

Visual Target Tracking Techniques for Autonomous Vehicle Navigation

Chanda L. Mwansa¹, Mukuka N. Bwalya², and Thandiwe P. Zulu³

¹Department of Electrical Engineering, University of Zambia, Lusaka, Zambia

²Department of Computer Science, Copperbelt University, Kitwe, Zambia

³Department of Civil Engineering, Mulungushi University, Kabwe, Zambia

Abstract:

Recent years have witnessed a significant surge of interest in the design and development of local algorithms for solving various network tasks, particularly due to their critical importance in the study and application of wireless networks. These local algorithms, characterized by their reliance on limited and localized information, are essential for ensuring efficient and scalable network operations. In this work, we specifically focus on the behavior and performance of algorithms operating under very low locality constraints, where each node has access to minimal information about its immediate neighborhood. Despite the inherent limitations imposed by such restricted locality, we successfully propose local constant-ratio approximation algorithms for solving several fundamental problems in location-aware Unit Disk Graphs (UDGs). These problems include the Minimum Dominating Set (MDS), Connected Dominating Set (CDS), Maximum Independent Set (MIS), and Minimum Vertex Cover (MVC). The proposed algorithms demonstrate that even with constrained locality, it is possible to achieve effective approximations for these challenging combinatorial problems in the context of wireless networks. Furthermore, this study contributes to the theoretical understanding of local algorithms by establishing the first-ever lower bounds on their performance for the aforementioned problems within a location-aware UDG setting. These bounds provide valuable insights into the trade-offs between locality and the achievable solution quality, offering a deeper understanding of the limitations and potential of local algorithms in practical scenarios. By addressing both algorithmic solutions and theoretical limits, this work not only advances the state of the art in local algorithm design for wireless networks but also lays a foundation for future research aimed at further exploring and optimizing locality-aware approaches in distributed systems.

Keywords: Unit Disk Graphs; Location Awareness; Local Algorithms; Approximation Algorithms; Lower Bounds

1. Introduction

In wireless networks, especially in ad hoc settings, the absence of a central authority to coordinate network traffic necessitates decentralized approaches. Devices in such networks must form structures, such as communication backbones, by exchanging information locally. Due to the typically vast size of these networks, it is impractical for any single node to have complete knowledge of the entire network. This constraint highlights the importance of **local algorithms**, where the status of a vertex v (e.g., its inclusion in a dominating set, independent set, etc.) depends solely on the vertices within a fixed number of edges (hops) from v . Importantly, this locality parameter remains independent of the overall network size.

Wireless networks are often modeled using **Unit Disk Graphs (UDGs)**—undirected graphs in which two vertices are connected by an edge if and only if their Euclidean distance is at most one unit. UDGs effectively capture the geometric and spatial constraints of wireless networks. In our model, each node is assumed to have knowledge of its geographic position in the plane, a reasonable assumption given the growing prevalence of positioning systems like GPS.

In wireless networks, **dominating sets** play a crucial role in topology control. Nodes are often grouped into clusters, with a designated cluster-head managing communication within the cluster and facilitating interactions with other clusters. The set of these cluster-heads forms a **dominating set** in the network graph. For efficient communication between clusters, it is essential for this dominating set to be **connected**, leading to the concept of a **connected dominating set**. To enhance network efficiency and reduce interference, it can be beneficial for cluster-heads to be non-adjacent, forming a **maximal independent set**. Additionally, the **minimum vertex cover problem** is another critical graph-theoretic challenge relevant to wireless networks.

While numerous algorithms have been proposed to address these tasks, there is a notable gap in our understanding of lower bounds for the approximation ratios achievable by local algorithms. Such bounds are essential for assessing the feasibility and limitations of local algorithms under specific locality constraints. In this paper, we aim to bridge this gap by investigating the fundamental restrictions imposed on algorithmic performance when locality is constrained. This

work not only contributes to the theoretical foundation of local algorithms but also provides insights into their practical implications for the design of efficient and scalable wireless network solutions.

