

Fast multipole networks

Steve Huntsman¹

BAE Systems FAST Labs
 steve.huntsman@baesystems.com

Abstract. Two prerequisites for robotic multiagent systems are mobility and communication. *Fast multipole networks* (FMNs) enable both ends within a unified framework. FMNs can be organized very efficiently in a distributed way from local information and are ideally suited for motion planning using artificial potentials. We compare FMNs to conventional communication topologies, and find that FMNs offer competitive communication performance (including higher network efficiency per edge at marginal energy cost) in addition to advantages for mobility.

1 Introduction

A multirobot system [23] is a group of autonomous, networked robots. In order to achieve a complex goal such as swarming [6], the system requires distributed coordination of both mobility and communication, among other objectives. This is nontrivial, and “[e]fficient networking of many-robot systems is considered one of the grand challenges of robotics” [30]. The respective enabling technologies for mobility and communication are path planning and mobile *ad hoc* networks (MANETs) [27]. While networks are inevitably analyzed from the perspective of graph theory [2], path planning may be considered in either graph-theoretical [29] or continuous settings. Meanwhile, because geometrical considerations such as distance and motion strongly influence the structure of MANETs, it is natural to try to address mobility and communication for multirobot systems together, e.g., as in [37]. Much effort has focused on connectivity maintenance in situations where, e.g. multirobot systems maintain periodic connectivity [15] or communicate by physically meeting [19] while pursuing a motion objective, or maintain continuous connectivity relative to a fixed set of access points [11,17,18]. Additionally, co-optimization of communication and motion or coverage for an individual robot have been considered in [41,30]. More recently, tree-based approaches for connectivity maintenance have been considered in [24,28,39].

In this paper, we *assume* connectivity is possible (by using more energy if necessary) without any optimization, and we introduce a class of network backbones that can be trivially formed using an efficient local motion planning technique. These *fast multipole networks* (FMNs) to support both mobility and communication within a unified framework. The basic idea is to follow common practice in modeling robots, goals, and obstacles as (superpositions of) charged particles satisfying the Laplace equation $\nabla^2\phi = 0$ [7,21,34] and exploit the *fast multipole method* (FMM), an efficient algorithm for simulating particle dynamics

[3,4,12], to simultaneously determine a sparse network topology that supports efficient communication. The animating principle that the far-field behavior of point charges [16] should determine a communication topology is geometrically natural. More surprisingly, we shall demonstrate that it leads to network topologies that perform well in their own right, with higher network efficiency per edge (at marginal energy cost) than standard topologies that ignore mobility.

After briefly reviewing the artificial potential approach to path planning in §2 and the FMM in §3, we introduce FMNs in §4, and compare them to conventional MANET topologies in §5 before making concluding remarks in §6.

2 Artificial potentials

The use of artificial potentials in motion and path planning has a long history, most frequently identified as beginning with [20]. The basic idea is to design a potential ϕ such that the equation of motion $m\ddot{x} = -\nabla\phi$ results in a desired trajectory x . Towards this end, goals and obstacles are respectively modeled by attractive and repulsive terms contributing to the total potential ϕ . Depending on circumstances, we may choose to model the robots as “sources” with potentials of their own (e.g., to avoid collisions), or as passive “targets” that simply move along the gradient of an ambient potential.

In general, we might consider essentially arbitrary forms for each term to produce very detailed behavior. Alternatively, we might rely on a single simple form for all the terms. Our approach is in the latter vein. The relative strengths and spatial distribution of these terms are chosen to establish priorities, spatially extended features, etc. In order to represent sufficiently complex spatial relationships along these lines, it is helpful to have an algorithmic framework that scales better in total computational effort, parallelism, and locality than evaluating $O(N^2)$ interactions, since the number N of terms in the potential can be much larger than the number of robots involved.

Besides these computational concerns, a problem with using artificial potentials that was identified at an early stage is the possible presence of local minima in the potential field that can trap agents [22]. To remedy this by construction, the notion of a *navigation function* that has a single minimum at the goal was developed, along with algorithms for constructing such functions [36]. A particularly simple way to avoid local minima while using a single form for all the potential terms is with a superposition of harmonic potentials [7,21,34], i.e., solutions to the Laplace equation $\nabla^2\phi = 0$, with a dominant term at the goal.

