each node may be much higher than $c(n)$, in the case that the node was selected as a neighbor by many other nodes. We refer to $\mathrm{BT}(r(n), c(n))$ as the resulting (undirected) graph.

Previous studies on the connectivity properties of $\mathrm{BT}(r(n), c(n))$ have considered only the case where each node is able to see a constant fraction of all other nodes, that is, the visibility range $r(n)$ is a constant. For this particular case, the experimental analysis conducted in [5] has shown that setting $c(n)$ to a small constant is sufficient to yield connectivity for $\mathrm{BT}(r(n), c(n))$ almost always. The experimental evidence has been later substantiated by the analysis in [6], which shows that, for constant $r(n)$, $c(n) = 2$ is sufficient to achieve connectivity with high probability. Also, in [7] it was proved that constant $c(n)$ (though much larger, in the order of the millions) is also sufficient to guarantee linear expansion of $\mathrm{BT}(r(n), c(n))$. These results suggest that device discovery can be performed efficiently whenever the network is sufficiently small (even though not necessarily a PAN). However, the assumption of constant $r(n)$ becomes quickly unfeasible as the number of devices to be connected increases, which would be the case when adopting Bluetooth for building large ad hoc networks.

In this paper we extend the above studies by providing both analytical and experimental evidence that, when the visibility range is a vanishing function of n, the device discovery step in Bluetooth can still be performed efficiently while guaranteeing connectivity, by letting each device discover only a"small", although non constant, number of neighbors. In particular, we prove that if $r(n) = \Omega(\sqrt{\ln n / n})$, then $\mathrm{BT}(r(n), c(n))$ is connected with high probability as long as $c(n) = \Omega(\ln(1/r(n)))$. We remark that the lower bound on $r(n)$ cannot be improved since it is known that when $r(n) \leq \delta\sqrt{\ln n / n}$, for some constant $0 < \delta < 1$, the *visibility graph* where each node is connected to *all* nodes in its visibility range is disconnected with high probability [8]. A challenging open question is whether the lower bound on the value of $c(n)$ required for connectivity is tight. We give a partial analytical answer to this question by showing that in fact $c(n) = 3$ is sufficient to attain connectivity with high probability, as long as $r(n) \geq n^{-\epsilon}$, for some constant $0 < \epsilon < 1/2$, but each node must choose two of the three neighbors sufficiently close to it.

In the paper we also report on a massive set of experiments conducted in order to assess the real performance of the two previously described protocols. Quite surprisingly, the experiments indicate that, even when the visibility range function is close to the aforementioned lower bound, the number of neighbors needed for connectivity exhibits an extremely weak dependence on $r(n)$: in fact, $c(n) = 3$ suffices almost always, *independently of how the neighbors are chosen*. Moreover, the experiments show that the expected maximum total degree featured by the topologies obtained by choosing three neighbors for each node is much smaller than the one featured by the visibility graph, while the diameter is only slightly larger.

Even though our results were mainly motivated by the question of whether Bluetooth is suitable as a large-scale ad hoc network technology, we believe that they may be of interest for other wireless network scenarios [9].

The rest of the paper is organized as follows. Section 2 analyzes the connectivity of $\mathrm{BT}(r(n), c(n))$ when $c(n)$ is $\Theta(\log(1/r(n)))$. Section 3 analyzes the case of $c(n) = 3$ under further constraints on neighbor selection. Section 4 reports the results of our experiments.

2 Connectivity of $\mathrm{BT}(r(n), c(n))$

Consider a set V of n nodes randomly and uniformly distributed in a unit-side square. Each node $v \in V$ has a visibility range of $r(n)$, i.e., v can "see" all nodes u at Euclidean distance $d(v, u) \leq r(n)$. Let the unit square be tessellated into k^2 square *cells* of side $1/k$, where $k = \lceil \sqrt{5}/r(n) \rceil$. Consequently, any two nodes residing in the same or in adjacent cells are at distance at most $r(n)$: hence, they are visible from one another. Most of our results hold *with high probability* (*w.h.p.* for short) by which we mean that the probability of the stated event is at least $1 - 1/\mathrm{poly}(n)$, where $\mathrm{poly}(n)$ denotes some polynomial function of n.