1.1. Related Work

The first results about local algorithms were given by Linial [2]. He gives the first bounds on locality distances for constructing a maximal independent set and a 3-coloring in an n -cycle. He also proves that at least $d - \sqrt{d}$ colors are needed for coloring a d -regular tree with radius r when the locality is restricted to $2r/3$. In [3] Naor and Stockmeyer provide a framework for local algorithms for Locally Checkable Labeling Problems (LCL). All problems which we discuss in this paper are LCLs. In his book [4], Peleg gives a locality sensitive perspective of distributed algorithms.

In general graphs all problems which we study - dominating and connected dominating set, independent set and vertex cover - are NP-hard [5]. Apart from vertex cover they do not even admit constant ratio approximation algorithms [6, 7]. For vertex cover there are several 2-approximation algorithms known, e.g., [8]. When restricting the case to unit disk graphs, the problems remain NP-hard [9], but constant ratio approximations [10] and PTASs [11, 12, 13] are known. However, all these algorithms are global in the sense that in order to determine whether a given vertex is in the computed set, we need knowledge of the entire graph. When looking for local algorithms, Kuhn et al. [14] proposed local approximation schemes for maximum independent set and dominating set for growth-bounded-graphs. This class of graphs includes UDGs. However, in these algorithms the status of a vertex depends on the vertices which are up to $O(\log |V|)$ hops away from it, which is not local in our sense. Gfeller and Vicari presented a distributed PTAS for dominating set with the same locality properties [15]. In [16] the authors provide a marking process in conjunction with two dominant pruning rules in order to reduce the size of a dominating set.

In [17] Czyzowicz et al. presented the first local constant ratio approximation algorithms for our setting with performance ratios 5 and $7.453 + \epsilon$ for dominating and connected dominating set, respectively. The locality distances of the dominating set algorithm is 11, the locality for the connected dominating set algorithm is not given in the paper. In [18] Wiese and Kranakis presented local PTASs for these two problems. However, their locality distances, though constant, can be very large.

1.2. Results of this paper

In this paper we try to assess the impact of locality on the algorithmic design of important computational tasks, like dominating set, connected dominating set, vertex cover, and independent set, in wireless networks. For arbitrary locality distances we show the first lower bounds for possible approximation ratios of local algorithms for all problems mentioned above in the setting of location aware nodes.

We also investigate the power of algorithms with very low localities. It turns out that despite the fact we are looking only at locality one we can still design constant ratio approximation algorithms for dominating set, independent set, and vertex cover. We prove that for connected dominating set there is no constant ratio algorithm with locality distance one. But we present such an algorithm with locality distance two.

In the proofs for the lower bounds we mostly employ unit disk graphs on a line (unit line graphs). In order to assess our bounds for each problem we present a local algorithm with locality one in unit line graphs. These algorithms achieve significantly better approximation ratios than the local algorithms for general unit disk graphs with this locality.

All results including trade-offs between approximation ratios and bounds are presented.

1.3. Organization of the paper

The remainder of this paper is organized as follows: First we give some basic definitions and preliminaries. In the following sections we discuss one problem per section: dominating set, connected dominating set, independent set, and finally vertex cover. For each problem we prove our lower bounds for approximation ratios of local algorithms with arbitrary locality distance. We also give our algorithms for unit disk graphs and unit line graphs. For our dominating set algorithm on unit disk graphs we give an example which shows that the analysis of the approximation ratio is tight.

1.4. Preliminaries

A Unit Disk Graph (UDG) is an undirected graph which has an embedding in the plane such that two vertices are connected by an edge if and only if their Euclidean distance is at most one.

Definition 1.

A Unit Line Graph (ULG) is a unit disk graph in which all vertices have the same y -coordinate.