This is most readily achieved through a discrete (if perhaps quasi-continuous) superposition of *point charges*, i.e. potentials of the form $-qV(|x - x_0|)$ (the sign is for physical reasons), where the *fundamental solution* $V(|x|)$ to the Laplace equation is defined by $\nabla^2V(|x|) = \delta(x)$, and as usual δ indicates the Dirac delta distribution [38]. For \mathbb{R}^d , it turns out that $V'(r) = 1/A_d(r)$, where $A_d(r)$ is the Minkowski content (i.e., generalized perimeter, surface area, etc.) of the sphere of radius r in \mathbb{R}^d . Choosing the most convenient constants of integration, for $d = 2$ we have $V(r) = \frac{1}{2\pi} \log r$, and for $d = 3$ we have $V(r) = -1/4\pi r$.

3 The fast multipole method

Naive simulation of N interacting point charges (e.g., the goals and obstacles modeled in Figure 1) requires computing the interactions of each pair of charges, and hence $O(N^2)$ operations per time step, which is prohibitive for large-scale N -body simulations. The FMM [3,4,12] enables the simulation cost to be reduced to $O(N)$ with an extremely high degree of locality and parallelism [13].¹

The key ideas underlying the FMM are

- i) a specification of accuracy (for truncating expansions in a controlled way);
- ii) decomposing space hierarchically to get *well-separated* charge clusters;²
- iii) representing well-separated clusters of point charges with multipole expansions that maintain a desired approximation error ε with as few ($\lceil \log_2(1/\varepsilon) \rceil$) terms as possible, leaving nearby particles to interact directly.

In particular, the FMM recursively builds a quad-tree (Figure 1; in three dimensions, an octo-tree is used instead) whose leaves are associated with boxes and truncated multipole expansions. This tree approximates a (typically much) finer tree whose leaves are associated with individual point charges that are well-separated and their monopoles. Importantly, the FMM tree topology essentially ignores the values of charges, depending only on the desired level of accuracy ε ³ and the locations of the charges.

The computationally expedient part of the FMM is to manipulate the origins and coefficients of controlled series approximations to far-field potentials for clusters of point charges that are well-separated. More general incarnations of the FMM (see, e.g., [26,42,43]) amount to a very efficient scheme for computing sums of the form $\sum_{j=1}^N K(x_i, \xi_j)\psi(\xi_j)$ for a given *kernel* K : i.e., the FMM and its generalizations are essentially specialized matrix multiplication algorithms. From this perspective, item iii) in the list above separates into [3]

- a far-field expansion of the kernel $K(x, \xi)$ that decouples the influence of the evaluation/target point x and the source point ξ ;
- (optionally) a conversion of far-field expansions into local ones.

The FMM's remarkable scaling performance has enabled petascale simulations of turbulence [45], molecular dynamics [32], and cosmological dynamics [35], and will also enable future exascale simulations across hundreds of thousands of nodes [44]. This performance makes the FMM a natural choice for large scale path planning using artificial potentials.

¹ For the calculations in this paper, we used the very user-friendly library FMM-LIB2D, available at <https://cims.nyu.edu/cmcl/fmm2dlib/fmm2dlib.html>.

² Two clusters of points $\{x_j\}$ and $\{y_k\}$ are well-separated iff there exist x_0, y_0 such that $\{x_j\} \subset B_{x_0}^\circ(r)$ and $\{y_k\} \subset B_{y_0}^\circ(r)$ with $|x_0 - y_0| > 3r$: here $^\circ$ denotes interior. Two squares with side length r are well-separated iff they are at distance $\geq r$.

³ Though in principle the desired level of accuracy can be affected by charge values, this situation is sufficiently pathological that we can safely disregard it in practice.