The following proposition can be easily proved using Chernoff's bound [10].

Proposition 1. *Let $\alpha = 9/10$ and $\beta = 11/10$. There exists a constant $\gamma_1 > 0$ such that for every $r(n) \geq \gamma_1 \sqrt{\ln n/n}$ the following events occur w.h.p.:*

1. *Every cell contains at least $\alpha n/k^2$ and at most $\beta n/k^2$ nodes.*
2. *Every node has at least $(\alpha/4)\pi n r^2(n)$ and at most $\beta \pi n r^2(n)$ nodes in its visibility range.*

The rest of the section is devoted to the proof of the following theorem.

Theorem 1. *There exist two positive real constants γ_1, γ_2 such that, if $r(n) \geq \gamma_1 \sqrt{\ln n/n}$ and $c(n) = \gamma_2 \ln(1/r(n))$ then $\mathrm{BT}(r(n), c(n))$ is connected w.h.p.*

Let $\epsilon = 1/8$. In the proof of the theorem we distinguish between the case $r(n) \leq n^{-\epsilon}$ and the case $r(n) > n^{-\epsilon}$, which are dealt with separately in the following subsections. Moreover, in both cases we condition on the events expressed by Proposition 1, which occur with high probability.

2.1 Case $\gamma_1\sqrt{\ln n/n} \le r(n) \le n^{-\epsilon}$

We fix the lower bound for $r(n)$ to be the same under which Proposition 1 holds. In the range of $r(n)$ considered in this case, we have that $c(n) = \gamma_2 \ln(1/r(n)) = \Theta(\ln n)$. Let Q be an arbitrary cell and let G_Q denote the subgraph of $\mathrm{BT}(r(n), c(n))$ formed by nodes and edges internal to Q. We first show that every G_Q is connected and then prove that for every pair of adjacent cells there exists an edge in $\mathrm{BT}(r(n), c(n))$ whose endpoints are in the two cells.

Lemma 1. *With high probability, every G_Q is connected.*

Proof. Fix an arbitrary cell Q and let A_Q be the event that, for every partition of the nodes in Q into two nonempty subsets, there is at least an edge with endpoints in distinct subsets. Observe that the subgraph $G_Q \subseteq \mathrm{BT}(r(n), c(n))$ is connected if and only if A_Q occurs. Then:

$$\begin{aligned}
&1 - \Pr(A_Q) \le \\
&\le \sum_{s=1}^{\beta n/(2k^2)} \binom{\beta n/k^2}{s} \left(1 - \frac{\alpha n/k^2 - s}{\beta\pi n r^2(n)}\right)^{sc(n)} \left(1 - \frac{s}{\beta\pi n r^2(n)}\right)^{(\alpha n/k^2 - s)c(n)} \\
&\le \sum_{s=1}^{\beta n/(2k^2)} \exp\left(s \ln \frac{e\beta n}{sk^2} - \frac{2sc(n)}{\beta\pi n r^2(n)}\left(\frac{\alpha n}{k^2} - s\right)\right) \\
&\le \frac{1}{n^2}
\end{aligned}$$

where the last inequality holds by choosing the constant γ_2 in the expression for $c(n)$ large enough. The lemma follows by applying the union bound over all k^2 cells. □

Lemma 2. *With high probability, for every pair of adjacent cells Q_1 and Q_2 there is an edge $(u, v) \in \mathrm{BT}(r(n), c(n))$ such that u resides in Q_1 and v resides in Q_2.*