Let $G = (V, E)$ be an undirected graph. A set $D \subseteq V$ is called a *dominating set* if each vertex in V is either in D or adjacent to a vertex in D . A set $CD \subseteq V$ is a *connected dominating set* if it is a dominating set such that the subgraph induced by CD is connected. We call a set $I \subseteq V$ an *independent set* if it does not contain two adjacent vertices. Finally a set $V_C \subseteq V$ is called a *vertex cover* if for every edge $e = (u, v)$ it holds that either $u \in V_C$ or $v \in V_C$.

Definition 2.

For two vertices u and v let $d(u, v)$ be the hop-distance between u and v , that is the number of edges on a shortest path between these two vertices.

For a set of vertices $V' \subseteq V$ we define $d_{V'}(u, v)$ to be the hop-distance between $u \in V'$ and $v \in V'$ in the subgraph induced by V' . The hop-distance is not necessarily the geometric distance between two vertices. Denote by $N^r(v) = \{u \in V \mid d(u, v) \leq r\}$ the r -neighborhood of a vertex v . For ease of notation we define $N^0(v) := \{v\}$, $N(v) := N^1(v)$ and for a set $V' \subseteq V$ we define $N(V') = \bigcup_{v \in V'} N(v)$. Note that $v \in N(v)$.

Denote by the *locality distance* (or short the *locality*) of an algorithm the minimum k such that the status of a vertex v (e.g., whether or not v is in a dominating or connected dominating set) depends only on the vertices in $N^k(v)$.

2. Dominating Set

We propose a local approximation algorithm that achieves a factor-12 approximation for the **dominating set problem** on unit disk graphs. The algorithm's correctness, approximation ratio, and locality—exactly one hop—are rigorously proven. Furthermore, we demonstrate the tightness of our approximation factor analysis by providing a corresponding lower bound. To broaden our understanding, we establish lower bounds for local algorithms with an arbitrary locality distance k , using **unit line graphs** as the basis for our analysis. Lastly, we present a local algorithm specifically designed for dominating sets on unit line graphs, with a one-hop locality and an approximation factor of 3.

2.1 Tiling of the Plane

The foundation of our algorithm involves a systematic **tiling of the plane** using hexagons, designed to meet the following criteria:

1. Each vertex in the graph is contained within exactly one hexagon.
2. All vertices within a single hexagon are connected by edges.

To achieve these properties, the plane is partitioned into hexagons with a diameter of one. Ambiguities arise when vertices lie on the borders of hexagons; these are resolved using a defined rule set, illustrated in **Figure 1 (b)**: the right borders of hexagons, excluding their upper and lower apices, are included as part of the hexagon, while the remainder of the border is excluded. The tiling begins with the origin $(0,0)$, positioned at the center of a hexagon classified as **class 1**. This structured tiling ensures a clear geometric foundation, enabling the algorithm to operate with precision and simplicity while maintaining its locality constraint. The hexagonal layout is particularly advantageous for ensuring connectivity and achieving the desired approximation guarantees within the algorithm's framework.




Figure 1. (a) A part of the tiling of the plane into hexagons. (b) One hexagon of the tiling. The bold lines indicate the parts of its border that belong to this hexagon.

This is similar to the tiling used in [17] but in contrast to their tiling our hexagons do not have any information that distinguishes them from one another (class numbers, etc.).

2.2. Algorithm for Unit Disk Graphs

Now we present a local algorithm for dominating set on unit disk graphs (UDGs) with locality one. Let $G = (V, E)$ be a unit disk graph. The algorithm works as follows: In each hexagon h the vertex v which is closest to the center of h is assigned to the dominating set D . Ambiguities are resolved by e.g., choosing the vertex with the smallest x -coordinate. We refer to this as Algorithm 1.

We prove the correctness of Algorithm 1 in Theorem 1.

Theorem 1.

Let G be a unit disk graph. Algorithm 1 has the following properties:

1.
The computed set D is a dominating set for G .
2.
Let D_{OPT} be an optimal dominating set. It holds that $|D| \leq 12 \cdot |D_{OPT}|$.
3.
Whether or not a vertex v is in D depends only on the vertices at most one hop away from v , i.e. Algorithm 1 is local.
4.
The processing time for a vertex v is linear in the number of vertices adjacent to v .