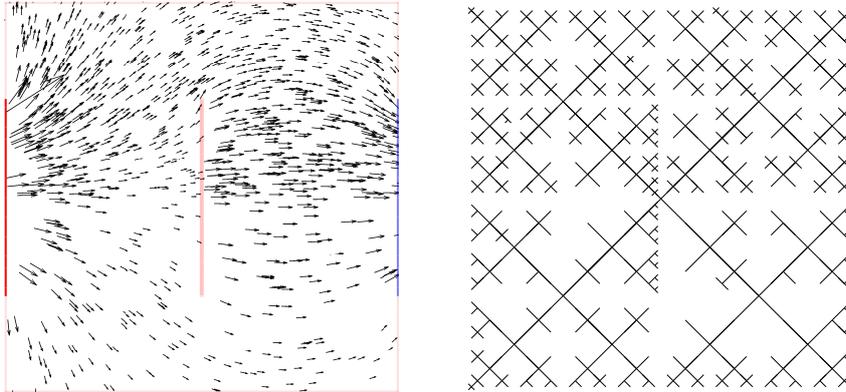


Fig. 1. (L) A toy scenario in $[-1, 1]^2$. Goals are modeled by negative charges and shown in blue; obstacles are modeled by positive charges and shown in red. Opacity indicates relative magnitude. 10^3 robots are modeled by test points (versus, e.g., test charges of small positive sign) and their locations and velocities indicated by black gradient vectors of the artificial potential. The target locations are distributed as $\frac{4}{5}U$ (top half) + $\frac{1}{5}U$ (bottom half), where here U indicates a uniform distribution. (R) The quad-tree associated to the scenario on the right. Varying the desired precision in the FMM has very little effect on this tree; as a practical matter it can be assumed unique.

Equally important for the considerations of this paper, however, are the hierarchical and spatial locality properties that the FMM exploits in order to communicate internally. The FMM’s patterning of a logical intra-algorithm communication network after the spatial distribution of particles suggests that it can be used not only for large-scale multirobot path planning in complex geometries, but also to help organize the communications between robots in a distributed way. Furthermore, although the FMM’s hierarchical properties might seem to imply centralization, the computational load is small enough that these functions can be easily duplicated among robots with low overhead, i.e., the FMM tree does not impose centralization.

4 Fast multipole networks

We construct the *fast multipole network* $FMN(\xi)$ corresponding to a configuration of points $\xi_j \in \mathbb{R}^2$ as follows: vertices correspond to the charge locations and we introduce edges that

- connect all vertices in the same FMM leaf box;
- connect nearest vertices in adjacent leaf boxes;
- connect otherwise isolated vertices to their nearest neighbors.

These edges are respectively colored blue, cyan, and red in Figure 2.⁴

By construction, $FMN(\xi)$ is connected, and the information required to generate it is automatically produced by the FMM. We note that while $FMN(\xi)$ is constructed using the quad- or octo-tree of the FMM, it is very far from a tree. Rather, the FMM tree and its corresponding coarse-graining of space determines which nodes are *permitted* to communicate directly.⁵ Within a clique of permitted communications corresponding to a leaf of the FMM tree, we may further restrict communications to avoid quadratic bandwidth overhead and/or energy, though we do not consider such tactics further here.

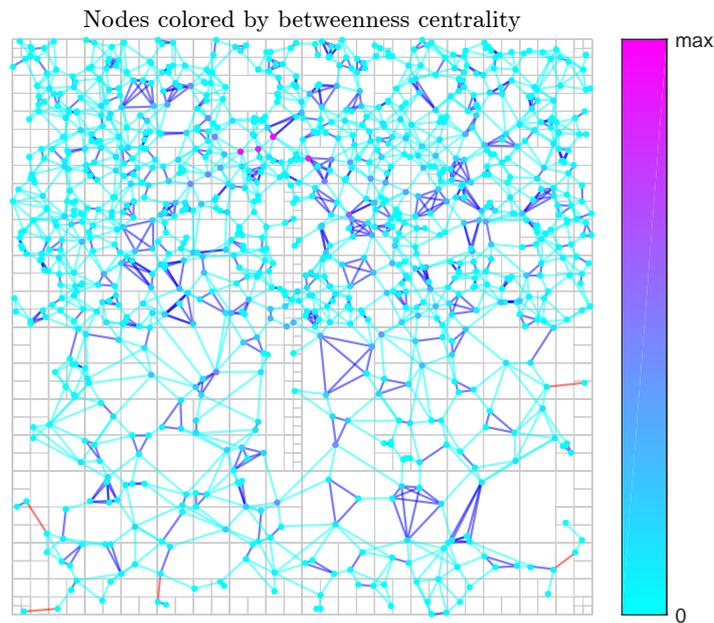


Fig. 2. The FMN corresponding to the scenario in Figure 1. Nodes are colored by betweenness centrality according to the colorbar on the right. The spatial decomposition from Figure 1 is shown in gray for reference. Edges within a FMM box are blue, while edges connecting nearest nodes in adjacent boxes are cyan and edges connecting otherwise isolated nodes to their nearest neighbors are red.