Proof. Consider an arbitrary pair of adjacent cells Q_1 and Q_2 and let B_{Q_1,Q_2} denote the event that there is at least one edge in $\mathrm{BT}(r(n), c(n))$ between the two cells. Since we are conditioning on the events described in Proposition 1, we have that

$$\begin{aligned}
1 - \Pr(B_{Q_1,Q_2}) &\le \left(1 - \frac{\alpha n/k^2}{\beta\pi n r^2(n)}\right)^{2c(n)\alpha n/k^2} \\
&\le \exp\left(-\frac{\alpha n/k^2}{\beta\pi n r^2(n)}(2c(n)\alpha n/k^2)\right) \\
&\le \exp(-\zeta \ln^2 n),
\end{aligned}$$

where ζ is a positive constant. The lemma follows by applying the union bound over all $O(n)$ pairs of adjacent cells. □

For the case $\gamma_1\sqrt{\ln n/n} \le r(n) \le \delta n^{-\epsilon}$, Theorem 1 follows by combining the results of the above two lemmas.

2.2 Case $r(n) > n^{-\epsilon}$

We generalize and simplify the argument used in [6] for the case $r(n) = \Theta(1)$. Specifically, we first show that $\mathrm{BT}(r(n), c(n))$ contains a large connected component C, and then we show that for every node v there is a path from v to C. We condition on the events that the number of nodes in each cell and in the visibility range of each node are within the bounds stated in Proposition 1, which occur with high probability.

Lemma 3. *For $r(n) > \delta n^{-\epsilon}$ and $c(n) \geq 2$, $\mathrm{BT}(r(n), c(n))$ contains a connected component of size $n/(8k^2)$, w.h.p.*

Proof. The argument is identical to the one used in the proof of Proposition 3 in [6]. In particular, the probability of the existence of the connected component is at least

$$1 - \frac{8k^2 \log_2^2 n}{n} - \frac{1}{9^{\log_2 n}},$$

which is $1 - o(n^{-2/3})$ by our choice of k. □

Let C be the connected component of size at least $n/(8k^2)$ which, by the above lemma, exists w.h.p. By the pigeonhole principle there must exist a cell Q containing at least $n/(8k^4)$ nodes of C. Let $V(Q, C)$ the set of nodes residing in Q and belonging to C. We have:

Lemma 4. *With high probability, for each node u there exists a path in $\mathrm{BT}(r(n), c(n))$ from u to some node in $V(Q, C)$.*

Proof. Consider a directed version of $\mathrm{BT}(r(n), c(n))$ where an edge (u, v) is directed from u to v if u selected v during the neighbor selection process. Since we are conditioning on the event stated in the second point of Proposition 1, our choice of $c(n)$ implies that the outdegree of each node is *exactly* $c(n)$ w.h.p. Pick an arbitrary node u and run a sequential breadth-first exploration from u in such a directed version of $\mathrm{BT}(r(n), c(n))$. Stop the exploration as soon as m nodes have been discovered but not yet explored (m is a suitable value that will be chosen later). We say that a failure occurs when the edge (v_1, v_2) is considered during the exploration of v_1, and node v_2 had been previously discovered. It is easy to see that at the moment when m nodes are discovered but not yet explored, if at most $c(n) - 1$ failures have occurred, then the total number of nodes discovered up to that moment is at most $2m$. Also, if at most $c(n) - 1$ failures occur before reaching m unexplored nodes, then at most m nodes have been explored. Therefore, from the second point of Proposition 1 it follows that the probability of not reaching m unexplored nodes with less than $c(n) - 1$ failures is at most

$$\binom{m \cdot c(n)}{c(n)} \left(\frac{2m}{(\alpha/4)\pi n r^2(n)} \right)^{c(n)} \leq \left(\frac{8m^2 e}{\alpha \pi n r^2(n)} \right)^{c(n)} . \tag{1}$$