Algorithm 1: Local algorithm for finding a dominating set in a unit disk graph

1 // Algorithm is executed independently by each node v ;

2 // Denote by V_h all vertices in h ;

3 Find all vertices in $N(v)$ and compute V_h ;

4 if v is the vertex closest to the center of h among all $v' \in V_h$ then become part of D else Do not become part of D

Proof.

We first prove that D is indeed a dominating set for G . Let h be a non-empty hexagon and denote by V_h all vertices in h . As $V_h \neq \emptyset$ it holds that one vertex $v \in V_h$ is the vertex which is closest to the center of h . So $v \in D$ and v dominates all vertices in h .

Now we prove that for an optimal dominating set D_{OPT} it holds that $|D| \leq 12 \cdot |D_{OPT}|$. In the following we prove that for any vertex in D_{OPT} there can exist at most 12 hexagons that have a vertex at (Euclidean) distance at most one from it. (In fact this holds for all vertices but we need the claim only for vertices in D_{OPT} .) Consider a vertex $v \in D_{OPT}$ in a hexagon h . Without loss of generality we assume that v is in the gray area of h or in the center of h . In both cases for a vertex $v' \in D$ that is adjacent to v it holds that v' must be in one of the 9 or 12 hexagon in the respective figure (for the case of v being in the center of h check the resolving method for ambiguities at the border of hexagons. Since at most one vertex per hexagon was assigned to D there are at most 12 vertices in D that are adjacent to v . (Note that in the case where v is in the gray area the resolving method for ambiguities does not come into play.)




Figure 2. A vertex v in the center of hexagon h can only be adjacent to vertices in these 9 hexagons. Note that due to the resolving method for vertices at the border of hexagons, the vertices v' and v'' belong to the hexagons h' and h'' respectively (and are adjacent to v).




Figure 3. A vertex in the gray area excluding the center of h can only be adjacent to vertices in these 12 hexagons.

So for each vertex $v \in D_{OPT}$ it holds that it is adjacent to at most 12 vertices in D . Since D_{OPT} is a dominating set all vertices in D must be adjacent to at least one vertex in D_{OPT} . This leads to $|D| \leq 12 \cdot |D_{OPT}|$.

Next we determine the locality of Algorithm 1. In the algorithm, a vertex v in a hexagon h only needs to explore all vertices which are at most one hop away from v in order to determine whether v is the vertex closest to the center of h . So whether or not $v \in D$ depends only on the vertices which are at most one hop away from v .

Now we prove the processing time of Algorithm 1. Let v be a vertex and define $n_1(v) := |N^1(v)|$. In the algorithm v needs to determine for each vertex $v' \in N^1(v)$ whether $v' \in h$ and whether v' is closer to the center of h than v . This can be done in $O(n_1(v))$ time. $\square$

2.3. Tightness of Approximation Factor

We give an example which shows that our analysis of Algorithm 1 is tight. denote it by G (we omitted the edges which are not adjacent to v in order to make the figure clearer). Note that the vertex v is on the crossing of three hexagons but is assigned to the hexagon h . The vertex w is on a crossing of three hexagons as well but by our resolving method it is assigned to the hexagon h' . Also note that the vertex u is directly in the center of its hexagon. We observe that the vertex v alone is sufficient to dominate all other vertices in G . Since each hexagon contains only one vertex, Algorithm 1 assigns every vertices in the graph to D . So Algorithm 1 achieves a competitive ratio of 12 in G .




Figure 4. Tight example for Algorithm 1. To make the figure clearer the edges which are not adjacent to v are omitted. We can enlarge this construction to an arbitrary size by putting several copies of G together. The copies of G are indeed connected since the length of an edge of a hexagon is $1/2$ and the diameter of a hexagon is 1.




Figure 5. The enlargement of our construction of the tight example. It is possible to extend this to an arbitrarily large graph.

