CONCURRENCY AND COMPUTATION: PRACTICE AND EXPERIENCE
Concurrency Computat.: Pract. Exper. 2000; **00**:1–7 *Prepared using* cpeauth.cls *[Version: 2002/09/19 v2.02]*

On the Connectivity of Bluetooth-Based Ad Hoc Networks†

P. Crescenzi‡, C. Nocentini‡, A. Pietracaprina§, and G. Pucci§,*

‡*Dipartimento di Sistemi e Informatica, Università di Firenze, Italy.*
E-mails: {piluc,nocentini}@dsi.unifi.it

§*Dipartimento di Ingegneria dell'Informazione, Università di Padova, Italy.*
E-mails: {capri,geppo}@dei.unipd.it

SUMMARY

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.

KEY WORDS: Ad Hoc Networks, Bluetooth, Random Graphs, Connectivity

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

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

*Correspondence to: Dipartimento di Ingegneria dell'Informazione, Università di Padova, Via Gradenigo 6/B, 35131, Padova, Italy. Phone: +39 0498277951, Fax: +39 0498277799

Received January 2008
Revised

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 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 (observe that $\mathrm{BT}(r(n), c(n))$ is a generalization of the well-studied random geometric graph [6, 7, 8] which can be obtained by setting $c(n) \geq n - 1$).

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 [9], which shows that, for constant $r(n)$, $c(n) = 2$ is sufficient to achieve connectivity with high probability. Also, in [10] 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 [6, 11]. 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, for values of n up to the hundreds of thousands $c(n) = 3$ 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 are 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 [12].

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, while in Section 5 we conclude with some final considerations and proposals for further research.

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 $d(v, u) \leq r(n) \leq 1$.[†]

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 (i.e., cells sharing a side) 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

[†]In fact, the maximum significant range in the case of the unit square is $\sqrt{2}$. Placing an upper bound of 1 allows us to simplify some of the proofs, however, all of the results in the paper would still hold up to the maximum range.

event is at least $1 - 1/\text{poly}(n)$, where $\text{poly}(n)$ denotes some polynomial function of n. We need the following technical fact.

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 together 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 $\beta \pi n r^2(n)$ other nodes in its visibility range.*

Proof. It is sufficient to show that any given cell contains at least $\alpha n/k^2$ and at most $\beta n/k^2$ nodes with probability greater than or equal to $1 - 1/(2n^2)$, and that any given node has at least $(\alpha/4)\pi n r^2(n)$ and at most $\beta \pi n r^2(n)$ other nodes in its visibility range, with probability greater than or equal to $1 - 1/(2n^2)$. Since there are less than n cells and exactly n nodes, the proposition will follow by a simple application of the union bound. Fix a cell Q. By using Chernoff's bound [13] and the fact $k \leq 4/r(n)$, we obtain that the probability that Q contains more that $\beta n/k^2$ nodes is less than

$$\left(\frac{e^{\beta-1}}{\beta^{\beta}}\right)^{{\gamma_1}^2 \ln n/16}.$$

A symmetrical argument shows that the probability that Q contains less than $\alpha n/k^2$ nodes is less than

$$\left(\frac{e^{1-\alpha}}{(2-\alpha)^{2-\alpha}}\right)^{{\gamma_1}^2 \ln n/16}.$$

By choosing a suitable constant γ_1, both upper bounds can be made smaller than $1/(4n^2)$. Therefore, the probability that the number of nodes in Q is between $\alpha n/k^2$ and $\beta n/k^2$ is at least $1 - 1/(2n^2)$.

The probability bound regarding the number of nodes in the visibility range of any fixed node can be proved in a similar fashion by making the further observation that for $r(n) \leq 1$, the visibility range of any node covers an area of the unit square which is at least $(\pi/4)r^2(n)$ and at most $\pi r^2(n)$. □

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 $\text{BT}(r(n), c(n))$ 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 $\gamma_1\sqrt{\ln n/n} \leq r(n) \leq n^{-\epsilon}$

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:

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

Note that for the values of s in the summation range, we have that $\ln(e\beta n/(sk^2)) = O(\ln n)$ and $((\alpha n/k^2) - s)/(\beta \pi n r^2(n)) = \Theta(1)$, therefore by choosing the constant γ_2 in the expression for $c(n)$ large enough, the summation is dominated by its first term and can be made as small as $1/n^2$. 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

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

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

2.2. Case $n^{-\epsilon} < r(n) \leq 1$

We generalize and simplify the argument which was used in [9] 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. Again, we condition on the events stated in Proposition 1, which occur with high probability.

