

On the Expansion and Diameter of Bluetooth-Like Topologies*

Alberto Pettarin, Andrea Pietracaprina, and Geppino Pucci

Department of Information Engineering,
University of Padova, Padova, ITALY
{pettarin,capri,geppo}@dei.unipd.it

Abstract The routing capabilities of an interconnection network are strictly related to its bandwidth and latency characteristics, which are in turn quantifiable through the graph-theoretic concepts of expansion and diameter. This paper studies expansion and diameter of a family of subgraphs of the random geometric graph, which closely model the topology induced by the device discovery phase of Bluetooth-based ad hoc networks. The main feature modeled by any such graph, denoted as $BT(r(n), c(n))$, is the small number $c(n)$ of links that each of the n devices (vertices) may establish with those located within its communication range $r(n)$. First, tight bounds are proved on the expansion of $BT(r(n), c(n))$ for the whole set of functions $r(n)$ and $c(n)$ for which connectivity has been established in previous works. Then, by leveraging on the expansion result, tight (up to a logarithmic additive term) upper and lower bounds on the diameter of $BT(r(n), c(n))$ are derived.

1 Introduction

Random graph models have been employed in the literature for the analytical characterization of topological properties of ad hoc wireless networks governed by a variety of network-formation protocols. One such case concerns networks based on the *Bluetooth* technology [1,2]. A Bluetooth network connects n devices, each endowed with a wireless transmitter/receiver able to communicate within a certain *visibility range*. The network is obtained by means of the following process: 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*. Subsequently, a hierarchical organization is superimposed on this initial topology. Since requiring each device to discover *all* of its neighbors is too time-consuming [3], the device discovery phase is terminated by a suitable time-out, hence only a limited number of neighbors are actually discovered.

The following random graph model for the Bluetooth topology has been proposed in [4] and subsequently generalized in [5]. The devices are represented

* Support for the authors was provided in part by the European Union under the FP6-IST/IP Project AEOLUS.

by n nodes, whose coordinates are randomly chosen within the unit square $[0, 1]^2$; each node selects $c(n)$ neighbors among all visible nodes, that is, among all nodes within Euclidean distance $r(n)$, where $r(n)$ models the visibility range, which is assumed to be the same for all devices. The resulting graph, called $BT(r(n), c(n))$, is the one where there is an undirected edge for each pair of neighbors. Note that such a graph is a subgraph of the well-known random geometric graph in $[0, 1]^2$, where all pairs of visible nodes are connected by an edge [6]. Experimental evidence shows that $BT(r(n), c(n))$ is a good model for the Bluetooth topology [4]. Moreover, $BT(r(n), c(n))$ may be employed as a model for network scenarios where nodes are constrained to maintain a small number of simultaneous connections, because of limited resources, both energetic and computational, or where establishing links to every visible node is, by far, too costly either in time or energy.

Properties of $BT(r(n), c(n))$ have been investigated in a number of recent works. In [7] the authors show that for any fixed constant $r > 0$ there exists a (large) constant c such that $BT(r, c)$ is an expander with high probability. In [8] it is proved that with high probability $BT(r, c)$ is connected for any fixed constant $r > 0$ and $c \geq 2$ whenever n becomes sufficiently large. These results require that the visibility range be a constant, which implies that every node can choose its neighbors among a constant fraction of all of the nodes in the system. Such an assumption becomes unfeasible as the number of devices grows large.

To overcome the latter problem, a more general setting has been analyzed in [5], where it has been proved that $BT(r(n), c(n))$ stays connected, with high probability, also for vanishing values of $r(n)$ (as $n \rightarrow \infty$), as long as each node selects a suitable number of neighbors. Precisely, if $r(n) = \Omega(\sqrt{\log n/n})$, just allowing $c(n) = O(\log(1/r(n)))$ neighbor selections per node ensures that the resulting graph be connected with high probability. The lower bound on $r(n)$ cannot be improved: in fact, when $r(n) \leq \delta\sqrt{\log n/n}$, for some constant $0 < \delta < 1$, the visibility graph obtained connecting every node to all visible ones (i.e., the random geometric graph $RGG(r(n))$ of [6] with radius $r(n)$) is disconnected with high probability [9]. The tightness of the lower bound on $c(n)$ is instead an open problem.

Most of the previous research focuses on the connectivity of the Bluetooth topology, with the exception of the expansion result of [7] which only considers the extreme case of constant visibility range. In this paper, we contribute to a deeper understanding of the Bluetooth topology by providing upper and lower bounds for two crucial structural properties, namely, expansion and diameter, for the values of $r(n)$ and $c(n)$ for which connectivity has been established by previous works. All of our bounds are tight, except for an additive logarithmic term in the upper bound on the diameter. To emphasize the relevance of our results, observe that the bandwidth and latency characteristics of a network, which determine its ability to perform efficient routing, are closely related to the expansion and diameter properties of its underlying topology [10].