⁴ The key difference between FMNs and the networks considered in [48] is that the latter are formed by inserting and permanently linking nearby charges, then dynamically evolving to obtain small-world features, whereas FMNs are (re)formed by linking nearby charges in a way that partially anticipates the next timestep of dynamical evolution. However, both types of networks exhibit aspects of small-world behavior (see §5 and [25]).

⁵ Limiting permission for direct communication in FMNs can be enforced by, e.g., cognitive radios [46] whose spectrum allocation cooperates with the FMM tree.

5 Evaluation

We now introduce several families of graphs for evaluation purposes.

Let $\xi_j \in \mathbb{R}^2$ for $1 \leq j \leq N$, and let $r > 0$. The *random geometric graph* or *disk graph* (RGG; Figure 3) $RGG(\xi; r)$ has vertices ξ_j and edges $E(RGG(\xi; r)) := \{(\xi_j, \xi_k) : d(\xi_j, \xi_k) \leq r\}$ [14,33]. By construction, a RGG is both the most effective network topology from the point of view of information exchange, and the least effective network topology from the point of view of infrastructure costs.

A more conservative topology is based on subgraphs of the *Delaunay graph*. The Delaunay graph $D(\xi)$ has vertices ξ_j and edges defined from a triangulation of the vertices such that no vertex is interior to a circle circumscribed about a triangle [5,9,10].⁶

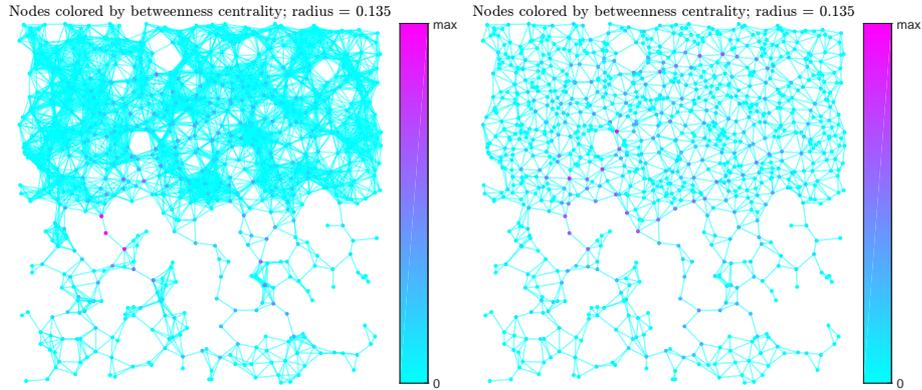


Fig. 3. (L) $RGG(\xi; r)$ for ξ corresponding to the scenario in Figure 1 and $r = 0.135$, slightly above the threshold for connectivity. (R) $RD(\xi; r)$.

The *Gabriel graph* $G(\xi)$ [29,31] is the unique (for the general position case) subgraph of the Delaunay graph such that each edge corresponds to the diameter of a disk that does not contain any other vertices; it is frequently considered as a potential candidate for “virtual backbones” in MANETs. It is worth noting however that $G(\xi)$ and $D(\xi)$ are more computationally expensive to construct than $FMN(\xi)$, and parallelism does not change this.

Because the Delaunay and Gabriel graphs do not have an intrinsic range parameter that will give a granular mechanism for evaluating their performance, we shall focus our attention on the (*minimal*) *restricted Delaunay graph* (Figure 3) $RD(\xi; r) := D(\xi) \cap RGG(\xi; r)$ [1] and the *restricted Gabriel graph* (Figure 4) $RG(\xi; r) := G(\xi) \cap RGG(\xi; r)$. Similarly, we shall consider the *restricted FMN* (Figure 4) obtained along the lines $RFMN(\xi; r) := FMN(\xi) \cap RGG(\xi; r)$.