Lemma 3. *For $n^{-\epsilon} < r(n) \leq 1$ 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 a simple adaptation of the one used in the proof of Proposition 3 in [9], which we highlight in the following for the sake of completeness. Starting from an arbitrary node u, consider a sequential discovery procedure where each node chooses two random neighbors out of those nodes in its visibility range. A simple application of the Chernoff bound [13] shows that each node has at least $n/(2k^2)$ other nodes in its visibility range with probability at least $1 - k^2 e^{-n/(8k^2)}$. This implies that the probability of reaching a previously discovered node in the first $2\log_2 n - 2$ neighbor selections is at most $8k^2 \log_2^2 n/n = O\left(\log^2 n/n^{1-2\epsilon}\right) = O\left(\log^2 n/n^{3/4}\right)$. Hence, by stopping the sequential discovery after $2\log_2 n - 2$ neighbor selections have been executed, we get, with probability $1 - O\left(\log^2 n/n^{3/4}\right)$, a full binary tree rooted at u with $\log_2 n$ leaves. Now, from these leaves we run $\log_2 n$ independent sequential discoveries until a total of $n/(8k^2)$ nodes are discovered. By reasoning as in [9] we argue that this process stochastically dominates a branching process [14] beginning with $\log_2 n$ individuals and binomial offspring distribution with parameters 2 and 3/4. We conclude that the probability of failing to reach $n/(8k^2)$ nodes is bounded from above by the probability that the branching process dies out, which is at most $(1/9)^{\log_2 n}$.

Putting it all together, the probability that the component grown from u has size at least $n/(8k^2)$ is at least

$$1 - O\left(\frac{\log^2 n}{n^{3/4}} + \frac{1}{n^{\log_2 9}}\right),$$

which proves the lemma. □

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 $\mathrm{BT}(r(n), c(n))$. With respect to this exploration, we say that a failure occurs whenever an edge (v_1, v_2) is considered during the exploration of v_1, but node v_2 has been already discovered. Let m be a suitable value, to be fixed later by the analysis. We stop the exploration as soon as one of the following events happen: (a) the $c(n)$-th failure occurs; or (b) m nodes are discovered but not yet explored. We now prove an upper bound on the probability that the exploration stops due to event (a). In this case, it is easy to show that any time before the $c(n)$-th failure occurs, the tree formed by the nodes discovered so far has at most one internal node of degree one, hence the tree contains less than m leaves (i.e., the unexplored nodes) and less than m internal nodes, for a total of less than $2m$ nodes altogether.

From the second point of Proposition 1 it follows that the probability that event (a) happens prior to event (b) is at most

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

where the binomial coefficient bounds from above the number of ways of fixing $c(n)$ failures in the node explorations, which are less than m, while the subsequent factor bounds from above the probability of a fixed configuration of $c(n)$ failures when less than $2m$ nodes have been discovered.

We will choose m so to make the upper bound given in Equation 1 vanishingly small in n. Therefore, we may condition on the event that m unexplored nodes, say $w_1, w_2, \ldots, w_m$, are reached via breadth-first exploration from u before $c(n)$ failures occur. 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 using the bounds in Proposition 1 we have that

$$\begin{aligned} p &\geq \left(1 - \left(1 - \frac{\alpha n/k^2 - 3m}{\beta \pi n r^2(n)} \right)^{c(n)} \right) \\ q &\geq \frac{n/(8k^4)}{\beta \pi n r^2(n)} = \frac{1}{8\beta \pi k^4 r^2(n)}. \end{aligned}$$

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

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

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

$$\left(\frac{2em^2}{(\alpha/4)\pi n r^2(n) - c(n)}\right)^{c(n)} + \left(1 - \frac{\tau}{\sigma k^2}\right)^m .$$

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. □

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 $c(n) = O(\ln(1/r(n)))$ visible neighbors at random is sufficient to enforce global connectivity for all values 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. Such a phenomenon provides evidence that the $O(\ln(1/r(n)))$ bound on $c(n)$ is not likely to be tight.