The rest of the paper is organized as follows. Section 2 introduces key definitions and properties which will be used throughout the paper. The lower and

upper bounds on the expansion of $BT(r(n), c(n))$ are presented in Section 3, while those on the diameter are obtained in Section 4. Section 5 concludes the paper with some final remarks.

2 Preliminaries

In this section we formally define the Bluetooth topology, illustrate the notation and recall some facts for later use.

Definition 1 (Bluetooth topology). *Given an integer n , a real-valued function $r(n) \rightarrow (0, \sqrt{2}]$ and a positive integer function $c(n)$, the Bluetooth topology, denoted by $BT(r(n), c(n))$, is the undirected random graph $G = (V_n, E_n)$, defined as follows.*

- The vertex set V_n is a set of n points chosen uniformly and independently at random in $[0, 1]^2$.
- The edge set E_n is obtained through the following process: independently, each node selects a random subset of $c(n)$ neighbors among all nodes within distance $r(n)$ (all of them, if they are less than $c(n)$). An edge $\{u, v\} \in E_n$ exists if and only if u has selected v , or viceversa.

In the next sections, we assume the following setting. Consider the standard tessellation of $[0, 1]^2$ into k^2 square cells of side $1/k$ where $k = \lceil \sqrt{5}/r(n) \rceil$. We say that two cells are *adjacent* if they share a side. Thus, any two nodes residing in the same or in two adjacent cells are at distance at most $r(n)$ (i.e., each node is within the range of the other) and we say that they can *see* each other. When the context is clear, with a slight abuse of notation, we identify a cell with the set of nodes residing therein.

Recall that an event occurs *with high probability* (in brief, *w.h.p.*) if its probability is at least $1 - 1/\text{poly}(n)$. Let $m = n/k^2$ be the expected number of nodes residing in a cell. The following proposition will be exploited several times.

Proposition 1 ([5]). *Let $\alpha = 9/10$, $\beta = 11/10$. There exists a constant $\gamma_1 > 0$ such that for every $r(n) \geq \gamma_1 \sqrt{\log n/n}$ the following two events occur w.h.p.:*

1. every cell contains at least αm and at most βm nodes;
2. every node has at least $(\alpha/4)\pi nr^2(n)$ and at most $\beta\pi nr^2(n)$ nodes in its visibility range.

Let $G = (V, E)$ be an undirected graph. Below, we define the quantities at the center of our analysis.

Definition 2 (Neighborhood). *Given a set of vertices $X \subseteq V$, its neighborhood is the set $\Gamma(X) = \{u \in V(G) : \exists e = \{u, v\} \in E(G), v \in X\}$.*

Definition 3 (Expansion). *The expansion of G is a function $\lambda(s)$, for $1 \leq s \leq |V|/2$, such that*

$$\lambda(s) = \min_{S \subseteq V : |S|=s} \frac{|\Gamma(S) - S|}{|S|}.$$

We remark that, in some works, the term “expansion” is used to refer to a global property of the graph, that is, the minimum value of the function $\lambda(s)$ [10]. In contrast, in this paper we offer a finer characterization of the expansion properties of $BT(r(n), c(n))$ by proving explicit bounds on all values of $\lambda(s)$.

Definition 4 (Diameter). *The diameter of G , denoted as $\text{diam}(G)$, is the maximum distance between any two nodes $u, v \in V$, where the distance between two nodes is the number of edges of a shortest path connecting them.*

Observe that, under any reasonable cost model for communication, the maximum latency to be expected of a point-to-point communication in a network is proportional to the diameter of its underlying topology.

In the rest of the paper we focus on $BT(r(n), c(n))$ and we study its expansion and diameter for those ranges of the parameters for which the connectivity is guaranteed by the results of [5], that is, $r(n) \geq \gamma_1 \sqrt{\log n/n}$ and $c(n) = \gamma_2 \log(1/r(n))$ for two suitable positive real constants γ_1, γ_2 .

3 Expansion of $BT(r(n), c(n))$

In this section we study the node expansion of $BT(r(n), c(n))$. Specifically, in Section 3.1 and Section 3.2 we establish, respectively, a lower bound and an upper bound to the node expansion of this family of random graphs. Recall that $m = n/k^2 = \Theta(nr^2(n))$ denotes the expected number of nodes residing in a cell.

3.1 Lower Bound

The main result of this section is the following theorem.