⁶ For ξ_j in general position, the Delaunay graph is unique.

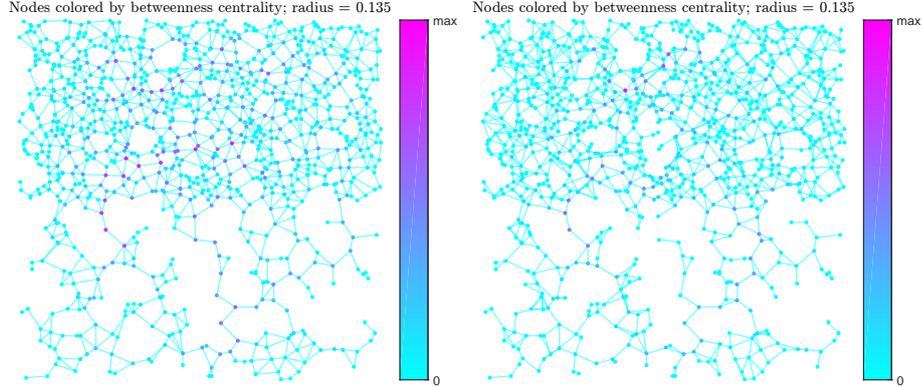


Fig. 4. (L) $RG(\xi; r)$ for ξ corresponding to the scenario in Figure 1 and $r = 0.135$, slightly above the threshold for connectivity. (R) $RFMN(\xi; r)$.

The basic evaluation metric we use is the *efficiency* of a graph $G = (V(G), E(G))$, defined as the average inverse distance between distinct vertices, i.e.

$$\text{eff}(G) := \binom{|V(G)|}{2}^{-1} \sum_{\substack{j, k \in V(G) \\ j \neq k}} \frac{1}{d_{jk}}, \quad (1)$$

where the distance d_{jk} between vertices j and k is computed in the obvious way from a given distance on edges (by default, we may always choose the *hop metric* assigning 1 to each edge). While the efficiency characterizes how well a network supports information flow [25], it neglects costs (e.g., bandwidth, energy, etc.) associated to edges as infrastructure. For this reason we will also consider the *efficiency per edge*, i.e. $\text{eff}(G)/|E(G)|$. Although other normalizations may be more appropriate in certain situations, this particular one strikes a good balance between convenience/generalality and detail, especially for the hop metric.

Figure 5 shows the metrics above for 100 simulations of 10^3 uniformly distributed test points in $[-1, 1]^2$ subject to the ambient potential from Figure 1. It is apparent from the figure that FMNs and their range-restricted versions are worthy candidates for network backbones in their own right even before accounting for their mobility-specific advantages. Furthermore, although there exist efficient local and parallel algorithms for constructing Delaunay graphs [5,9,10], their computation and communication complexity and scaling behavior are still inferior to the FMM.

Figure 6 shows metrics relating to degree distributions and *efficiency per unit energy*, i.e., $\text{eff}(G)/\text{energy}_\bullet(G)$, where (ignoring an irrelevant constant of proportionality) the unidirectional energy for a metric graph G is