3. Connected Dominating Set

In this section we present a local approximation algorithm that computes a factor 216 approximation for connected dominating set. We prove its correctness, its approximation ratio and its locality distance of two hops. We prove that in contrast to the other problems discussed in this paper there is no constant ratio approximation algorithm with locality distance one. We give lower bounds for the approximation ratio of local algorithms with arbitrary locality distance k . Finally we give an approximation algorithm for connected dominating set for unit line graphs which achieves an approximation factor of 6.

3.1. Algorithm for Unit Disk Graphs

We present our approximation algorithm for connected dominating set in unit disk graphs with locality two. Let $G = (V, E)$ be a connected unit disk graph. We use the same tiling of the plane as introduced. The algorithm works as follows: Consider one hexagon h . We consider each hexagon h' such that there are pairs of adjacent vertices v and v' with $v \in h$, $v' \in h'$. We compute the two adjacent vertices $v \in h$, $v' \in h'$ which are closest to each other (ties are resolved by some resolving method, e.g., by choosing the pair where the vertex in h has the smallest x-coordinate etc.). We assign v and v' to the connected dominating set CD . Do this for all hexagons h' with the above property. If there are no such hexagons h' we assign the vertex $v \in h$ to CD which is closest to the center of h (ties are broken like in Algorithm 1). The above description is presented in Algorithm 3.

We prove the correctness and the other properties of Algorithm 3 in Theorem 4.

Theorem 4.

Let G be a unit disk graph. Algorithm 3 has the following properties:

The computed set CD is a connected dominating set for G .

Let CD_{OPT} be an optimal connected dominating set. It holds that $|CD| \leq 216 \cdot |CD_{OPT}|$.

Whether or not a vertex v is in CD depends only on the vertices, i.e. Algorithm 3 is local.

The processing time for a vertex v is quadratic in the number of vertices adjacent to v .

Proof.

First we prove that the set CD is indeed a connected dominating set for G . From the construction it follows that in each hexagon which contains vertices of G at least one vertex is assigned to CD . So CD is a dominating set for G . If in G a vertex v in a hexagon h is adjacent to a vertex v' in a hexagon h' , then the algorithm ensures that for two adjacent vertices v and v' with $v \in h$ and $v' \in h'$ it holds that $v \in CD$ and $v' \in CD$. So in G there is an edge between two vertices in two different hexagons if and only if in G restricted to CD there is an edge between these hexagons. It follows that CD is connected since G is connected.

6. Conclusion

In this work, we analyzed the role of locality in algorithms designed for **dominating set**, **connected dominating set**, **independent set**, and **vertex cover** in the framework of location-aware unit disk graphs. We achieved several significant milestones:

1. **Lower Bounds for Local Approximation Algorithms** We established the first-ever lower bounds for local approximation algorithms tackling these problems. These bounds are directly tied to the locality distance of the algorithms, providing foundational insights into the trade-offs between locality and approximation quality.
2. **Computational Power of Low-Locality Algorithms** For problems such as dominating set, independent set, and vertex cover, we demonstrated that a locality distance of just **one hop** suffices to guarantee a **constant approximation ratio**. However, for the connected dominating set problem, we proved that no constant ratio approximation algorithm exists with locality one. To address this, we designed an algorithm that achieves a constant approximation ratio with a locality distance of two.
3. **Open Gaps and Challenges** Despite these advancements, significant gaps remain between the lower bounds we derived and the approximation factors of our algorithms (for locality distances one and two). Closing these gaps remains an open challenge. A promising direction is to establish **tight bounds for the special case of unit line graphs**, which were extensively utilized in our lower-bound proofs. Extending these constructions to **unit disk graphs in two dimensions** may yield better lower bounds, providing a more comprehensive understanding of the problems on general unit disk graphs.
4. **Future Directions for Locality and Approximation** Another open problem is to explore the performance of local algorithms with **small locality distances $k > 1$** . It is plausible that even a slight increase in locality distance could significantly enhance approximation ratios. A related avenue is to investigate these problems on a **grid with integer coordinates** and nodes with arbitrary transmission radii r . This setup closely approximates unit disk graphs while ensuring that all derived lower bounds remain applicable to general unit disk graphs.
5. **Beyond Small Locality** For larger locality distances, the vast open challenge is to determine **tight bounds for the approximation ratios of local algorithms** with arbitrary locality distances k . These investigations could uncover new insights into the computational power and limitations of local algorithms, providing a deeper understanding of the fundamental constraints of locality in approximation.