Theorem 1. *Consider an instance of $BT(r(n), c(n))$ with $r(n) \geq \gamma_1 \sqrt{\log n/n}$ and $c(n) = \gamma_2 \log(1/r(n))$, for two suitable positive constants γ_1 and γ_2 . With high probability, for every integer s , $1 \leq s \leq n/2$, we have*

$$\lambda(s) = \begin{cases} \Omega(\min\{c(n), m/s\}) & \text{if } 1 \leq s \leq \alpha m \\ \Omega(\sqrt{m/s}) & \text{if } \alpha m < s \leq n/2. \end{cases}$$

The proof of Theorem 1 relies on three technical lemmas, which characterize the expansion of certain types of node subsets confined within a single cell. Consider a given subset of vertices S of size s . For any cell Q , we call the set $P = S \cap Q$ the *pocket* of S in Q .

Lemma 1. *Let α' and ε' be two suitable positive constants, with $\alpha' \leq \min\{\varepsilon', \alpha\}$. When $|P| \geq \log n$ or $r(n) = O(n^{-1/8})$, then with high probability, for any cell Q and for every pocket $P \subseteq Q$, with $1 \leq |P| \leq \alpha' m$, we have $|\Gamma(P) - P| \geq \varepsilon' \min\{c(n)|P|, m\}$.*

Proof. Fix a cell Q and a size p for P , with $1 \leq p \leq \alpha' m$. We bound the probability that, for every $P \subseteq Q$ its neighborhood is contained in $P \cup T$, where T is a set of nodes not belonging to P with a certain (small) size t . For notational convenience, we abbreviate $c = c(n) = \gamma_2 \log(1/r(n))$ and introduce the following quantities:

- q is the number of nodes in Q ;
- v is the total number of nodes visible by at least one node in Q ;
- w is the minimum number of nodes visible by any node;
- w' is the maximum number of nodes visible by any node;
- z is the minimum number of nodes visible by all nodes in P .

Conditioning on the events of Proposition 1, we have that $q, v, w, w', z = \Theta(m)$.

Let \mathcal{E} be the union, over all the cells Q and all the choices of the pocket $P \subseteq Q$, of the events $|\Gamma(P) - P| \leq t$. We can bound the probability of \mathcal{E} :

$$\begin{aligned} \Pr[\mathcal{E}] &\leq \binom{q}{p} \binom{v}{t} \left(\frac{\binom{t+p}{c}}{\binom{w}{c}} \right)^p \left(\frac{\binom{w'-p}{c}}{\binom{w'}{c}} \right)^{z-(t+p)} \\ &\leq \left(\frac{eq}{p} \right)^p \left(\frac{ev}{t} \right)^t \left(\frac{t+p}{w} \right)^{cp} \left(\frac{w'-p}{w'} \right)^{c(z-(t+p))} \\ &\leq \left(\frac{eq}{p} \right)^p \left(\frac{ev}{t} \right)^t \left(\frac{t+p}{w} \right)^{cp} e^{-\frac{cp}{w'}(z-(t+p))}. \end{aligned}$$

We distinguish between two cases, depending on the value of p .

Case 1: $1 \leq p \leq m/c$. Let $t = \varepsilon' cp$. We rewrite the bound on $\Pr[\mathcal{E}]$ as

$$\Pr[\mathcal{E}] \leq \left(\left(\frac{eqc}{cp} \right)^{1/c} \left(\frac{ev}{\varepsilon' cp} \right)^{\varepsilon'} \left(\frac{\varepsilon' cp}{aw} \right) \right)^{cp},$$

where a is a positive constant, since $p = O(t)$ and $(z - (t + p))/w' = \Theta(1)$. By regrouping the factors, we obtain:

$$\Pr[\mathcal{E}] \leq \left(\frac{c^{1/c} (eq)^{1/c} (ev)^{\varepsilon'}}{a \varepsilon'^{\varepsilon'} w} (cp)^{1-\varepsilon'-1/c} \varepsilon' \right)^{cp} < \frac{1}{n^3},$$

where the last inequality holds for a sufficiently large γ_2 in $c = \gamma_2 \log(1/r(n))$, and for a sufficiently small ε' , since $cp = \Omega(\log n)$. The claim follows by invoking the union bound over the $O(n)$ cells and the $O(n)$ choices of $p = |P|$.