$$\text{energy}_{\text{uni}}(G) := \sum_{(j,k) \in E(G)} d_{jk}^2 \quad (2)$$

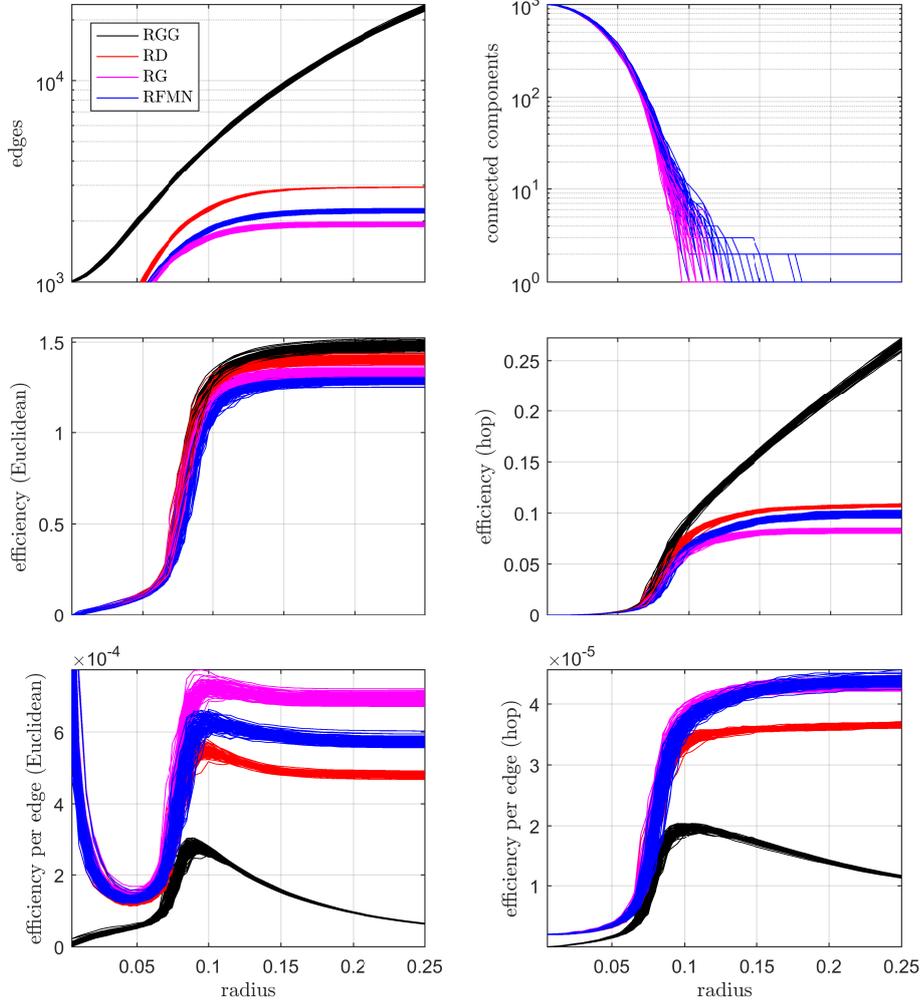


Fig. 5. Network metrics for $RGG(\xi; r)$ (in black), $RD(\xi; r)$ (in red), $RG(\xi; r)$ (in magenta), and $RFMN(\xi; r)$ (in blue) for 100 simulations of $N = 10^3$ uniformly distributed test charges in $[-1, 1]^2$. Although $RGG(\xi; r)$ is most efficient, this network performance comes at a high cost in edges, and $RFMN(\xi; r)$ performs well (and for hop efficiency per edge, the best) for all measures of efficiency. Note that $RFMN(\xi; r) = FMN(\xi)$ for sufficiently large r within the range shown. We also have that, e.g. $RD(\xi; r') = D(\xi)$, and though the corresponding r' is outside the range shown, the residual effects are minimal.

and the omnidirectional energy is

$$\text{energy}_{\text{omni}}(G) := \sum_{j \in V(G)} \left(\max_{\substack{k \in V(G) \\ (j,k) \in E(G)}} d_{jk} \right)^2. \quad (3)$$

These quantities model the total energy budgets required to transmit uni- and omnidirectional signals, respectively. Figure 6 highlights that FMNs continue to perform marginally better than Delaunay graphs and marginally worse than Gabriel graphs for energy-normalized measures of network efficiency.

6 Remarks

By virtue of calculating potentials and forces, the FMM/FMN approach enables dynamic and predictable network topology reconfiguration with minimal cost and effort. In other words, as robots use the FMM to efficiently compute their motion according to a navigation function supplied by a superposition of point charges, the FMN is easily updated and efficiently represented.