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, this *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.[‡]

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, with high probability, 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

[‡] 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.

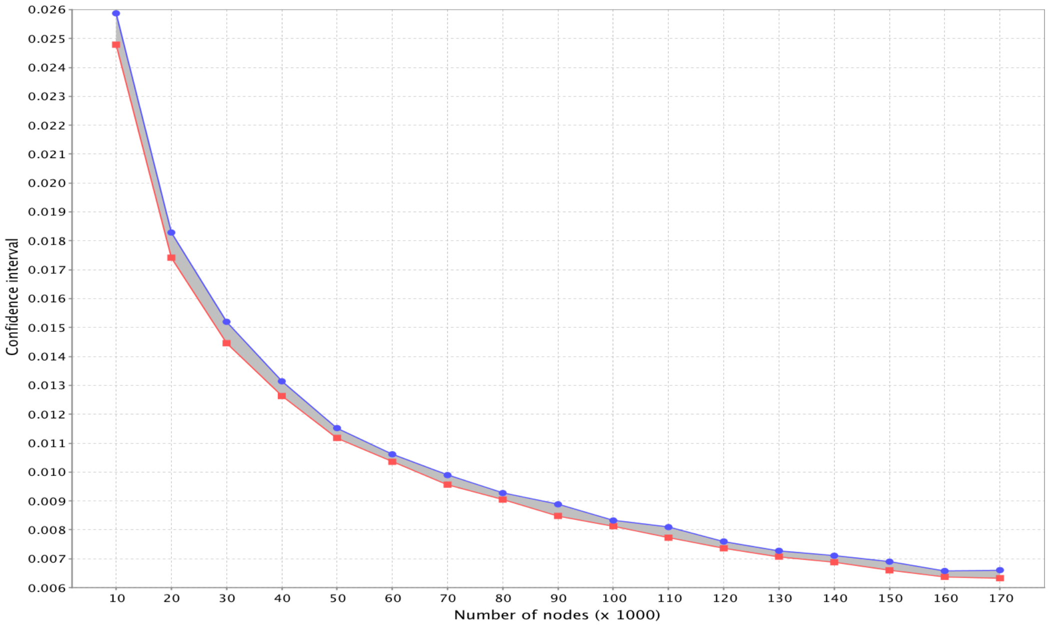

Figure 1. The 95% confidence intervals of the minimum range of the visibility graph

of [9] 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

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

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.§ In a first set of experiments, for values of n ranging from 10000 to 170000 with step 10000, we have performed 20 times the following binary search of the minimum range that guarantees connectivity of the visibility graph associated with the placement (i.e., the graph where each node connects to all its visible neighbors).

1. Set the search interval to $\left[r_{\text{left}} = 0.1\sqrt{\ln/n}, r_{\text{right}} = 0.99\sqrt{\ln/n}\right]$.
   Comment: r_{left} (resp., r_{right}) is a value which guarantees, in practice, that the corresponding visibility graph is always disconnected (resp., connected).
2. If the length of the search interval is less than $0.005 r_{\text{left}}$, then return r_{right}.
3. Generate 50 placements of n nodes in the unit square and verify whether for all of them the range $r = (r_{\text{left}} + r_{\text{right}})/2$ guarantees connectivity of the visibility graph associated with the placement. If this the case, then set $r_{\text{right}} = r$ otherwise set $r_{\text{left}} = r$. Go to step 2.

Figure 1 plots, for each value of n, the lower (r_{lb}) and upper (r_{ub}) endpoints of the 95% confidence intervals yielded by the 20 estimates of the minimum range provided by the above experiments.

