Noname manuscript No.
(will be inserted by the editor)

On the Expansion and Diameter of Bluetooth-like Topologies

Alberto Pettarin · Andrea Pietracaprina · Geppino Pucci

Received: date / Accepted: date

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 $BT\left(r(n), c(n)\right)$, 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 $BT\left(r(n), c(n)\right)$ 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, nearly-tight upper and lower bounds on the diameter of $BT\left(r(n), c(n)\right)$ are derived. In particular, we show asymptotically tight bounds on the diameter when the communication range is near the minimum needed for connectivity, the typical scenario considered in practical applications.

Keywords Wireless Networks · Expansion · Diameter · Random Graphs · Bluetooth Protocol

---

An extended abstract of this work [1] was presented at the *17th European Symposium on Algorithms, ESA 2009.*

A. Pettarin
Department of Information Engineering, University of Padova, Padova, ITALY
E-mail: pettarin@dei.unipd.it

A. Pietracaprina
Department of Information Engineering, University of Padova, Padova, ITALY
E-mail: capri@dei.unipd.it

G. Pucci
Department of Information Engineering, University of Padova, Padova, ITALY
E-mail: geppo@dei.unipd.it

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 [2,3]. 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 to discover other devices contained within its visibility range and to establish 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 [4], 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 [5] and subsequently generalized in [6]. The devices are represented 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 the 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 $BT\left(r(n), c(n)\right)$, 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 [7] in two dimensions. Experimental evidence shows that $BT\left(r(n), c(n)\right)$ is a good model for the Bluetooth topology [5]. Moreover, we remark that the $BT\left(r(n), c(n)\right)$ graph may be employed as a model for other real ad hoc 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 too costly either in time or energy.

Properties of $BT\left(r(n), c(n)\right)$ have been investigated in a number of recent works. In [8] 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 [9] 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. It must be stressed that 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 rapidly unfeasible as the number of devices grows large.

To overcome the latter problem, a more general setting has been analyzed in [6], where it has been proved that $BT\left(r(n), c(n)\right)$ stays connected, with high probability, also for vanishing values of $r(n)$ (as $n \to \infty$), as long as each node selects a suitable number of neighbors. Precisely, if $r(n) = \Omega\left(\sqrt{\log n / n}\right)$, just allowing $c(n) = O\left(\log\left(1/r(n)\right)\right)$ neighbor selections per node ensures the connectivity of the resulting graph with high probability. The lower bound on

$r(n)$ cannot be asymptotically 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 [7] with radius $r(n)$) is disconnected with high probability [10].

A different model to represent the device discovery phase of Bluetooth networks, dubbed Blue Pleiades, has been proposed in [11]. In this model nodes attempt to establish connections with up to $c = \Theta\,(1)$ devices in a number of asynchronous rounds, where each connection might fail if the polled node has already reached a maximum allowed degree $c^\star$. In this framework the transmission radius is still a constant $0 < r < \sqrt{2}$, and the authors show that, for a suitable constant $c^\star$, connectivity is guaranteed with high probability. However, the neighbor selection protocol requires $\Theta\,(\log n)$ rounds to complete.

Very recently Broutin et al. [12] have studied the minimum number of neighbor choices needed to achieve a connected Bluetooth Topology when the transmission radius matches the connectivity threshold for the corresponding $RGG\,(r(n))$, that is, for $r(n) = \Theta\left(\sqrt{\log n/n}\right)$. They show that $c(n) = \Theta\left(\sqrt{\log n/\log\log n}\right)$ choices yield a connected graph with high probability, and that this threshold is sharp. However, to the best of our knowledge, given an arbitrary transmission radius $r(n) = \omega\left(\sqrt{\log n/n}\right)$, establishing the minimum $c(n)$ such that the resulting $BT\,(r(n), c(n))$ is connected is still an open problem. Additionally, for $c(n) = \Theta\left(\sqrt{\log n}\right)$, the authors of [12] prove that the diameter of the resulting Bluetooth Topology is $\Theta\,(1/r(n))$.

It has to be remarked that most of the aforementioned works concentrate on studying the connectivity of the Bluetooth topology, with the exception of the expansion result of [8] which only considers the extreme case of constant visibility range and the result on the diameter of [12] which only applies to $r(n)$ at the connectivity threshold. 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 [13].

The rest of the paper is organized as follows. Section 2 introduces a number of 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. Section 4 provides the analogous bounds on the expansion of $RGG\,(r(n))$, thus showing that the Bluetooth topology, despite of being a (possibly very sparse) subgraph of the random geometric graph with the same visibility radius, exhibits roughly the same expansion properties. Through a general technique of independent interest which leverages on the expansion result, Section 5 characterizes the diameter of $BT\,(r(n), c(n))$, while Section 6 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 a positive integer n, a real-valued function $r(n) \to (0, \sqrt{2}]$ 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 node 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 vice versa.

We say that two nodes *see each other* if they are within distance $r(n)$. 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 $k = \left\lceil \sqrt{5}/r(n) \right\rceil$. Consequently, any two nodes residing in the same or in two adjacent cells are at distance at most $r(n)$, hence they 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* if its probability is at least $1 - 1/\operatorname{poly}(n)$. Let $m = n/k^2 = \Theta\left(nr^2(n)\right)$ be the expected number of nodes residing in a cell. The following proposition will be used several times throughout the paper.

Proposition 1 ([6]) *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 with high probability:*

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, connected graph. Below, we define the quantities at the core 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 $\lambda(s)$ taken over all subset sizes $1 \leq s \leq |V|/2$ (see [13]). In contrast, we offer a finer characterization of the expansion properties of $BT\,(r(n), c(n))$ by proving explicit bounds on $\lambda(s)$ for *all* values of s.

Definition 4 (Diameter) The *diameter* of G, denoted as $\mathrm{diam}(G)$, 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 [6], 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 and γ_2.