Case 2: $m/c < p \leq \alpha' m$. Note that in this case $cp > m$, whence we set $t = \varepsilon' m$. We rewrite the upper bound on $\Pr[\mathcal{E}]$ as

$$\begin{aligned} \Pr[\mathcal{E}] &\leq \left(\frac{eq}{p} \right)^p \left(\frac{ev}{\varepsilon' m} \right)^{\varepsilon' m} \left(\frac{\varepsilon' m + p}{aw} \right)^{cp} \\ &\leq \left(\left(\frac{eq}{p} \right)^{1/c} \left(\frac{ev}{\varepsilon' m} \right)^{\varepsilon' m/(cp)} \left(\frac{\varepsilon' m + p}{aw} \right) \right)^{cp}. \end{aligned}$$

The first and the second factor of the latter bound are bounded by a constant, for a suitable choice of c and ε' . By our choice of α' , letting ε' be a sufficiently small value, we can make the product of the three factors at most a constant less than 1, so that $\Pr[\mathcal{E}] < \frac{1}{n^3}$ since $cp = \Omega(\log n)$. The claim then follows by applying the union bound as done for Case 1. \square

Lemmas 2 and 3 are proved via counting arguments which are similar in spirit to the one employed in the proof of Lemma 1, and are omitted due to space limitations. Detailed proofs are reported in [11].

Lemma 2. *Let $r(n) = \Omega(n^{-1/8})$, and $c(n) \geq 3$. With high probability, for any cell Q and for every pocket $P \subseteq Q$, with $|P| < \log n$, we have $|\Gamma(P)| > \frac{1}{3}c(n)|P|$.*

Lemma 3. *Let α'' and ε'' be two suitable positive constants, with $\alpha'' \leq \alpha/(2(1+\varepsilon''))$. With high probability, for any pair of distinct adjacent cells Q and Q' and for every pocket $P \subseteq Q$, with $m/c(n) \leq |P| \leq \alpha''m$, we have $|\Gamma(P) \cap Q'| \geq (1+\varepsilon'')|P|$.*

We are now ready to prove the main result of this section.

Proof (Theorem 1). Throughout the proof, we condition on the events stated in Proposition 1 and in the three previous lemmas. We also define $\bar{\alpha} = \min\{\alpha', \alpha''\}$ and $\bar{\varepsilon} = \min\{\varepsilon', \varepsilon'', 1/3\}$ where $\alpha', \alpha'', \varepsilon', \varepsilon''$ are the constants appearing in the statements of Lemma 1 and Lemma 3, respectively. Consider an arbitrary set S of s vertices of $BT(r(n), c(n))$, with $1 \leq s \leq n/2$. We classify the cells according to the size of the pockets of S that they contain: namely, a cell Q such that $Q \cap S \neq \emptyset$ is said to be *black* if it contains at least $\bar{\alpha}m$ nodes of S , and *gray* otherwise. Two cases are possible: either a majority of nodes of S resides in black cells or a majority of nodes of S resides in gray cells.

In the first case, there are at least $\lceil s/(2\beta m) \rceil$ black cells. By well-known topological properties of two-dimensional meshes [12], we have that at least $\Omega(\sqrt{s/m})$ black cells are adjacent to distinct non-black cells. From Lemma 3, every subset $P' \subseteq S$ of size $\bar{\alpha}m$ contained in one of these black cells expands into the corresponding adjacent non-black cell Q' of at least $(1+\bar{\varepsilon})$ times its cardinality, hence $|\Gamma(P) - S| \geq |(\Gamma(P') \cap Q') - S| \geq \bar{\varepsilon}\bar{\alpha}m$, and thus $\lambda(s) = \Omega(\sqrt{m/s})$, which is the correct bound since $s = \Omega(m)$ in this case.

In the second case, we resort to a proof strategy inspired by the one employed in [7]. Referring to the tessellation of $[0, 1]^2$ into k^2 cells, let us index the cells as Q_{ij} , with $1 \leq i, j \leq k$. Define the *sector* \mathcal{S}_{ij} of a cell Q_{ij} as

$$\mathcal{S}_{ij} = \bigcup_{\substack{\max\{i-6,1\} \leq x \leq \min\{i+6,k\} \\ \max\{j-6,1\} \leq y \leq \min\{j+6,k\}}} Q_{xy}.$$

The *active area* \mathcal{A}_{ij} of sector \mathcal{S}_{ij} is defined as

$$\mathcal{A}_{ij} = \bigcup_{\substack{\max\{i-3,1\} \leq x \leq \min\{i+3,k\} \\ \max\{j-3,1\} \leq y \leq \min\{j+3,k\}}} Q_{xy}.$$

Cell Q_{ij} is called the *center* of both sector \mathcal{S}_{ij} and its active area \mathcal{A}_{ij} . Note that the neighborhood of the pocket $P_{ij} = Q_{ij} \cap S$ is entirely contained in \mathcal{A}_{ij} and that the definition of a sector ensures that given two sectors \mathcal{S}_{ij} and $\mathcal{S}_{i'j'}$, with $Q_{i'j'} \cap \mathcal{S}_{ij} = \emptyset$, their active areas are non-overlapping.