Now suppose that the above event occurs and consider the m unexplored nodes, say $w_1, w_2, \ldots, w_m$, reached via breadth-first exploration from u. We now estimate the probability that $\mathrm{BT}(r(n), c(n))$ contains a path from w_i to a node in $V(Q, C)$. Observe that from the cell containing w_i there is a sequence of at most $2k$ pairwise adjacent cells ending at Q. Specifically, we estimate the probability that $\mathrm{BT}(r(n), c(n))$ contains a path from w_i to $V(Q, C)$ following such a sequence of cells, with the constraint that the path contains one node per cell and these nodes do not belong to the set of at most $2m$ nodes initially discovered from u or to the $m-1$ paths constructed for any other w_j, with $j \neq i$. This probability is at least $p^{2k}q$, where p is the probability of extending the path one cell further, and q is the probability of ending, in the last step, in a node of $V(Q, C)$. By Proposition 1 we have that

$$p \geq \left(1 - \left(1 - \frac{\alpha n/k^2 - 3m}{\beta \pi n r^2(n)}\right)^{c(n)}\right)$$

$$q \geq \frac{n/(8k^4)}{\beta \pi n r^2(n)} = \frac{1}{8\beta \pi k^4 r^2(n)}.$$

Recall that $c(n) = \gamma_2 \ln(1/r(n)) = \Theta(\ln k)$. If we take $m = o(n/k^2)$ and γ_2 large enough, we have that

$$p^{2k} \geq \tau$$

for some constant $0 < \tau < 1$. It follows that the probability that all of the w_is fail to reach $V(Q, C)$ is at most

$$(1 - \tau q)^m \leq \left(1 - \frac{\tau}{8\beta \pi k^4 r^2(n)}\right)^m = \left(1 - \frac{\tau}{\sigma k^2}\right)^m, \qquad (2)$$

for some positive constant σ.

By combining Equations 1 and 2, we get that the probability that u is not connected to $V(Q, C)$ is at most

$$\left(\frac{8m^2 e}{\alpha \pi n r^2(n)}\right)^{c(n)} + \left(1 - \frac{\tau}{\sigma k^2}\right)^m.$$

Now, since $r(n) > n^{-1/8}$, we have that $k = O(n^{1/8})$. If we choose $m = \Theta(n^{1/3})$ we have that $m = o(n/k^2)$, as required above, and $m = \omega(k^2 \ln n)$. This, combined with the choice of $c(n)$, ensures that the above probability is smaller than $1/n^2$. The lemma follows by applying the union bound over all nodes u. $\square$

For the case $r(n) > n^{-\epsilon}$, Theorem 1 follows by combining the results of the above two lemmas.

3 Achieving $c(n) = 3$ using a double choice protocol

In the previous section we showed that selecting $c(n) = \Theta(\ln(1/r(n)))$ visible neighbors at random is sufficient to enforce global connectivity for all ranges of $r(n)$ which guarantee connectivity of the visibility graph. Whether these many neighbors are necessary remains a challenging open question. As a step towards this objective, we show that, at least for large enough (yet nonconstant) radii, $c(n) = 3$ always suffices under a slightly different neighbor selection protocol where each node is required to direct the selection of some neighbors within a certain geographical region. More formally, consider again the tessellation of the unit square into k^2 square cells of side $1/k$, with $k = \lceil \sqrt{5}/r(n) \rceil$. Define $\mathrm{BT}(r(n), 2, 1)$ to be the undirected graph resulting by letting each node select two neighbors at random among the nodes residing in its cell, and another neighbor at random among all visible nodes.

Observe that if applied in a practical scenario, the above *double-choice* protocol would require each node to infer geographical information about its location and the location of the nodes in its visibility range. For example, this information could be provided by a GPS device.[3]

Theorem 2. *There exists a constant* ϵ*,* $0 < \epsilon < 1/2$ *such that if* $r(n) = \Omega(n^{-\epsilon})$*, then* $\mathrm{BT}(r(n), 2, 1)$ *is connected w.h.p.*