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

In this section we study the expansion of $BT\,(r(n), c(n))$. Specifically, in Section 3.1 we establish a lower bound on the expansion of this family of random graphs. As an application of the latter result, in Section 3.2 we prove an upper bound to the flooding time of a message in a dynamic system closely related to the Bluetooth topology. Finally, Section 3.3 provides an upper bound on the expansion of $BT\,(r(n), c(n))$, matching the above lower bound.

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*

$$\lambda(s) = \begin{cases} \Omega\left(\min\{c(n), m/s\}\right) & \text{if} \quad s = O(m)\,, \\ \Omega\left(\sqrt{m/s}\right) & \text{if} \quad s = \Omega(m)\,, \end{cases}$$

where m is the expected number of nodes in a cell.

Observe that, by setting $s = n/2$ in Theorem 1, when $r(n) = \Theta(1)$ (hence, $c(n) = \Theta(1)$), we obtain as a by-product the result in [8].

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 shows that a pocket is highly expanding either when the visibility range is $r(n) = O\left(n^{-\delta}\right)$ for some constant $\delta > 0$, or when we consider a sufficiently large pocket. Lemma 2 covers the case in which the visibility range is large but the pocket is of at most logarithmic size. Finally, for all radii, Lemma 3 assures that a pocket P containing a sufficiently large constant fraction of the nodes of its cell Q, expands roughly linearly into any adjacent cell Q' or in Q itself.

Lemma 1 *Let α' and ε' be two suitable positive constants, with $\alpha' \leq \min\{\varepsilon', 1/2\}$. Then, with high probability, for any cell Q and for every pocket $P \subseteq Q$ such that $c(n)\,|P| = \Omega(\log n)$ and $|P| \leq \alpha' m$, we have $|\Gamma(P) - P| \geq \varepsilon' \min\{c(n)\,|P|, m\}$, where m is the expected number of nodes in a cell.*

Proof Fix a cell Q and a pocket P whose size $p = |P|$ satisfies the hypotheses of the lemma. We bound the probability that the entire neighborhood of P 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 event that $|\Gamma(P) - P| \leq t$. 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:

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

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

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

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 $\Pr[\mathcal{E}] < 1/n^3$ since $cp = \Omega(\log n)$. The claim then follows by applying the union bound as done for Case 1.

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

Proof Let $c = c(n) = \gamma_2 \log(1/r(n)) \ge 3$ and let $p = |P|$. Since $|\Gamma(P)| \ge c(n)$ for all pockets P, then the lemma is trivial when $p < 3$. Suppose now that $p \ge 3$, and fix a subset T of t possible neighbors of P such that $t = \frac{1}{3}cp \ge c$. Let m be the expected number of nodes in a cell. There are at most $\binom{\beta m}{p}$ ways of choosing P and at most $\binom{24\beta m}{\frac{1}{3}cp}$ ways of choosing T. Since a node in Q can choose its neighbors from at least 3 cells, the probability that all cp neighbor choices from P are within nodes of T is at most

$$\left(\frac{\binom{t}{c}}{\binom{3\alpha m}{c}}\right)^p = \left(\frac{t!\,(3\alpha m - c)!}{(3\alpha m)!\,(t-c)!}\right)^p \le \left(\frac{\frac{1}{3}cp}{3\alpha m}\right)^{cp}.$$

Define $\mathcal{E}$ to be the event that all of the sets $P \subseteq Q$, with $|P| < \log n$, choose a number of neighbors in T less than $\frac{1}{3}c\,|P|$. Then,

$$\Pr[\mathcal{E}] \le \sum_{p=1}^{\log n - 1} \binom{\beta m}{p}\binom{24\beta m}{\frac{1}{3}cp}\left(\frac{\frac{1}{3}cp}{3\alpha m}\right)^{cp}$$
$$\le \sum_{p=1}^{\log n - 1} \left(\frac{e\beta m}{p}\right)^p \left(\frac{e24\beta m}{\frac{1}{3}cp}\right)^{\frac{1}{3}cp} \left(\frac{\frac{1}{3}cp}{3\alpha m}\right)^{cp}$$
$$\le \sum_{p=1}^{\log n - 1} \left(\frac{24e^2\beta^2 m^2}{\frac{1}{3}cp^2}\right)^{\frac{1}{3}cp} \left(\frac{\frac{1}{3}cp}{3\alpha m}\right)^{cp}$$
$$\le \sum_{p=1}^{\log n - 1} \left(\tau\frac{c^{2/3}p^{1/3}}{m^{1/3}}\right)^{cp} = O\left(\left(\frac{\log n}{m^{1/3}}\right)^c\right)$$

where $\tau = \left(72e^2\beta^2\right)^{1/3} / (9\alpha)$ is a positive constant. Since $r(n) = \Omega\left(n^{-1/8}\right)$, $\Pr[\mathcal{E}] < 1/n^3$ for a convenient choice of γ_2 in the definition of c. The result follows from the union bound over $k^2 = O\left(n^2\right)$ cells.

Lemma 3 *Let α' be the constant defined in Lemma 1 and let m be the expected number of nodes in a cell. With high probability, for any pair of cells Q and Q', with either $Q' = Q$ or Q' adjacent to Q, and for every pocket $P \subseteq Q$, with $|P| = \alpha' m$, we have $|\Gamma(P) \cap Q'| \geq (1/2 + \epsilon'')\, m$, for a suitable constant $\epsilon'' > 0$.*

Proof Let $c = c(n) = \gamma_2 \log(1/r(n))$. To ease the argument, we suppose that nodes choose their neighbors by picking uniformly at random a node within distance $r(n)$ for $c(n)$ times. Clearly, this process is stochastically dominated by the actual one since in the latter a node cannot be chosen multiple times.