Let G be the set of at least $s/2$ nodes of S belonging to gray cells. To estimate the expansion of S , we first execute a greedy procedure, which selects a number of gray cells which are centers of nonoverlapping active areas, and then obtain a lower bound on the expansion by adding up the contributions related to these selected cells. The selection of the centers is obtained via the following marking strategy. Initially all of the gray cells are unmarked. Then, iteratively, the center of the next active area is selected as the unmarked gray cell Q containing the largest pocket of S , and all of the unmarked cells of the sector centered at Q are marked. The procedure terminates as soon as every gray cell becomes marked. The procedure is described by the following pseudocode, where sets I and U maintain, respectively, the indices of the selected centers and the indices of unmarked cells, and subroutine `LARGESTPOCKET(U)` returns the pair (i, j) corresponding to the unmarked cell with the largest pocket (ties broken arbitrarily).

```

CENTER SELECTION
 $I \leftarrow \emptyset; U \leftarrow \{(i, j) : Q_{ij} \text{ is a gray cell}\}$ 
while  $U \neq \emptyset$  do
     $(i, j) \leftarrow \text{LARGESTPOCKET}(U)$ 
     $I \leftarrow I \cup (i, j)$ 
    for each  $Q_{xy} \in \mathcal{S}_{ij}$  do  $U \leftarrow U - (x, y)$ 

```

Let $\langle c_1, c_2, \dots, c_w \rangle$ be the list of w centers picked by `CENTER SELECTION`, where $c_t = (i_t, j_t)$ was chosen at the t -th iteration of the **while** loop. Let $p_t = |P_{c_t}|$, and let g_t be the number of nodes residing in unmarked gray cells of \mathcal{S}_{c_t} at the beginning of iteration t . Clearly, we have that $\sum_{t=1}^w g_t = |G|$ and, by the greedy choice of the centers, $g_t \leq 169p_t$.

In order to lower bound the expansion of S , we proceed as follows. For each t , with $1 \leq t \leq w$, we determine a suitable set of nodes $N_t \subseteq \Gamma(S)$, which belong to gray cells of \mathcal{A}_{c_t} . We distinguish between two different cases. If \mathcal{A}_{c_t} contains a black cell, since Q_{c_t} is gray, there must exist a pair of adjacent black-gray cells in \mathcal{A}_{c_t} , and we pick N_t as a set of $(1 + \bar{\epsilon})\bar{\alpha}m$ nodes in the gray cell reached by nodes of S in the black cell, which exists by virtue of Lemma 3. Otherwise, we let $N_t = \Gamma(P_{c_t}) - P_{c_t}$ and observe that by Lemmas 1 and 2, $|N_t| \geq \bar{\epsilon} \min\{c(n)p_t, m\}$. Note that the N_t 's are all disjoint, but the sum of their sizes does not immediately yield a lower bound to $|\Gamma(S) - S|$, since each N_t may itself contain nodes of S , which have to be subtracted from the overall count.

Let us first consider the special case when no active area \mathcal{A}_{c_t} contains black cells. In this case, the number of *external neighbors* of S (i.e., nodes of $\Gamma(S) - S$) accounted for by the N_t 's is

$$\left(\sum_{t=1}^w |N_t| \right) - |G| = \sum_{t=1}^w (|N_t| - g_t) \geq \sum_{t=1}^w (|N_t| - 169p_t).$$

Since $p_t \leq \bar{\alpha}m$ and $|N_t| \geq \bar{\varepsilon} \min \{c(n)p_t, m\}$, then for a sufficiently large choice of γ_2 in $c(n) = \gamma_2 \log(1/r(n))$ and a sufficiently small value of $\bar{\alpha}$, we have that $|N_t| - 169p_t \geq \mu |N_t|$ for a certain constant $\mu > 0$. Hence,

$$\sum_{t=1}^w (|N_t| - 169p_t) = \Omega \left(\sum_{t=1}^w \bar{\varepsilon} \min \{c(n)p_t, m\} \right) = \Omega(\min \{c(n)s, m\}),$$

and the theorem follows.