Proof. We employ the same approach used in Subsection 2.1. Specifically, we first argue that w.h.p. for all cells Q, the graph G_Q induced by the nodes in Q is connected, and that for every pair of adjacent cells there is an edge with endpoints in the two cells. Since by the first point of Proposition 1, each cell Q contains $\Omega(n^{1-2\epsilon})$ nodes w.h.p., the main result of [6] implies that two neighbors selected by each node in Q suffice to guarantee connectivity of G_Q with probability at least $1 - 1/n^{\delta(1-2\epsilon)}$, for a suitable positive constant $\delta < 1$. Then, choosing ϵ smaller than $\delta/(2(1+\delta))$ and applying the union bound, all cells will be internally connected with high probability. In order to prove connectivity between adjacent cells, we proceed as in the proof of Lemma 2. In particular, consider an arbitrary pair of adjacent cells Q_1 and Q_2, and let B_{Q_1,Q_2} denote the event that there is at least one edge in $\mathrm{BT}(r(n), 2, 1)$ between the two cells. By conditioning on the events described in Proposition 1, we have that

$$
\begin{aligned}
1 - \Pr(B_{Q_1,Q_2}) &\le \left(1 - \frac{\alpha n/k^2}{\beta \pi n r^2(n)}\right)^{2\alpha n/k^2} \\
&\le \exp\left(-\frac{\alpha n/k^2}{\beta \pi n r^2(n)}(2\alpha n/k^2)\right) \\
&\le \exp(-\zeta n^{1-2\epsilon}),
\end{aligned}
$$

where ζ is a positive constant. The theorem follows by applying the union bound over all $O(n)$ pairs of adjacent cells. □

4 Experiments

We have designed an extensive suite of experiments aimed at comparing the connectivity and other topological properties of the graphs analyzed in the previous sections.[4] In a

[3] A full discussion on the feasibility of this approach is outside the scope of this paper, since the analysis of the double-choice protocol is mostly meant to provide evidence that the selection of very few neighbors may suffice in order to build a connected topology.

[4] The implemented code makes use of the Boost Graph Libraries [11] for computing the number of connected components.

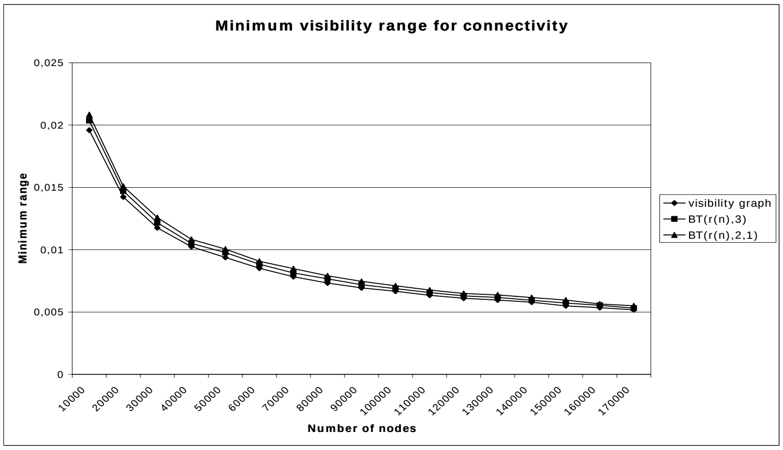

Fig. 1. Comparison of the minimum ranges $r_{\rm lb}$, $r_{\rm sc}$, and $r_{\rm dc}$ yielding connectivity in $\mathrm{BT}(r(n), 3)$, $\mathrm{BT}(r(n), 2, 1)$, and in the visibility graph with $r(n) = r_{\rm lb}$, respectively