Consider a particular pair of adjacent cells Q and Q' (or $Q' = Q$) and a subset $P \subseteq Q$ of size $p = |P| = \alpha' m$. First we show that with high probability a constant fraction of the cp neighbor choices from nodes of P goes toward nodes of Q'. Then, conditioning on this event, we prove the lemma allowing a suitably large constant γ_2 in the definition of $c(n)$.

The probability q that a node $u \in P$ chooses a neighbor inside Q' is at least $q \geq \alpha m / (\beta \pi n r(n)^2)$, since there are at least αm nodes inside Q' and at most $\beta \pi n r(n)^2$ nodes within distance $r(n)$ from u. By our definition of $m = n/k^2$, we have that $q \geq \alpha/(5\pi\beta) = \Theta(1)$. Let L denote the number of neighbor choices from nodes of P which select a node inside Q'. Note that L does not count the number of distinct nodes of Q' reached from P but it is instead the number of edges from P to Q' (counted with repetitions). By the linearity of expectation, we have $\mathrm{E}[L] = qpc = \Omega(\log n)$ since $p = \alpha' m = \Omega(\log n)$. Applying the standard Chernoff bound [16], we get

$$\Pr\left[L < \frac{1}{2}\mathrm{E}[L]\right] \leq e^{-\frac{1}{8}qpc}.$$

Since there are no more than $5n$ pairs of cells (Q, Q') to be accounted for, and at most $\binom{\beta m}{p} \leq \left(\frac{e\beta m}{\alpha' m}\right)^p$ ways of choosing a pocket P of size $p = \alpha' m$ inside Q, we can conclude that, for any such pair and any such pocket P,

$$\Pr\left[L < \frac{1}{2}qpc\right] \leq 5n\left(\frac{e\beta}{\alpha'}\right)^p e^{-\frac{1}{8}qpc} \leq \frac{1}{n},$$

where the last inequality holds by allowing a sufficiently large constant γ_2 in the definition of $c(n)$. The above inequality implies that, with high probability, $L \geq \sigma m$ for any constant $\sigma > 0$. In the following, we will make use of such a fact for a specific value σ to be determined by the analysis.

Since the neighbor choices are independent and uniform, we can model the neighbor selection process as an instance of the classical balls-and-bins problem, where L balls are thrown inside $b = |Q'|$ bins. Let Z_v be the indicator variable of the event "node $v \in Q'$ was not chosen by any node of P" (i.e., it is an empty bin) and let $Z = \sum_{v \in Q'} Z_v$ denote the total number of nodes of Q' which were not chosen by any node of P (i.e., the total number of empty bins). Conditioning on the event $L \geq \sigma m$, by the linearity of expectation we have

$$\mathrm{E}[Z] \leq b\left(1-\frac{1}{b}\right)^{\sigma m} \leq b e^{-\frac{\sigma m}{b}} \leq \beta e^{-\frac{\sigma}{\beta}} m,$$

since $b \leq \beta m$. The Z_v variables are not independent but they satisfy the Lipschitz condition with bound 1. Therefore, we can apply the Azuma-Hoeffding concentration bound [16], obtaining

$$\Pr\left[Z > 2\beta e^{-\frac{\sigma}{\beta}} m\right] \leq \Pr\left[Z > \mathrm{E}[Z] + \beta e^{-\frac{\sigma}{\beta}} m\right] \leq e^{-\frac{2\left(\beta e^{-\frac{\sigma}{\beta}} m\right)^2}{\sigma m}} \leq \frac{1}{n^2},$$

where the last inequality holds for any value of σ, provided that we choose a sufficiently large constant $\gamma_1 > 0$ in the definition of $r(n)$, since $m = \Theta\left(nr^2(n)\right)$.

Therefore, the number of distinct nodes of Q' reached by nodes of P is

$$|\Gamma(P) \cap Q'| \geq (\alpha m - 2\beta e^{-\frac{\sigma}{\beta}} m) = (1/2 + \epsilon'')m$$

by letting $\sigma = \beta \log \frac{2\beta}{\alpha - 1/2 - \epsilon''}$, and this can be achieved by selecting a suitably large constant γ_2 in the definition of $c(n)$. Invoking the union bound over $O(n)$ pairs of adjacent cells concludes the proof.

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 define $\bar{\varepsilon} = \min\{\varepsilon', \varepsilon''/2, 1/3\}$ where $\varepsilon', \varepsilon''$ are the constants appearing in the statements of Lemma 1 and Lemma 3, respectively, and let $\alpha' \leq \min\{\varepsilon', 1/2\}$ be a suitable small constant (thus, consistent with the constraint posed by the aforementioned lemmas). 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 *red* if it contains at least $\alpha' m$ nodes of S, and *blue* otherwise. Two cases are possible: either a majority of nodes of S resides in red cells or a majority of nodes of S resides in blue cells.

In the first case, $N_{\mathrm{R}} > s/2$ nodes of S belong to red cells. We further subdivide the red cells into two groups, depending on the number of nodes of S they contain. We say that a red cell is "dark red" if it contains at least $(1/2 + \varepsilon''/2)m$ nodes of S; otherwise we call it "light red".

Suppose that a majority of the N_{R} nodes resides inside dark red cells. There are at most $n/((1+\varepsilon'')m)$ dark red cells since $s \leq n/2$, and thus at least $\varepsilon'' n/((1+\varepsilon'')m)$ cells are not dark red. Hence, by well-known topological properties of two-dimensional meshes [15], there are $\Omega\left(\sqrt{s/m}\right)$ disjoint pairs of adjacent cells (Q, Q'), where Q is dark red and Q' is not. Consider one such pair: applying Lemma 3 by taking a pocket $P \subseteq Q \cap S$ of size $|P| = \alpha' m$, we have that

$$|\Gamma(Q) \cap Q' - S| \geq |\Gamma(P) \cap Q' - S| \geq \varepsilon'' m/2 \geq \bar{\varepsilon} m,$$

since, by definition, Q' contains less than $(1/2 + \varepsilon''/2)m$ nodes of S. Summing over the $\Omega\left(\sqrt{s/m}\right)$ disjoint pairs of cells (Q, Q'), we get $\lambda(s) = \Omega\left(\sqrt{m/s}\right)$.