We repeated the above procedure to estimate the minimum ranges r_{sc} and r_{dc} which guarantee connectivity of $\text{BT}(r_{\text{sc}}, 3)$ and $\text{BT}(r_{\text{dc}}, 2, 1)$, respectively, so to appreciate whether there is a significant discrepancy between these minimum ranges and the one obtained for the visibility graph. As before, the binary search procedure has been repeated 20 times. At each execution of Step 3, we check whether connectivity is achieved for all of 50 graphs, each obtained by a random placement of n nodes and the neighbor selection protocol, and restrict the search interval accordingly.

Table I reports, for each value of n, the confidence interval $[r_{\text{lb}}; r_{\text{ub}}]$ for the minimum range of the visibility graph, and the values of r_{sc} and r_{dc}, averaged over the 20 estimates. According to these experiments, r_{sc} is very close to r_{lb} (always within 4% for all values of n) and it is almost always within the confidence interval itself (apart from the three starred cases). Also, r_{dc} features a very similar similar behavior. In fact, interestingly, connectivity of $\text{BT}(r(n), 2, 1)$ does not seem to require that $r(n) \in \Omega(\sqrt{1/n^{\epsilon}})$ as implied by the analysis, since it is attained for values of $r(n)$ close to r_{lb}.

In a second set of experiments we measured the maximum degree of the graphs $\text{BT}(r(n), 3)$ and $\text{BT}(r(n), 2, 1)$, 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. In particular, the average maximum degrees have been computed over 50 random placements

§The implemented code makes use of the Boost Graph Libraries [15] for computing the number of connected components and for performing a breadth first search of a graph.

n	$[r_{\text{lb}}; r_{\text{ub}}]$	r_{sc}	r_{dc}
10000	[0.0247868;0.0258708]	0.0251758	0.0253024
20000	[0.0174184;0.0182865]	0.0178253	0.0177519
30000	[0.0144605;0.0151994]	0.0146205	0.0149556
40000	[0.0126296;0.0131409]	0.0129900	0.0126957
50000	[0.0111809;0.0115222]	0.0115894*	0.0112902
60000	[0.0103549;0.0106147]	0.0105789	0.0108026*
70000	[0.0095643;0.0098931]	0.0097572	0.0098735
80000	[0.0090486;0.0092731]	0.0092001	0.0091857
90000	[0.0084843;0.0088834]	0.0087328	0.0087249
100000	[0.0081272;0.0083245]	0.0082948	0.0082109
110000	[0.0077308;0.0080951]	0.0078290	0.0080540
120000	[0.0073631;0.0075985]	0.0076215*	0.0075632
130000	[0.0070734;0.0072776]	0.0072350	0.0072433
140000	[0.0068810;0.0071065]	0.0069794	0.0070066
150000	[0.0066020;0.0069010]	0.0068197	0.0067747
160000	[0.0063759;0.0065772]	0.0066330*	0.0066525*
170000	[0.0063293;0.0066083]	0.0063751	0.0063459

Table I. Comparison between the ranges r_{sc} and r_{dc} which guarantee connectivity for $\text{BT}(r_{\text{sc}}, 3)$ and $\text{BT}(r_{\text{dc}}, 2, 1)$, and the 95% confidence intervals for the minimum range that guarantees connectivity for the visibility graph. Starred values highlight outliers.

for which the values $(r_{\text{lb}} + r_{\text{ub}})/2$, r_{sc}, and r_{dc}, respectively, guaranteed connectivity of the corresponding graphs. The results of these experiments are depicted in Figure 2 for each value of n. It can be seen that $\text{BT}(r(n), 2, 1)$ exhibits a slightly smaller maximum degree than $\text{BT}(r(n), 3)$, and, clearly, both graphs have a much smaller maximum degree than the visibility graph whose expected maximum degree can be analytically shown to be $\Theta(\ln n)$ when $r(n) = \Theta(\sqrt{(\ln n)/n})$.