Incorporating resilient routing reconfiguration [40,8] on $(F + 1)$ -connected local subgraphs of the FMN can be done with reasonable computational effort (e.g., the key linear program is quickly and easily solved in MATLAB for realistic networks of ≈ 50 nodes). This enables virtually instantaneous failover and rerouting in the presence of $\leq F$ link failures. Combining this local approach with a separate (perhaps similar) routing protocol to handle wide-area network traffic and obstacle potentials that prevent deterioration of basic connectivity can ensure network integrity and basic quality of service (QoS). These features can render our framework competitive with the approach of [37], which centers on the higher-level functions of network integrity and QoS, and which uses a convex program instead of an algorithmically simpler linear program. Along similar lines, [47] shows how to construct artificial potentials that discourage loss of connectivity. Although these fields are not harmonic, it is plausible that this idea can be adapted to the present context.

It is worth pointing out that there are FMM variants for non-harmonic potentials, e.g. power laws, (generalized) multiquadratics [3,43], or more general kernels [26], and many of these have actually been applied in the context of interpolation and/or physical simulation. However, using non-harmonic potentials eliminates the automatic guarantee that there are no metastable local minima. We note in particular that the kernel-independent FMM variant of [42] exploits the existence and uniqueness of solutions to elliptic boundary value problems [38] to represent clustered sources in far field based on their behavior on a suitable boundary. This perspective suggests an extension of FMNs to sources modeled by fundamental solutions of elliptic partial differential equations.

Acknowledgements

We thank Brendan Fong, Marco Pravia, and David Spivak for their comments.

References

1. Avin, C. “Fast and efficient restricted Delaunay triangulation in random geometric graphs.” *Internet Math.* **5**, 195 (2008).
2. Barrat, A., Barthélemy, M., and Vespignani, A. *Dynamical Processes on Complex Networks*. Cambridge (2008).
3. Beatson, R. and Greengard, L. “A short course on fast multipole methods.” In *Wavelets, Multilevel Methods, and Elliptic PDEs*. Ainsworth, M., *et al.*, eds. Oxford (1997).
4. Board, J. and Schulten, L. “The fast multipole algorithm.” *Comp. Sci. Eng.* **2**, 76 (2000).
5. Chen, R. and Gotsman, C. “Localizing the Delaunay triangulation and its parallel implementation.” ISVD (2012).
6. Chung, S.-J., *et al.* “A survey on aerial swarm robotics.” *IEEE Trans. Robotics* **34**, 837 (2018).
7. Connolly, C. I., Burns, J. B. and Weiss, R. “Path planning using Laplace’s equation.” ICRA (1990).
8. DeCleene, B. and Huntsman, S. “Wireless resilient routing reconfiguration.” arXiv:1904.04865 (2019).
9. Fuetterling, V., Lojewski, C., and Pfreundt, F.-J. “High-performance d -D Delaunay triangulations for many-core computers.” HPG (2014).
10. Funke, D. and Sanders, P. “Parallel d -D Delaunay triangulations in shared and distributed memory.” ALENEX (2017).
11. Ghaffarkhah, A. and Mostofi, Y. “Communication-aware motion planning in mobile networks.” *IEEE Trans. Auto. Control* **56**, 2478 (2011).
12. Greengard, L. and Rokhlin, V. “A fast algorithm for particle simulations.” *J. Comp. Phys.* **73**, 325 (1987).
13. Greengard, L. and Gropp, W. D. “A parallel version of the fast multipole method.” *Comp. Math. Appl.* **20**, 63 (1990).
14. Haenggi, M. *Stochastic Geometry for Wireless Networks*. Cambridge (2013).
15. Hollinger, G. and Singh, S. “Multi-robot coordination with periodic connectivity.” ICRA (2010).
16. Jackson, J. D. *Classical Electrodynamics*. 3rd ed. Wiley (1998).
17. Kantaros, Y. and Zavlanos, M. M. “Distributed communication-aware coverage control by mobile sensor networks.” *Automatica* **63**, 209 (2016).
18. Kantaros, Y. and Zavlanos, M. M. “Global planning for multi-robot communication networks in complex environments.” *IEEE Trans. Robotics* **32**, 1045 (2016).
19. Kantaros, Y., Guo, M., and Zavlanos, M. M. “Temporal logic task planning and intermittent connectivity control of mobile robot networks.” *IEEE Trans. Auto. Control* **64**, 4105 (2019).
20. Khatib, O. “Real-time obstacle avoidance for manipulators and mobile robots.” ICRA (1985).
21. Kim, J.-O. and Khosla, P. K. “Real-time obstacle avoidance using harmonic potential functions.” *IEEE Trans. Robotics and Automation* **8**, 501 (1992).
22. Koren, Y. and Borenstein, J. “Potential field method and their inherent limitations for mobile robot navigation.” ICRA (1991).
23. Knorn, S., Chen, Z., and Middleton, R. H. “Overview: collective control of multi-agent systems.” *IEEE Trans. Cont. Net. Sys.* **3**, 334 (2016).
24. Krupke, D., *et al.* “Distributed cohesive control for robot swarms: maintaining good connectivity in the presence of exterior forces.” IROS (2015).