first set of experiments, for values of n ranging from 10000 to 170000 with step 10000, we have generated 50 placements of n nodes in the unit square. For each placement, we have determined (through binary search) an approximation to the minimum range $r_{\rm lb}$ which guarantees connectivity of the visibility graph associated with the placement (i.e., the graph where each node connects to all its visible neighbors). Moreover, for the same placement we have determined the minimum range $r_{\rm sc}$ such that the graph $\mathrm{BT}(r_{\rm sc}, 3)$ turns out to be connected in all of 30 repetitions of the neighbor selection protocol. Finally, for the same placement we have determined the minimum radius $r_{\rm dc}$ such that the graph $\mathrm{BT}(r_{\rm dc}, 2, 1)$ turns out to be connected in all of 30 repetitions of the neighbor selection protocol. The results of these experiments are depicted in Figure 1 where for every $10000 \leq n \leq 170000$ the values of $r_{\rm lb}$, $r_{\rm sc}$, and $r_{\rm dc}$, averaged over the 50 placements, are shown. According to the experiments, $r_{\rm sc}$ is very close to $r_{\rm lb}$ (within 5% for all values of n). Moreover, $r_{\rm dc}$ features a similar behavior with a slightly larger value than $r_{\rm sc}$. Observe that, interestingly, connectivity of $\mathrm{BT}(r(n), 2, 1)$ does not seem to require that $r(n) \in \Omega(1/n^{\epsilon})$ as required by the analysis since it is attained for values of $r(n)$ close to $r_{\rm lb}$.

In a second set of experiments we measured the maximum degree of the graphs $\mathrm{BT}(r(n), 3)$ and $\mathrm{BT}(r(n), 2, 1)$, and of the visibility graph with visibility range $r(n)$, where $r(n)$ is chosen to be an approximation of the smallest value which guarantees connectivity in all three cases. The results of these experiments are depicted in Figure 2 where, as before, for each value of n the reported values represent the averages over 50 placements. It can be seen that $\mathrm{BT}(r(n), 2, 1)$ exhibits a slightly smaller maximum degree than $\mathrm{BT}(r(n), 3)$, and, clearly, both graphs have a much smaller maximum degree than the visibility graph whose expected maximum degree can be shown to be $\Theta(\ln n)$, when $r(n) \in \Theta(\sqrt{(\ln n)/n})$ is used.

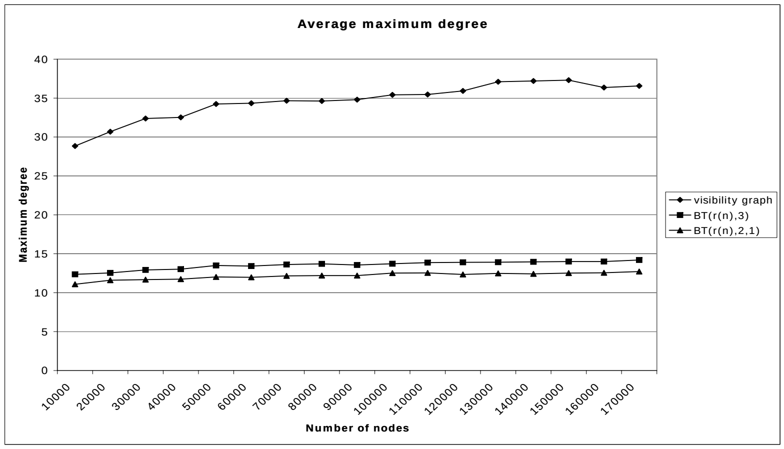

Fig. 2. Comparison of the maximum degree of $\mathrm{BT}(r(n), 3)$, $\mathrm{BT}(r(n), 2, 1)$, and of the visibility graph with range $r(n)$

One last set of experiments concerned the estimation of the average diameter of $\mathrm{BT}(r(n), 3)$ and $\mathrm{BT}(r(n), 2, 1)$, and of the visibility graph with visibility range $r(n)$, where $r(n)$ is chosen to be an approximation of the minimum value which guarantees connectivity in all three cases. The results of these experiments are depicted in Figure 3, once again reporting for each n the averages over 50 placements. It can be seen that $\mathrm{BT}(r(n), 3)$ has a smaller diameter than $\mathrm{BT}(r(n), 2, 1)$, and that it has only a slightly larger diameter than the one of the visibility graph.