One last set of experiments concerned the estimation of the average diameter of $\text{BT}(r(n), 3)$ and $\text{BT}(r(n), 2, 1)$, and of the visibility graph with visibility range $r(n)$, where $r(n)$ is chosen as in the previous experiments. In order to avoid expensive all-pairs shortest paths computations, we have chosen to approximate the diameter as twice the maximum height of 30 breadth-first search trees rooted at randomly chosen nodes. The results of these experiments are depicted in Figure 3, once again reporting, for each n the averages over 50 estimates. It can be seen that $\text{BT}(r(n), 3)$ has a diameter which is smaller than the one of $\text{BT}(r(n), 2, 1)$, and not considerably larger 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

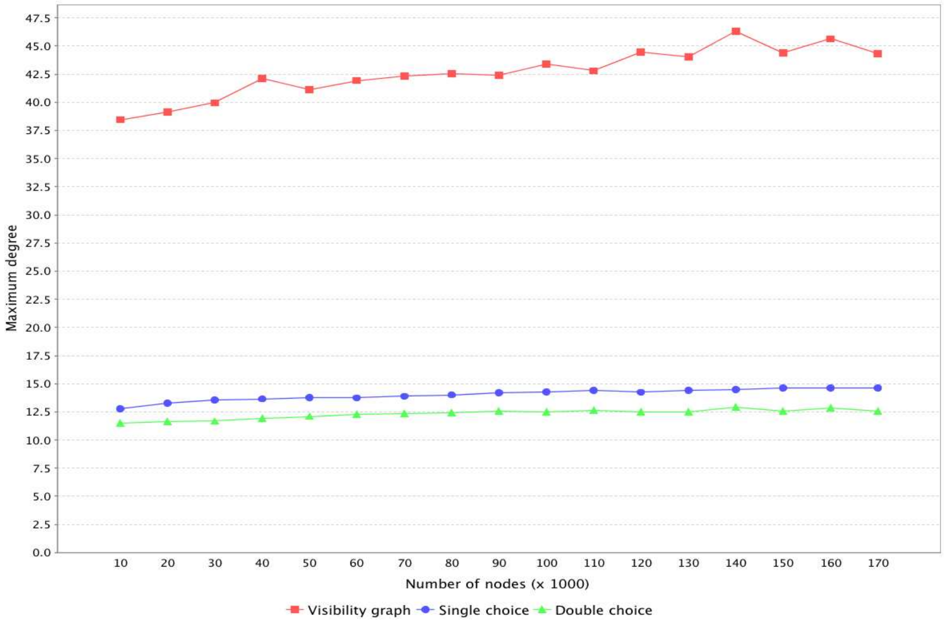

Figure 2. Comparison of the average maximum degree of $\mathrm{BT}(r(n), 3)$, $\mathrm{BT}(r(n), 2, 1)$, and of the visibility graph with range $r(n)$

introduce and analyze 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, from the simulation results, we can conclude that, within the range of 10000 to 170000 nodes, the three protocols seem to be statistically indistinguishable for what concerns connectivity; on the other hand, if the degree of a node is the main concern, then the protocols based on the choice of fewer neighbors offer a substantial advantage, while if the diameter is the main concern, then the visibility graph is slightly superior (recall however that requiring each node to discover all of its visible neighbors may be unfeasible in the practical Bluetooth scenario). From a theoretical point of view, the experimental results substantiate that the best avenue for future research is to tighten the bound on $c(n)$ that guarantees connectivity of $\mathrm{BT}(r(n), c(n))$, since it appears that the double choice protocol (which is in fact of limited practical applicability) does not provide any significant advantage in practice.