The other subcase where a majority of the N_{R} nodes resides inside light red cells is easier to deal with since we can just consider the expansion of light red cells into themselves. Mimicking the application of Lemma 3 as in the previous subcase, we have that, for any light red cell Q,

$$|\Gamma(Q) \cap Q - S| \geq \varepsilon'' m/2 \geq \bar{\varepsilon} m.$$

Since there are at least $s/((1 + \epsilon'')m)$ light red cells, we immediately obtain $\lambda(s) = \Omega(1) = \Omega\left(\sqrt{m/s}\right)$, which is the correct bound since $s = \Omega(m)$ in this case.

To analyze the second case, where at least $s/2$ nodes of S belong to blue cells, we resort to a proof strategy inspired by the one employed in [8]. 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 B be the set of at least $s/2$ nodes of S belonging to blue cells. To estimate the expansion of S, we first execute a greedy procedure, which selects a number of blue cells which are centers of non-overlapping 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 blue cells are unmarked. Then, iteratively, the center of the next active area is selected as the unmarked blue 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 blue cell becomes marked. The procedure is described by the following pseudo code, where sets I and U maintain, respectively, the indices of the selected centers and the indices of unmarked cells, and subroutine LARGEST-POCKET(U) returns the pair (i, j) corresponding to the unmarked cell with the largest pocket (ties broken arbitrarily).

Algorithm 1 CENTERSELECTION

$I \leftarrow \emptyset$
$U \leftarrow \{(i, j) : Q_{ij}$ is a blue cell$\}$
while $U \neq \emptyset$ **do**
    $(i, j) \leftarrow$ LARGESTPOCKET(U)
    $I \leftarrow I \cup (i, j)$
    for all $Q_{xy} \in \mathcal{S}_{ij}$ **do**
        $U \leftarrow U - \{(x, y)\}$
    end for
end while

Let $\langle c_1, c_2, \ldots, c_w \rangle$ be the list of w centers picked by CENTERSELECTION, 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 b_t be the number of nodes residing in unmarked blue cells of $\mathcal{S}_{c_t}$ at the beginning of iteration t. Clearly, we have that $\sum_{t=1}^{w} b_t = |B|$ and, by the greedy choice of the centers, $b_t \leq 169 p_t$.

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 blue cells of the active area $\mathcal{A}_{c_t}$. We distinguish between two different cases. First, consider the case $s < \alpha' m$, which implies that no red cell exists. Let $N_t = \Gamma(P_{c_t}) - P_{c_t}$ and observe that by Lemmas 1 and 2, $|N_t| \geq \bar{\varepsilon} \min\{c(n) p_t, m\}$. Note that the N_t's are all disjoint, but the sum of their sizes does not immediately yield a lower bound on $|\Gamma(S) - S|$, since each set N_t may itself contain nodes of S, which have to be subtracted from the overall count. Specifically, the number of *external neighbors of S* (i.e., nodes of $\Gamma(S) - S$) accounted for by the N_t's is

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

Since $p_t < \alpha' m$ 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 α', we have that $|N_t| - 169 p_t \geq \mu |N_t|$ for a certain constant $\mu > 0$. Hence,

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

and the theorem follows.

Consider now the case $s \geq \alpha' m$. Observe that $\sum_{t=1}^{w} |N_t| = \Omega(|B|) = \Omega(s)$, 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 $I = \{1, 2, \ldots, t\}$ into two disjoint subsets B_1 and B_2, such that $t \in B_1$ if $\mathcal{A}_{c_t}$ contains no red 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 $(1/2 + \varepsilon'')m$ nodes, and at least $(1/2 + \varepsilon'' - \alpha')m$ of these nodes are external neighbors of S. Hence, the total number of external neighbors of S is at least

$$\sum_{t\in B_2} (1/2+\varepsilon''-\alpha')m = \frac{1/2+\varepsilon''-\alpha'}{1/2+\varepsilon''} \sum_{t\in B_2} |N_t| \geq \frac{1/2+\varepsilon''-\alpha'}{1/2+\varepsilon''} \frac{\tau}{1+\tau} \sum_{t=1}^{w} |N_t| \,,$$

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

$$\begin{aligned}\left(\sum_{t=1}^{w} |N_t|\right) - |B| &= \sum_{t\in B_1} (|N_t| - 169p_t) + \sum_{t\in B_2} (|N_t| - 169p_t) \\ &\geq \sum_{t\in B_1} \mu\,|N_t| + \sum_{t\in B_2} ((1/2+\varepsilon'')m - 169\alpha' m) \\ &> \sum_{t\in B_1} \mu\,|N_t| - \sum_{t\in B_1} \left(\frac{169\alpha'}{1/2+\varepsilon''} - 1\right) \tau\,|N_t| \,.\end{aligned}$$

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

$$\sum_{t\in B_1} \mu\,|N_t| - \sum_{t\in B_1} \left(\frac{169\alpha'}{1/2+\varepsilon''} - 1\right) \tau\,|N_t| = \frac{\mu}{2} \sum_{t\in B_1} |N_t| = \Omega\left(\sum_{t=1}^{w} |N_t|\right) ,$$

and the theorem follows.

3.2 Flooding Time of the Stationary Dynamic Bluetooth Topology

We now turn our attention to a dynamic version of the $BT\,(r(n), c(n))$, where we allow both the positions of the nodes and the set of links to evolve over time. This framework has been adopted by Clementi *et al.* as a model for mobile agents in [14], where the primitive of interest is *flooding*, that is, the spreading of information from one agent to all the others. Building on the relation between graph expansion and flooding established in [14], in this section we study the flooding time of a system of mobile agents where the communication links are established through the neighbor selection protocol of the Bluetooth topology.