Consider now the general case where some active areas contain black cells, which implies that $s = \Omega(m)$. Observe that $\sum_{t=1}^w |N_t| = \Omega(|G|) = \Omega(s)$, and note that it is sufficient to show that the number of external neighbors of S is $\Omega(\sum_{t=1}^w |N_t|)$. Partition the index set $I = \{1, 2, \dots, t\}$ into two disjoint subsets B_1 and B_2 , such that $t \in B_1$ if \mathcal{A}_{c_t} contains no black cells, and $t \in B_2$ otherwise. Suppose that $\sum_{t \in B_2} |N_t| \geq \tau \sum_{t \in B_1} |N_t|$, for a suitable positive constant τ which will be specified later. For each $t \in B_2$ the set N_t contains $(1 + \bar{\varepsilon})\bar{\alpha}m$ nodes, and at least $\bar{\varepsilon}\bar{\alpha}m$ of these nodes are external neighbors of S . Hence, the total number of external neighbors of S is at least

$$\sum_{t \in B_2} \bar{\varepsilon}\bar{\alpha}m = \frac{\bar{\varepsilon}}{1 + \bar{\varepsilon}} \sum_{t \in B_2} |N_t| \geq \frac{\bar{\varepsilon}}{1 + \bar{\varepsilon}} \frac{\tau}{1 + \tau} \sum_{t=1}^w |N_t|,$$

and the theorem follows. Finally, if $\sum_{t \in B_2} |N_t| < \tau \sum_{t \in B_1} |N_t|$, the number of external neighbors of S accounted for by the nodes in the N_t 's is

$$\begin{aligned} \left(\sum_{t=1}^w |N_t| \right) - |G| &= \sum_{t \in B_1} (|N_t| - 169p_t) + \sum_{t \in B_2} (|N_t| - 169p_t) \\ &\geq \sum_{t \in B_1} \mu |N_t| + \sum_{t \in B_2} ((1 + \bar{\varepsilon})\bar{\alpha}m - 169\bar{\alpha}m) \\ &> \sum_{t \in B_1} \mu |N_t| - \sum_{t \in B_1} \left(\frac{169}{1 + \bar{\varepsilon}} - 1 \right) \tau |N_t|. \end{aligned}$$

By fixing τ such that $((169/(1 + \bar{\varepsilon})) - 1)\tau = \mu/2$, we get

$$\sum_{t \in B_1} \mu |N_t| - \sum_{t \in B_1} \left(\frac{169}{1 + \bar{\varepsilon}} - 1 \right) \tau |N_t| = \frac{\mu}{2} \sum_{t \in B_1} |N_t| = \Omega \left(\sum_{t=1}^w |N_t| \right),$$

and the theorem follows. \square

3.2 Upper Bound

We now prove that the lower bound of Theorem 1 is asymptotically tight.

Theorem 2. *Consider an instance of BT $(r(n), c(n))$ with $r(n) \geq \gamma_1 \sqrt{\log n/n}$ and $c(n) = \gamma_2 \log(1/r(n))$, for two suitable positive constants γ_1 and γ_2 . With*

high probability, for every integer s , $1 \leq s \leq n/2$, there exists a set of vertices S of size s whose expansion is

$$\lambda(s) = \begin{cases} O(\min\{c(n), m/s\}) & \text{if } 1 \leq s \leq \alpha m \\ O(\sqrt{m/s}) & \text{if } \alpha m < s \leq n/2. \end{cases}$$

Proof. Suppose that the events stated in Proposition 1 occur. If $s \leq \alpha m$, we can choose any subset S of the nodes in a single corner cell Q , so that a total of at most $13\beta m$ nodes are visible from S . Hence, $\lambda(s) = O(m/s)$. Consider a list $\langle v_1, v_2, \dots, v_n \rangle$ of the vertices of V , sorted by nondecreasing node degree. If we take $S = \{v_1, v_2, \dots, v_s\}$, we are guaranteed that the sum of the degrees of all nodes in S is not greater than $2c(n)s$, or otherwise the sum of the degrees of the n nodes would exceed $2c(n)n$, which is impossible. Combining the two cases above yields the thesis for the case $s \leq \alpha m$. Let $s > \alpha m$ and consider a set S which occupies an approximately square area of $\Theta(s/m)$ cells in a corner of $[0, 1]^2$. Since only the nodes in $O(\sqrt{s/m})$ cells are visible from S , we have that $\lambda(s) = O(\sqrt{ms}/s) = O(\sqrt{m/s})$, and the theorem follows. \square

We remark that the tight bounds on the expansion of $BT(r(n), c(n))$ provided by Theorems 1 and 2 extend the results in [7] from the specific case of $r(n) = \Theta(1)$ to any value of $r(n)$ which guarantees the connectivity of the graph. Note also that if we consider the minimum expansion $\lambda = \min_{1 \leq s \leq n/2} \lambda(s)$, we obtain that for the Bluetooth topology $\lambda = \Theta(r(n))$.

Similar techniques may be applied to prove that $RGG(r(n))$ features an expansion of $\lambda(s) = \Theta(m/s)$ for $1 \leq s \leq \alpha m$, and $\lambda(s) = \Theta(\sqrt{m/s})$ for $\alpha m < s \leq n/2$ (details are given in [11]). Hence, quite surprisingly, the expansion of $BT(r(n), c(n))$ is, within a constant factor, equal to the expansion of $RGG(r(n))$ whenever $s = \Omega(m/c(n))$.