REFERENCES

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

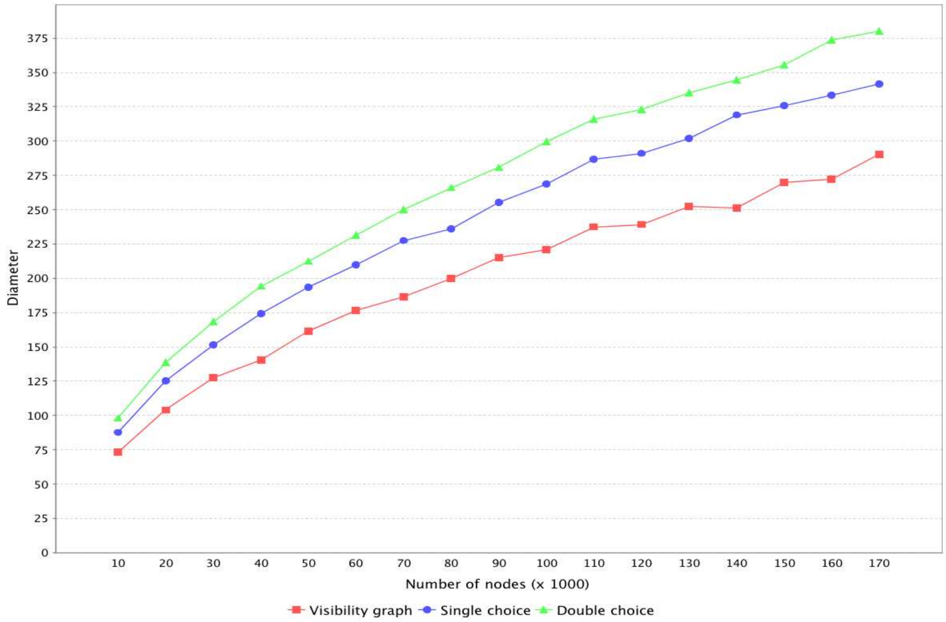

Figure 3. Comparison of the average diameter of $\mathrm{BT}(r(n), 3)$, $\mathrm{BT}(r(n), 2, 1)$, and the visibility graph with range $r(n)$

2. Stojmenovic I, Zaguia N. Bluetooth scatternet formation in ad hoc wireless networks. *Performance Modeling and Analysis of Bluetooth Networks*, Misic J, Misic V (eds.), Auerbach Publications, 2006; 147–171.
3. Basagni S, Bruno R, Mambrini G, Petrioli C. Comparative performance evaluation of scatternet formation protocols for networks of Bluetooth devices. *Wireless Networks* 2004; **10**(2):197–213.
4. Kettimuthu R, Muthukrishnan S. Is Bluetooth suitable for large-scale sensor networks? *Proc. of the 2005 Intl. Conf. on Wireless Networks*, 2005; 448–454.
5. Ferraguto F, Mambrini G, Panconesi A, Petrioli C. A new approach to device discovery and scatternet formation in Bluetooth networks. *Proc. of the 18th International Parallel and Distributed Processing Symposium*, 2004; 221–228.
6. Appel MJP, Russo RP. The connectivity of a graph on uniform points on $[0, 1]^d$. *Statistics & Probability Letters* 2002; **60**(4):351–357.
7. Gupta P, Kumar PR. *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W. H. Fleming*, chap. Critical power for asymptotic connectivity in wireless networks. Birkhäuser, 1999; 547–566.
8. Penrose MD. *Random Geometric Graphs*. Oxford University Press: New York, USA, 2003.
9. Dubhashi D, Johansson C, Häggstrom O, Panconesi A, Sozio M. Irrigating ad hoc networks in constant time. *Proc. of the 17th ACM Symp. on Parallel Algorithms and Architectures*, 2005; 106–115.
10. Panconesi A, Radhakrishnan J. Expansion properties of (secure) wireless networks. *Proc. of the 16th ACM Symp. on Parallel Algorithms and Architectures*, 2004; 281–285.
11. Ellis R, Jia X, Yan C. On random points in the unit disk. *Random Structures and Algorithms* 2005; **29**(1):14–25.
12. Akyildiz I, Su W, Sankarasubramaniam Y, Cayirci E. Wireless sensor networks: a survey. *Computer Networks* 2002; **38**:393–422.

13. Hagerup T, Rüb C. A guided tour of Chernoff bounds. *Information Processing Letters* Feb 1990; **33**(6):305–308.
14. Harris T. *The Theory of Branching Processes*. Springer: Berlin, Germany, 1963.
15. Siek JG, Lee L, Lumsdaine A. *Boost Graph Library, The: User Guide and Reference Manual*. Addison Wesley Professional: Reading MA, 2001.