Suppose that we have n agents moving along the nodes of a square grid of side 1 and edge-length[1] $\epsilon > 0$. Time is discrete and, in a time step, each agent moves to a grid node chosen uniformly at random among the grid nodes within the Euclidean distance $0 < \rho \leq \sqrt{2}$ from its current position. The parameter

[1] The edge-length ϵ can be made arbitrarily small, and it is introduced only to guarantee the technical condition that the state space of the Markov chain describing the system is finite. In fact, ϵ only affects the constants, hence it does not appear in the results expressed in asymptotic notation.

ρ can be interpreted as the maximum velocity that any agent can achieve. We suppose that all the moves are synchronous and independent. After reaching its new position, each agent establishes communication links with $c(n)$ other agents, chosen uniformly at random among those within the Euclidean distance $r(n)$, or with all of them if they are less than $c(n)$. For $t \in \mathbb{N}$, let G_t be the graph induced by the positions of the agents and the links established at time t. We can formally describe the evolution of the system resulting from this stochastic process as the sequence of graphs $\mathcal{G}(n, \rho, r(n), c(n), \epsilon) = \{G_t : t \in \mathbb{N}\}$, which we call *Dynamic Bluetooth topology.*

In the above dynamic scenario, we aim at upper bounding the flooding time of a message, that is, the minimum number of time steps required to inform all the agents in the system of a message originating from a source agent. When a link connecting an informed agent to an uninformed one is established, the latter becomes informed of the message. It is easily seen that $\mathcal{G}(n, \rho, r(n), c(n), \epsilon)$ constitutes a Markov chain, hence it is a *Markovian evolving graph* according to the Definition 2.1 of [14]. Moreover, when the positions of the nodes in G_0 are chosen according to the stationary distribution (in this case, $\mathcal{G}$ is referred to as *stationary Markovian evolving graph*), the flooding time can be bounded from above based on graph expansion, as established by the following proposition.

Proposition 2 ([14, Corollary 2.5]) *Let $\mathcal{M} = \{G_t : t \in \mathbb{N}\}$ be a stationary Markovian evolving graph. Assume a decreasing sequence $k_1 \geq k_2 \geq \ldots \geq k_{n/2}$ of positive real numbers exists such that, with probability at least $1 - 1/n^4$, G_t has expansion $\lambda(s) \geq k_s$, for every $s = 1, 2, \ldots, n/2$. Then the flooding time in $\mathcal{M}$ is with high probability*

$$O\left(\sum_{s=1}^{n/2} \frac{1}{s \log(1 + k_s)}\right).$$

In order to apply the above result to the Dynamic Bluetooth topology $\mathcal{G}(n, \rho, r(n), c(n), \epsilon)$, we need to lower bound the expansion of each constituent graph G_t. Unfortunately, we cannot directly apply the result of the previous section, since the agents in G_t are not distributed uniformly in $[0, 1]^2$, as it is the case for $BT\,(r(n), c(n))$. However, the stationary distribution of the positions of the agents is *quasi-uniform*, meaning that, with respect to a suitably defined tessellation of the domain into non-overlapping cells of equal size, the number of agents in any two cells differs by at most a constant factor, which is sufficient to obtain a lower bound on the expansion. In what follows we consider a stationary Dynamic Bluetooth topology $\mathcal{G}_{\text{STAT}}(n, \rho, r(n), c(n), \epsilon)$, where the positions of the agents in G_0 are chosen randomly according to the stationary distribution, and tessellate the domain into k^2 cells where $k = \lceil \sqrt{5}/r(n) \rceil$, and $m = n/k^2$. The following proposition is analogous to Proposition 1.

Proposition 3 (Quasi-Uniformity) *Consider a stationary Dynamic Bluetooth topology $\mathcal{G}_{\text{STAT}}(n, \rho, r(n), c(n), \epsilon)$ 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 and arbitrary positive constants ρ and ϵ. Let m be the expected number of agents in a cell. Then, with probability $1 - 1/n^5$, in the stationary Dynamic Bluetooth topology $\mathcal{G}_{\mathrm{STAT}}(n, \rho, r(n), c(n), \epsilon)$, the number of agents $N_Q(t)$ residing in cell Q at time t satisfies

$$m/\mu \leq N_Q(t) \leq \mu m,$$

for each cell Q, $0 \leq t < n$, and for some constant $\mu > 0$.

Proof Let $\pi(x)$ be the probability that an agent is located at grid node x in the stationary distribution. As noted in [14], for any two positions x and y, there exists a constant $\sigma > 0$, depending on ϵ and ρ, such that $1/\sigma \leq \pi(x)/\pi(y) \leq \sigma$. Let $N_Q(t)$ denote the number of agents residing inside cell Q at time t. Since the size of each cell is $1/k^2$ and the distribution of the agents' position is stationary, we have that $\mathrm{E}[N_Q(t)] \geq m/\sigma$ and $\mathrm{E}[N_Q(t)] \leq \sigma m$. Also, the agents move independently, so we can apply the Chernoff bound on $N_Q(t)$, obtaining that $\Pr[N_Q(t) \leq m/(2\sigma)] \leq 1/n^7$, and $\Pr[N_Q(t) \geq 3\sigma m/2] \leq 1/n^7$, for a sufficiently large constant γ_1. The proof follows by applying the union bound over $O(n)$ cells and the n time instants and setting $\mu = 2\sigma$.

The quasi-uniform distribution of the agents described in Proposition 3 allows us to characterize the expansion of each snapshot G_t of the process, as stated in the next theorem.