5 Conclusions

The main theoretical contribution of this paper is a proof of connectivity for the Bluetooth graph when the visibility range $r(n)$ is a vanishing function of the number n of nodes and each node selects only a logarithmic number of neighbors with respect to $1/r(n)$. Also, we introduced a novel neighbor selection protocol based on a double choice mechanism, which ensures connectivity when a total of only three neighbors are selected by each node. In the paper we also report the results of extensive experiments which validate the theoretical findings. In fact, the experiments suggest that the best avenue for future research is to tighten the analytical result on the connectivity yielded by the single choice protocol, while the double choice idea (which could be more complex to implement in practice) seems only needed for the analysis but does not outperform single choice in practice.

References

1. R.M. Whitaker, L. Hodge, and I. Chlamtac. Bluetooth scatternet formation: a survey. *Ad Hoc Networks*, 3:403–450, 2005.

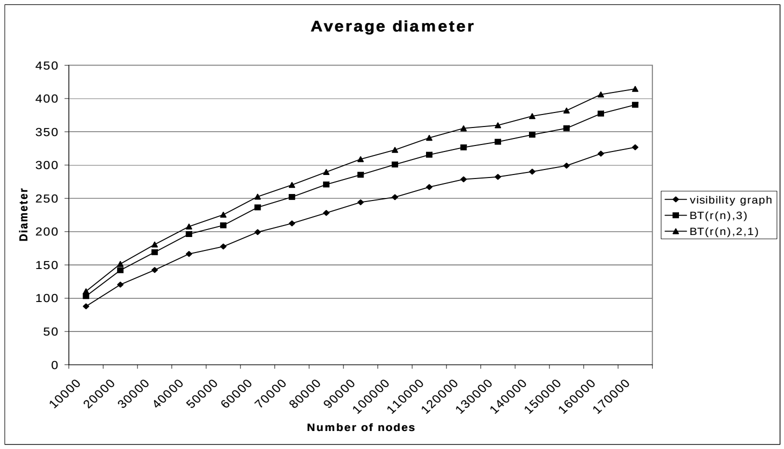

Fig. 3. Comparison of the average diameter of $\mathrm{BT}(r(n), 3)$, $\mathrm{BT}(r(n), 2, 1)$, and of the visibility graph with range $r(n)$

2. I. Stojmenovic and N. Zaguia. Bluetooth scatternet formation in ad hoc wireless networks. In J. Misic and V. Misic, editors, *Performance Modeling and Analysis of Bluetooth Networks*, pages 147–171. Auerbach Publications, 2006.
3. S. Basagni, R. Bruno, G. Mambrini, and C. Petrioli. Comparative performance evaluation of scatternet formation protocols for networks of Bluetooth devices. *Wireless Networks*, 10(2):197–213, 2004.
4. R. Kettimuthu and S. Muthukrishnan. Is Bluetooth suitable for large-scale sensor networks? In *Proc. of the 2005 Intl. Conf. on Wireless Networks*, pages 448–454, 2005.
5. F. Ferraguto, G. Mambrini, A. Panconesi, and C. Petrioli. A new approach to device discovery and scatternet formation in Bluetooth networks. In *Proc. of the 18th International Parallel and Distributed Processing Symposium*, pages 221–228, 2004.
6. D. Dubhashi, C. Johansson, O. Häggstrom, A. Panconesi, and M. Sozio. Irrigating ad hoc networks in constant time. In *Proc. of the 17th ACM Symp. on Parallel Algorithms and Architectures*, pages 106–115, July 2005.
7. A. Panconesi and J. Radhakrishnan. Expansion properties of (secure) wireless networks. In *Proc. of the 16th ACM Symp. on Parallel Algorithms and Architectures*, pages 281–285, 2004.
8. R.B. Ellis, X. Jia, and C. Yan. On random points in the unit disk. *Random Structures and Algorithms*, 29(1):14–25, 2005.
9. I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: a survey. *Computer Networks*, 38:393–422, 2002.
10. T. Hagerup and C. Rüb. A guided tour of Chernoff bounds. *Information Processing Letters*, 33(6):305–308, February 1990.
11. J. G. Siek, L.Q. Lee, and A. Lumsdaine. *Boost Graph Library, The: User Guide and Reference Manual*. Addison Wesley Professional, Reading MA, December 2001.