4 Diameter of $BT(r(n), c(n))$

In this section, we provide upper and lower bounds on the diameter of $BT(r(n), c(n))$ by leveraging on the expansion result of Section 3. Specifically, the upper bound relies on the following lemma, which relates diameter and expansion.

Lemma 4. *Given a connected undirected graph $G = (V, E)$ with n nodes and expansion $\lambda(s)$, for $1 \leq s \leq n/2$, consider the following recurrence:*

$$\begin{aligned} N_0 &= 1 \\ N_i &= (1 + \lambda(N_{i-1})) N_{i-1}. \end{aligned} \tag{1}$$

Define i^* as the smallest index such that $N_{i^*} > n/2$. Then, $\text{diam}(G) \leq 2i^*$.

Proof. Let $d = \text{diam}(G)$ and let u and v be two nodes at distance d in G . Consider a breadth-first tree rooted at u . For $0 \leq i \leq d$, let W_i denote the set of nodes at level i in the tree, and $Y_i = \bigcup_{\ell=0}^i W_\ell$. Note that the expansion properties of G imply that $|Y_i| \geq N_i$. Define now j^* as the smallest index such that $|Y_{j^*}| > n/2$, which implies that $j^* \leq i^*$. Also, w.l.o.g., we can assume that $j^* \geq \lceil d/2 \rceil$, or otherwise we repeat the argument considering the breadth-first tree rooted at v . Indeed, since u and v are at distance d , one of the two breadth-first trees must reach at most $n/2$ nodes within the first $\lceil d/2 \rceil - 1$ levels, or there would be a path shorter than d connecting u and v . The lemma follows. \square

Theorem 3. *Consider an instance of $BT(r(n), c(n))$ with $r(n) \geq \gamma_1 \sqrt{\log n/n}$ and $c(n) = \gamma_2 \log(1/r(n))$, for two suitable positive constants γ_1 and γ_2 . With high probability,*

$$\text{diam}(BT(r(n), c(n))) = O\left(\frac{1}{r(n)} + \log n\right).$$

Proof. We apply Lemma 4 by estimating the value i^* for the graph $BT(r(n), c(n))$, conditioning on the fact that the expansion of $BT(r(n), c(n))$ is $\lambda(s) = \Omega(\min\{c(n), m/s\})$ for $s \leq \alpha m$, and $\lambda(s) = \Omega(\sqrt{m/s})$ for $s > \alpha m$, which happens w.h.p. (see Theorem 1).

In order to account for these two different expansion regimes, we proceed as follows. Let $K(j) = \min\{i : N_i \geq 2^j\}$, so that $i^* = K(\log n - 1)$ and let j_1 be such that $2^{j_1} = \Theta(m)$. Since $\lambda(N_i) = \Omega(1)$ for $0 \leq i < K(j_1)$, it follows that $K(j_1) = O(\log n)$. Observe that for $i > K(j_1)$, there exists a constant σ such that $\lambda(N_i) \geq \sigma \sqrt{m/N_i}$. As a consequence, for $j > j_1$, we have:

$$\begin{aligned} N_{K(j)} &\geq N_{K(j-1)} \prod_{s=K(j-1)}^{K(j)-1} \left(1 + \frac{\sigma \sqrt{m}}{\sqrt{N_s}}\right) \\ &\geq N_{K(j-1)} \left(1 + \frac{\sigma \sqrt{m}}{\sqrt{N_{K(j-1)}}}\right)^{K(j)-K(j-1)} \\ &\geq 2^{j-1} \left(1 + \frac{\sigma \sqrt{m}}{2^{j/2}}\right)^{K(j)-K(j-1)}. \end{aligned}$$

Since $K(j)$ is defined as the smallest index for which $N_{K(j)} \geq 2^j$, from the above inequalities it follows that $K(j) - K(j-1) = O\left(\frac{2^{j/2}}{r(n)\sqrt{n}}\right)$. Therefore,

$$\begin{aligned} i^* = K(\log n - 1) &= \sum_{j=1}^{\log n - 1} (K(j) - K(j-1)) \\ &= \sum_{j=1}^{j_1} (K(j) - K(j-1)) + \sum_{j=j_1+1}^{\log n - 1} (K(j) - K(j-1)) \\ &= O(\log n) + O\left(\frac{1}{r(n)}\right), \end{aligned}$$

and the theorem follows from Lemma 4. \square

We now show that Theorem 3 gives a tight estimate for the diameter of $BT(r(n), c(n))$ when $r(n) = O(1/\log n)$.