Theorem 2 *Consider a stationary Dynamic Bluetooth topology $\mathcal{G}_{\mathrm{STAT}}(n, \rho, r(n), c(n), \epsilon)$, 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 and arbitrary positive constants ρ and ϵ. Then there exist positive constants $\delta_1, \delta_2, \delta_3, \alpha$ such that, with probability $1 - 1/n^4$, each graph $G_t \in \mathcal{G}_{\mathrm{STAT}}(n, \rho, r(n), c(n), \epsilon)$ has expansion*

$$\begin{aligned} \lambda(s) &\geq \delta_1 c(n) && \text{if} \quad 1 \leq s \leq m/c(n), \\ \lambda(s) &\geq \delta_2 m/s && \text{if} \quad m/c(n) < s \leq \alpha m, \\ \lambda(s) &\geq \delta_3 \sqrt{m/s} && \text{if} \quad s > \alpha m, \end{aligned}$$

for all $0 \leq t < n$, where m is the expected number of agents in a cell.

Proof The proof follows exactly the same steps of the proof of Theorem 1 for $BT(r(n), c(n))$. In fact, it suffices to observe that the result of Proposition 3 enables us to prove lemmas analogous to Lemmas 1, 2 and 3, where the constants involved in the pocket expansion become suitable functions of μ.

Theorem 3 *Consider a stationary Dynamic Bluetooth topology $\mathcal{G}_{\mathrm{STAT}}(n, \rho, r(n), c(n), \epsilon)$, 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 and arbitrary positive constants ρ and ϵ. Then, with high probability, the flooding time in $\mathcal{G}_{\mathrm{STAT}}(n, \rho, r(n), c(n), \epsilon)$ is*

$$T_{\mathrm{FL}} = O\left(\frac{1}{r(n)} + \log n\right).$$

Proof Let m be the expected number of agents in a cell. By plugging in the lower bounds on the expansion stated in Theorem 2 in the formula given in Proposition 2, we have that

$$T_{\mathrm{FL}} = O\left(\sum_{s=1}^{m/c(n)} \frac{1}{s\log(1+\delta_1 c(n))} + \sum_{s=m/c(n)+1}^{\alpha m} \frac{1}{s\log(1+\delta_2 m/s)} + \sum_{s=\alpha m+1}^{n/2} \frac{1}{s\log(1+\delta_3\sqrt{m/s})}\right).$$

We evaluate the three summations separately. The first summation easily yields

$$\sum_{s=1}^{m/c(n)} \frac{1}{s\log(1+\delta_1 c(n))} = \frac{1}{\log(1+\delta_1 c(n))} H\left(\frac{m}{c(n)}\right) = O(\log n).$$

For the second summation, since we can always presume that $\delta_2 \le \alpha$, we obtain

$$\begin{aligned}\sum_{s=m/c(n)+1}^{\alpha m} \frac{1}{s\log(1+\delta_2 m/s)} &\le 2\frac{\alpha}{\delta_2} \sum_{s=m/c(n)+1}^{\alpha m} \frac{1}{s\log(1+\delta_2 m/s)} \frac{\delta_2 m}{\alpha m + s} \\ &\le 2\frac{\alpha}{\delta_2} \int_{m/c(n)}^{\alpha m} \frac{1}{x\log(1+\delta_2 m/x)} \frac{\delta_2 m}{\delta_2 m + x} \mathrm{d}x \\ &= O(\log\log c(n)).\end{aligned}$$

Since $c(n) = O(\log n)$, we have that the second summation is bounded by $O(\log\log\log n)$. Finally, for the third summation we obtain

$$\begin{aligned}\sum_{s=\alpha m+1}^{n/2} \frac{1}{s\log(1+\delta_3\sqrt{m/s})} &\le \sum_{s=\alpha m+1}^{n/2} \frac{1+\delta_3/\sqrt{\alpha}}{\delta_3\sqrt{m}} \frac{1}{\sqrt{s}} \\ &\le 2\frac{1+\delta_3/\sqrt{\alpha}}{\delta_3} \frac{1}{\sqrt{m}} \int_{\alpha m}^{n/2} \frac{1}{\sqrt{x}} \mathrm{d}x \\ &= O\left(\frac{1}{r(n)}\right).\end{aligned}$$

Summing the three contributions concludes the proof of the theorem.

3.3 Upper Bound

In this section we prove that the lower bound established by Theorem 1 is asymptotically tight.

Theorem 4 *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*

$$\lambda(s) = \begin{cases} O\left(\min\{c(n), m/s\}\right) & \text{if} \quad s = O(m) \\ O\left(\sqrt{m/s}\right) & \text{if} \quad s = \Omega(m), \end{cases}$$

where m is the expected number of nodes in a cell.

Proof We fix the constant γ_1 so that Proposition 1 holds. If $s \leq \alpha m$, we can choose any subset S of the nodes in a single corner cell, so that a total of at most $13\beta m$ nodes are visible from S. Hence, $\lambda(s) = O(m/s)$. Consider a list $\langle v_1, v_2, \ldots, v_n \rangle$ of the vertices of V, sorted by non-decreasing degree. If we take $S = \{v_1, v_2, \ldots, v_s\}$, then we are guaranteed that the sum of the degrees of all nodes in S is no 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 proves the theorem for the case $s \leq \alpha m$.

Consider now the case $s > \alpha m$ and choose 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 $O\left(\sqrt{s/m}\right)$ cells are visible from S, we have that $\lambda(s) = O\left(\sqrt{ms}/s\right) = O\left(\sqrt{m/s}\right)$, and the theorem follows.

We remark that the tight bounds on the expansion of *BT* $(r(n), c(n))$ provided by Theorems 1 and 4 extend the results in [8] from the case $r(n) = \Theta(1)$ to arbitrary values of $r(n)$ that guarantee 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 $\lambda = \Theta(r(n))$ for the Bluetooth topology.

4 Expansion of *RGG* $(r(n))$

The analysis performed in the previous section for *BT* $(r(n), c(n))$ can also be applied to *RGG* $(r(n))$ with simpler technical arguments, due to the absence of the neighbor selection procedure. Indeed, the following theorem establishes the asymptotic order of the expansion $\lambda(s)$ of a random geometric graph, for all values of s.