By addressing these open challenges, we can bridge the gap between theory and practical implementations of algorithms in wireless networks, advancing our understanding of locality's impact on performance in distributed systems.

References:

1. Manasse, M.; McGeoch, L. A.; Sleator, D. Competitive algorithms for server problems. J. Algorithms 1990, 11, 208–230.
2. Koutsoupias, E.; Papadimitriou, C. On the k -server conjecture. J. ACM 1995, 42, 971–983.
3. Chrobak, M.; Karloff, H.; Payne, T. H.; Vishwanathan, S. New results on server problems. SIAM J. Discrete Math. 1991, 4, 172–181.
4. Bartal, Y.; Bollobas, B.; Mendel, M. A Ramsey-type theorem for metric spaces and its applications. Proc. 42nd FOCS 2001, 396–405.
5. Coppersmith, D.; Doyle, P. G.; Raghavan, P.; Snir, M. Random walks on weighted graphs. J. ACM 1993, 40, 421–453.
6. Karlin, A. R.; Kenyon, C.; Randall, D. Dynamic TCP acknowledgement and other stories. Proc. 33rd STOC 2001, 502–509.
7. Reingold, N.; Westbrook, J.; Sleator, D. D. Randomized competitive algorithms for the list update problem. Algorithmica 1994, 11, 15–32.
8. Fiat, A.; Karp, R.; Luby, M.; McGeoch, L. A.; Sleator, D.; Young, N. E. Competitive paging algorithms. J. Algorithms 1991, 12, 685–699.
9. Appel, K.; Haken, W. Every planar map is four colorable. Illinois J. Math. 1977, 21, 429–597.
10. Feige, U.; Goemans, M. X. Approximating the value of two prover proof systems. Proc. 3rd ISTCS 1995, 182–189.
11. Goemans, M. X.; Williamson, D. P. Improved approximation algorithms for satisfiability problems. J. ACM 1995, 42, 1115–1145.
12. Karloff, H.; Zwick, U. A $7/8$ -approximation algorithm for max 3SAT. Proc. 38th FOCS 1997, 406–417.
13. Trevisan, L.; Sorkin, G. B.; Sudan, M.; Williamson, D. P. Gadgets, approximation, and linear programming. SIAM J. Comput. 2000, 29, 2074–2097.
14. Seiden, S. S. On the online bin packing problem. J. ACM 2002, 49, 640–671.
15. Chrobak, M.; Larmore, L. L. The server problem and on-line games. DIMACS Series Discrete Math. Theor. Comput. Sci. 1992, 7, 11–64.