Theorem 4. *Consider an instance of $BT(r(n), c(n))$ with $r(n) \geq \gamma_1 \sqrt{\log n/n}$ and $c(n) \geq \gamma_2 \log(1/r(n))$, for two suitable positive constants γ_1 and γ_2 . With high probability,*

$$\text{diam}(BT(r(n), c(n))) = \Omega\left(\frac{1}{r(n)}\right).$$

Proof. Consider the natural tessellation introduced in Section 2. By Proposition 1, with high probability the top leftmost cell and the bottom rightmost cell contain at least one node each, hence the Euclidean distance between these two nodes is $\Theta(1)$. Therefore, any path in $BT(r(n), c(n))$ connecting them must contain at least $\Omega(1/r(n))$ nodes. \square

We point out that the lower bound for the case $r(n) = \Theta(1)$ can be improved to $\Omega(\log n / \log \log n)$. (Full details will be given in the full version of this extended abstract.)

5 Conclusions

The main result of this paper is a tight characterization of the node expansion properties of the Bluetooth topology. Since expansion is essentially a measure of bandwidth, being able to provide a quantitative estimate of this property is useful for the design and analysis of routing strategies [10]. Our result is valid for the entire set of visibility ranges $r(n)$ and number of neighbor choices $c(n)$ which are known to produce a connected graph, as opposed to the result of [7] which holds only for the extreme case $r(n) = \Theta(1)$.

By leveraging on the expansion properties, we also derive nearly tight bounds on the diameter of the same topology, which is again an important measure for routing, related to the latency of the network. Our bounds are tight for a large spectrum of visibility ranges (i.e., $r(n) = O(1/\log n)$), which includes “small ranges”, that is, those which are most interesting for the large scale deployment of the technology. In fact, for the larger ranges $r(n) = \omega(1/\log n)$ the upper and lower bounds differ by a mere additive logarithmic term. Closing this gap is still an open problem.

A consequence of our results is that for subsets of $s = \Omega(m/c(n))$ vertices, $BT(r(n), c(n))$ exhibits roughly the same expansion as the much denser random geometric graph $RGG(r(n))$ of [6]. Also, the diameters of the two graphs differ by at most a logarithmic additive term. These are important considerations for real ad hoc networks, especially for what concerns routing capabilities, since they imply that $BT(r(n), c(n))$ features similar bandwidth and latency characteristics of $RGG(r(n))$ at only a fraction of the costs.

Finally, we recall that it is still an open problem to establish, for every given visibility range $r(n) = \Omega(\sqrt{\log n/n})$, the *minimum* number $c(n)$ of neighbor

choices which yield connectivity and to assess the corresponding diameter and expansion properties.

References

1. Whitaker, R., Hodge, L., Chlamtac, I.: Bluetooth scatternet formation: a survey. *Ad Hoc Networks* **3** (2005) 403–450
2. Stojmenovic, I., Zaguia, N.: Bluetooth scatternet formation in ad hoc wireless networks. In Misić, J., Misić, V., eds.: *Performance Modeling and Analysis of Bluetooth Networks*, 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* **10**(2) (2004) 197–213
4. Ferraguto, F., Mambrini, G., Panconesi, A., Petrioli, C.: A new approach to device discovery and scatternet formation in Bluetooth networks. In: *Proc. of the 18th Int. Parallel and Distributed Processing Symposium*. (2004)
5. Crescenzi, P., Nocentini, C., Pietracaprina, A., Pucci, G.: On the connectivity of Bluetooth-based ad hoc networks. *Concurrency and Computation: Practice and Experience* **21**(7) (2009) 875–887
6. Penrose, M.: *Random Geometric Graphs*. Oxford University Press (2003)
7. Panconesi, A., Radhakrishnan, J.: Expansion properties of (secure) wireless networks. In: *Proc. of the 16th ACM Symp. on Parallelism in Algorithms and Architectures*. (2004) 281–285
8. Dubhashi, D., Johansson, C., Häggström, O., Panconesi, A., Sozio, M.: Irrigating ad hoc networks in constant time. In: *Proc. of the 17th ACM Symp. on Parallelism in Algorithms and Architectures*. (2005) 106–115
9. Ellis, R., Jia, X., Yan, C.: On random points in the unit disk. *Random Structures and Algorithms* **29**(1) (2005) 14–25
10. Leighton, T., Rao, S.: Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM* **46**(6) (1999) 787–832
11. Pettarin, A., Pietracaprina, A., Pucci, G.: On the expansion and diameter of Bluetooth-like topologies. Technical report available at <http://www.dei.unipd.it/~pettarin/pubbl/manu/PettarinPP09.pdf>.
12. Bilardi, G., Preparata, F.P.: Area-time lower-bound techniques with applications to sorting. *Algorithmica* **1**(1) (1986) 65–91