Theorem 5 *Consider an instance of RGG $(r(n))$ with $r(n) \geq \gamma_1 \sqrt{\log n/n}$, for a suitable constant γ_1. With high probability, for every integer s, $1 \leq s \leq n/2$, we have*

$$\lambda(s) = \begin{cases} \Theta(m/s) & \text{if} \quad s = O(m), \\ \Theta\left(\sqrt{m/s}\right) & \text{if} \quad s = \Omega(m), \end{cases}$$

where m is the expected number of nodes in a cell.

Proof We fix the constant γ_1 so that Proposition 1 holds. For the lower bound, consider a subset $S \subseteq V$, $|S| = s$, with $1 \leq s \leq n/2$. If $s \leq \alpha m$, recall that Proposition 1 implies that any node $u \in S$ has at least $(\alpha/4)\pi n r^2(n)$ neighbors in $RGG(r(n))$. Since $(\alpha/4)\pi n r^2(n) - s = \Omega(m)$, we have that $\lambda(s) = \Omega(m/s)$. If $s > \alpha m$, similarly to the proof of Theorem 1 we say that a cell Q such that $Q \cap S \neq \emptyset$ is *red* if it contains at least $(3/4)\alpha m$ nodes of S, and *blue* otherwise. Two subcases are possible: either a majority of nodes of S resides in red cells or a majority of nodes of S resides in blue cells. In the first case, since $s \leq n/2$, the number N_R of red cells satisfies $N_R \leq (2/(3\alpha))n/m$ and therefore the number of non-red cells is at least $n/m - N_R = \Omega(n/m)$. Therefore, at least $\Omega\left(\sqrt{n/m}\right) = \Omega\left(\sqrt{s/m}\right)$ red cells are adjacent to a non-red cell [15], and each of these (disjoint) pairs contributes $\Omega(m)$ nodes to the expansion of S, yielding $\lambda(s) = \Omega\left(\sqrt{m/s}\right)$. On the other hand, if a majority of nodes of S resides in blue cells, there are at least $s/((3/4)\alpha m)$ blue cells, and each contributes at least $(\alpha/4)m$ nodes to the expansion of S, since any node residing in that cell connects to every other node residing in the same cell. In this second case, we get $\lambda(s) \geq 1/3 = \Omega\left(\sqrt{m/s}\right)$ since $s > \alpha m$.

In order to complete the proof, we derive a matching upper bound. If $s \leq \alpha m$, we can pick the set S entirely contained in a single corner cell. Since the number of nodes visible from S is bounded by $13\beta m$, we have that $\lambda(s) = O(m/s)$. On the other hand, if $s > \alpha m$, consider the "densest" set S as in the proof of Theorem 4. We immediately conclude that the neighborhood of S has at most $O\left(m\sqrt{s/m}\right)$ nodes, and thus $\lambda(s) = O\left(\sqrt{m/s}\right)$.

Quite surprisingly, Theorems 1, 4, and 5 imply that the expansion $\lambda(s)$ of $BT(r(n), c(n))$ is, within a constant factor, equal to the expansion of $RGG(r(n))$, as soon as we consider a set of $s = \Omega(m/c(n))$ vertices, although $BT(r(n), c(n))$ is a (possibly very sparse) subgraph of $RGG(r(n))$.

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

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 $0 \le i \le d$, 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 $|Y_i| \ge N_i$. Define now $j^\star$ as the smallest index such that $|Y_{j^\star}| > n/2$, which implies that $j^\star \le i^\star$. Also, w.l.o.g., we can assume that $j^\star \ge \lceil d/2 \rceil$, 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 up to level $\lceil d/2 \rceil - 1$, or there would be a path shorter than d connecting u and v. The lemma follows.

Theorem 6 *Consider an instance of $BT(r(n), c(n))$ with $r(n) \ge \gamma_1 \sqrt{\log n/n}$ 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 Let m denote the expected number of nodes in a cell. We apply Lemma 4 by estimating the value $i^\star$ 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 $s = O(m)$, and $\lambda(s) = \Omega\left(\sqrt{m/s}\right)$ for $s = \Omega(m)$, which happens with high probability (see Theorem 1).

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

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

Since $K(j)$ is defined as the smallest index for which $N_{K(j)} \ge 2^j$, from the above inequalities it follows that $K(j) \le \min\{\ell \ : \ (1 + \sigma\sqrt{m}/2^{j/2})^{\ell-K(j-1)} \ge$

2}, hence $K(j) - K(j-1) = O\left(2^{j/2}/(r(n)\sqrt{n})\right)$. Therefore,

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

and the theorem follows from Lemma 4.

We now show that Theorem 6 gives a tight estimate for the diameter of $BT(r(n), c(n))$ when $r(n) = O(1/\log n)$.

Theorem 7 *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,*

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

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.

The above lower bound can be improved for large visibility radii $r(n) = \Omega(\log\log n/\log n)$, yielding a lower bound almost matching the $O(\log n)$ upper bound given by Theorem 6.

Theorem 8 *Consider an instance of $BT(r(n), c(n))$ with $\gamma_1 \log\log n/\log n \leq r(n) \leq \sqrt{2}$ and $c(n) \geq \gamma_2 \log(1/r(n))$, for two suitable positive constants γ_1 and γ_2. With high probability,*

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

Proof We show that, with high probability, each node of the graph has degree bounded by $\Delta = c(n) + 2\log_2 n$ and therefore the diameter of $BT(r(n), c(n))$ cannot be smaller than the diameter of a tree with arity Δ.