16. Suomela, J. Survey of local algorithms. *ACM Comput. Surv.* 2013, 45, 24:1–24:40.
17. Kuhn, F.; Moscibroda, T.; Wattenhofer, R. What cannot be computed locally! *Proc. 23rd PODC 2004*, 300–309.
18. Linial, N. Locality in distributed graph algorithms. *SIAM J. Comput.* 1992, 21, 193–201.
19. Lovász, L. On the Shannon capacity of a graph. *IEEE Trans. Inf. Theory* 1979, 25, 1–7.
20. Awerbuch, B.; Peleg, D. Sparse partitions. *Proc. 31st FOCS 1990*, 503–513.
21. Kuhn, F.; Wattenhofer, R. On the complexity of distributed graph coloring. *Proc. 25th PODC 2006*, 7–15.
22. Suomela, J. Locality and distributed computing. *Encyclopedia of Algorithms 2008*, 1–5.
23. Alon, N.; Spencer, J. H. *The Probabilistic Method*; Wiley: Hoboken, NJ, USA, 2000.
24. Har-Peled, S. *Geometric Approximation Algorithms*; AMS: Providence, RI, USA, 2011.
25. Clarke, N.; Colbourn, C. J. Unit disk graphs: Domination and independent sets. *Discrete Math.* 1998, 213, 55–72.
26. Kuhn, F.; Moscibroda, T.; Wattenhofer, R. Fault-tolerant clustering in wireless ad hoc networks. *Proc. 26th IEEE INFOCOM 2007*, 1–10.
27. Penrose, M. D. On k-connectivity for a geometric random graph. *Random Struct. Algorithms* 1999, 15, 145–164.
28. Hochbaum, D. S. Approximation algorithms for the set covering and vertex cover problems. *SIAM J. Comput.* 1982, 11, 555–556.
29. Shachnai, H.; Tamir, G. On two class-constrained versions of the minimum spanning tree problem. *Theor. Comput. Sci.* 2001, 295, 339–357.
30. Johari, R.; Tse, D. Distributed optimization in ad hoc networks. *IEEE/ACM Trans. Netw.* 2005, 13, 279–290.
31. Alon, N.; Bar-Noy, A.; Linial, N.; Peleg, D. A lower bound for radio broadcast. *J. Comput. Syst. Sci.* 1991, 43, 290–298.
32. Gandham, S. R.; Dawande, M.; Prakash, R. Energy-efficient schemes for wireless sensor networks. *IEEE Comput.* 2005, 38, 54–60.
33. Gupta, P.; Kumar, P. R. The capacity of wireless networks. *IEEE Trans. Inf. Theory* 2000, 46, 388–404.
34. Boukerche, A. *Algorithms and protocols for wireless sensor networks*. Wiley Series on Parallel and Distributed Computing 2008, 1–24.
35. Du, D.-Z.; Wan, P.-J. *Connected Dominating Sets in Wireless Ad Hoc Networks*; Springer: New York, NY, USA, 2012.
36. Estrin, D.; Govindan, R.; Heidemann, J.; Kumar, S. Next century challenges. *Proc. MobiCom 1999*, 263–270.
37. Johnson, D. B.; Maltz, D. A. Dynamic source routing in ad hoc wireless networks. *Mobile Computing* 1996, 5, 153–181.
38. Ramanathan, R.; Rosales-Hain, R. Topology control of multihop wireless networks. *Proc. 22nd INFOCOM 2000*, 404–413.
39. Kleinberg, J. M.; Rubinfeld, R. Short paths in expander graphs. *Proc. 31st STOC 1999*, 62–70.
40. Kuhn, F.; Wattenhofer, R.; Zollinger, A. Ad-hoc networks beyond unit disk graphs. *Proc. 1st DIALM-POMC 2003*, 69–78.
41. Schneider, J.; Wattenhofer, R. Distributed coloring depending on the chromatic number. *Proc. 24th PODC 2005*, 237–246.
42. Wan, P.-J.; Alzoubi, K. M.; Frieder, O. Distributed construction of connected dominating sets. *Proc. 21st INFOCOM 2002*, 1597–1604.
43. Dubhashi, D.; Panconesi, A. *Concentration of Measure for the Analysis of Randomized Algorithms*; Cambridge University Press: Cambridge, UK, 2009.
44. Nemhauser, G. L.; Trotter, L. E. Properties of vertex packing and covering. *Math. Program.* 1975, 8, 232–248.
45. Bar-Noy, A.; Peleg, D. Distributed priority scheduling and guaranteed end-to-end packet delays. *Proc. 14th PODC 1995*, 235–246.
46. Raghavan, P.; Thompson, C. D. Randomized rounding. *Combinatorica* 1987, 7, 365–374.
47. Kumar, R.; Rajaraman, R.; Sundaram, R. A polylogarithmic approximation for the group Steiner tree problem. *Proc. 30th STOC 1998*, 568–578.