25. Latora, V. and Marchiori, M. "Efficient behavior of small-world networks." *Phys. Rev. Lett.* **87**, 198701 (2001).
26. Létourneau, P.-D., Cecka, C., and Darve, E. "Cauchy fast multipole method for general analytic kernels." *SIAM J. Sci. Comp.* **36**, A396 (2014).
27. Loo, J., Mauri, J. L., and Ortiz, J. H., eds. *Mobile Ad Hoc Networks*. CRC (2016).
28. Majcherczyk, N., *et al.* "Decentralized connectivity-preserving deployment of large-scale robot swarms." IROS (2018).
29. Mesbahi, M. and Egerstedt, M. *Graph Theoretic Methods in Multiagent Networks*. Princeton (2010).
30. Minelli, M., *et al.* "Stop, think and roll: online gain optimization for resilient multi-robot topologies." DARS (2019).
31. Norrenbrock, C. "Percolation threshold on planar Euclidean Gabriel graphs." *Eur. Phys. J. B* **89**, 111 (2016).
32. Ohno, Y., *et al.* "Petascale molecular dynamics simulation using the fast multipole method on K computer." *Comp. Phys. Comm.* **185**, 2575 (2014).
33. Penrose, M. *Random Geometric Graphs*. Oxford (2003).
34. Pimenta, L. C. A., *et al.* "On computing complex navigation functions." ICRA (2005).
35. Potter, D., Stadel, J., and Teyssier, R. "PKDGRAV3: beyond trillion particle cosmological simulations for the next era of galaxy surveys." *Comp. Astrophys. Cosmol.* **4**, 2 (2017).
36. Rimon, E. and Koditschek, D. E. "Exact robot navigation using artificial potential functions." *IEEE Trans. Robotics and Automation* **8**, 501 (1992).
37. Stephan, J., *et al.* "Concurrent control of mobility and communication in multi-robot systems." *IEEE Trans. Robotics* **33**, 1248 (2017).
38. Taylor, M. E. *Partial Differential Equations: Basic Theory*. Springer (1996).
39. Varadharajan, V. S., Adams, B., and Beltrame, G. "The unbroken telephone game: keeping systems connected." AAMAS (2019).
40. Wang, Y. *et al.* "R3: resilient routing reconfiguration." SIGCOMM (2010).
41. Yan, Y. and Mostofi, Y. "Co-optimization of communication and motion planning of a robotic operation in fading environments." ACSSC (2011).
42. Ying, L., Biro, G., and Zorin, D. "A kernel-independent adaptive fast multipole algorithm in two and three dimensions." *J. Comp. Phys.* **196**, 591 (2004).
43. Ying, L. "A kernel-independent fast multipole algorithm for radial basis functions." *J. Comp. Phys.* **213**, 457 (2006).
44. Yokota, R. and Barba, L. A. "A tuned and scalable fast multipole method as a preeminent algorithm for exascale systems." *Int. J. High Perf. Comp. Appl.* **26**, 337 (2012).
45. Yokota, R., *et al.* "Petascale turbulence simulation using a highly parallel fast multipole method on GPUs." *Comp. Phys. Comm.* **184**, 445 (2013).
46. Yu, R. F., ed. *Cognitive Radio Mobile Ad Hoc Networks*. Springer (2011).
47. Zavlanos, M. M. and Pappas, G. J. "Potential fields for maintaining connectivity of mobile networks." *IEEE Trans. Robotics* **23**, 812 (2007).
48. Zitin, A. *et al.* "Spatially embedded growing small-world networks." *Sci. Rep.* **4**, 7047 (2015).