Consider an arbitrary node $u \in BT(r(n), c(n))$ and denote by $\deg(u)$ its degree. By the definition of Bluetooth topology, the number $\deg_{\mathrm{OUT}}(u)$ of neighbors chosen by u satisfies (deterministically) $\deg_{\mathrm{OUT}}(u) \leq c(n)$. For each node $v \neq u$ within distance $r(n)$ from u, let X_v be a 0/1 random variable, taking value 1 iff node v selects u as its neighbor. Observe that the number $\deg_{\mathrm{IN}}(u)$ of nodes which selected u as their neighbor can be written as

$$\deg_{\mathrm{IN}}(u) = \sum_v X_v,$$

where the summation ranges over all nodes within distance $r(n)$ from u. It is straightforward to see that $\Pr[X_v = 1] = c(n)/(N_v - 1)$, where N_v denotes the number of nodes in the visibility disk of v. By Proposition 1, we have that

$$\mathrm{E}[\deg_{\mathrm{IN}}(u)] \leq \frac{4\beta}{\alpha} c(n) = O(\log \log n).$$

Let $t = 2\log_2 n$. For a sufficiently large n, $t \geq 6\mathrm{E}[\deg_{\mathrm{IN}}(u)]$. Since the neighbor choices performed by different nodes are independent, we can apply the Chernoff bound [16] to $\deg_{\mathrm{IN}}$ to obtain that $\Pr[\deg_{\mathrm{IN}} \geq t] \leq 2^{-t} = 1/n^2$. Applying the union bound over the n nodes yields that with probability at least $1 - 1/n$ all the nodes have in-degree at most t and hence their degree is at most $\Delta = c(n) + t = O(\log n)$. The theorem follows by observing that $\mathrm{diam}(BT(r(n), c(n))) \geq \log_\Delta n = \Omega\left(\frac{\log n}{\log \log n}\right)$.

We conclude this section by noticing that the $BT(r(n), c(n))$ exhibits, for not too large visibility radii, the same asymptotic diameter of the denser random geometric graph with the same parameter $r(n)$. Indeed, [10] proves that $\mathrm{diam}(RGG(r(n))) = O(1/r(n))$, while Theorem 6 yields $\mathrm{diam}(BT(r(n), c(n))) = O(1/r(n))$ for $r(n) = O(1/\log n)$.

6 Conclusions

The main result of this paper is a tight characterization of the 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 [13]. 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 results of [8] and of [12] which hold only for the extreme cases $r(n) = \Theta(1)$ and $r(n) = \Theta\left(\sqrt{\log n/n}\right)$, respectively.

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. For the larger ranges $r(n) = \Omega(1/\log n)$ we provided a more sophisticated lower bound which matches the upper bound up to a $\log \log n$ factor.

A somewhat surprising consequence of our results is that for subsets of $s = \Omega(m/c(n))$ nodes, $BT(r(n), c(n))$ exhibits roughly the same expansion as the random geometric graph $RGG(r(n))$ of [7], which is a much denser supergraph of $BT(r(n), c(n))$. 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 $r(n) = \omega\left(\sqrt{\log n/n}\right)$, the *minimum* number $c(n)$ of neighbor choices which yield connectivity and to assess the corresponding diameter and expansion properties.

Acknowledgements Support for the authors was provided, in part, by the Italian "Ministero dell'Istruzione, dell'Università e della Ricerca" under project AlgoDEEP, and by the University of Padova under the Strategic Project STPD08JA32 and Project CPDA099949/09. Part of this work was done while the first author was visiting the Department of Computer Science of Brown University, partially supported by "Fondazione Ing. Aldo Gini", Padova, Italy.

References

1. Pettarin, A., Pietracaprina, A., Pucci, G.: On the Expansion and Diameter of Bluetooth-like Topologies. In: Proc. of the 17th European Symposium on Algorithms. (2009) 528–539
2. Whitaker, R., Hodge, L., Chlamtac, I.: Bluetooth scatternet formation: a survey. Ad Hoc Networks **3** (2005) 403–450
3. Stojmenovic, I., Zaguia, N.: Bluetooth scatternet formation in ad hoc wireless networks. In Misic, J., Misic, V., eds.: Performance Modeling and Analysis of Bluetooth Networks, Auerbach Publications (2006) 147–171
4. 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
5. 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 International Parallel and Distributed Processing Symposium. (2004) 221–228
6. 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
7. Penrose, M.: Random Geometric Graphs. Oxford University Press (2003)
8. Panconesi, A., Radhakrishnan, J.: Expansion properties of (secure) wireless networks. In: Proc. of the 16th ACM Symposium on Parallelism in Algorithms and Architectures. (2004) 281–285
9. Dubhashi, D., Johansson, C., Häggström, O., Panconesi, A., Sozio, M.: Irrigating ad hoc networks in constant time. In: Proc. of the 17th ACM Symposium on Parallelism in Algorithms and Architectures. (2005) 106–115
10. Ellis, R., Jia, X., Yan, C.: On random points in the unit disk. Random Structures and Algorithms **29**(1) (2005) 14–25
11. Dubhashi, D., Häggström, O., Mambrini, G., Panconesi, A., Petrioli, C.: Blue Pleiades, a new solution for device discovery and scatternet formation in multi-hop Bluetooth networks. Wireless Networks **13**(1) (2007) 107–125
12. Broutin, N., Devroye, L., Fraiman, N., Lugosi, G.: Connectivity threshold of Bluetooth graphs. E-print arXiv:1103.0351v1 [math.PR] (2011)
13. Leighton, T., Rao, S.: Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. Journal of the ACM **46**(6) (1999) 787–832
14. Clementi, A.E., Monti, A., Pasquale, F., Silvestri, R.: Information spreading in stationary Markovian evolving graphs. In: Proc. of the 23rd International Parallel and Distributed Processing Symposium. (2009) 1–12
15. Bilardi, G., Preparata, F.P.: Area-time lower-bound techniques with applications to sorting. Algorithmica **1**(1) (1986) 65–91

16. Mitzenmacher, M., Upfal, E.: Probability and Computing. Randomized Algorithms and Probabilistic Analysis. Cambridge University Press (2003)