

Visualization of Resource Allocation in Large-Scale Mobile Ad Hoc Networks

Alex Fridman, Dan Hennessey, David Breen, Steven Weber, and Moshe Kam

Drexel University
3141 Chestnut Street
Philadelphia, PA 19104
{af59@dph290,david@cs.,sweber@ece.,kam@ece.}@drexel.edu

Abstract. Resource allocation in ad hoc communication networks is a field of high complexity because of both *i*) the distributed nature of the interactions between the nodes, and *ii*) the large set of control variables for even the most primitive networks. Visual representation of this information across physical space and across layers of the network can greatly benefit the understanding of efficient allocation policies in these complex systems. Yet, very few software packages have been developed to address specifically this task, especially for large scale networks. We develop such a software system, and demonstrate some of its capabilities. The system illustrates that multi-layered visualization of the network state can be an effective tool at making network design decisions even in the face of uncertainty and for a large number of network interactions.

1 Introduction

A mobile ad hoc network is generally unstructured and highly dynamic, which makes the process of its visualization particularly important and challenging. In large-scale networks, the amount of resources allocated in a localized fashion is immense and requires graph drawing tools that optimize the trade-off between information density and visual clarity. We propose an abstraction of mobile ad hoc networks that achieves a reasonable balance between these two objectives. The four layers we define are: the mobility layer, the physical (PHY) layer, the medium access control (MAC) layer, and the transport layer. The mobility layer models the movement kinematics of the nodes as well as their movement plans. The PHY layer models the lowest level of the communication medium which includes simultaneous transmissions, radiation patterns, interference, signal quality, and capacity. The MAC layer models the temporal scheduling of node transmission and the aggregate channel state produced by a sequence of physical states. Finally, the transport layer models the flow of information on the network given the aggregate state computed at the MAC layer.

In this paper we *i*) present a multi-layer mobile ad hoc network model, and *ii*) discuss how each of the layers of that model can be visualized to present the essential information about the allocation of network resources. Although there

is both significant flexibility in parameter selection and significant model complexity, the visualization model we propose is sufficiently general to allow useful visual presentation of the impact of parameter selection on network performance.

2 Related Work

The technical literature on the visualization of ad hoc networks is significantly limited. To the best of our knowledge, existing work is restricted to small-scale examples that focus more on the attainment of the data that is visualized and less of the clarity, quality, and richness of information contained in the visual output [1]. Connectivity is a critical metric for performance of ad hoc networks, and is easy to visualize, even as the size of the network scales up. Therefore, most larger scale visual information processing efforts in the networking and visualization literature are focused on connectivity, such as in [2,3,4,5]. On the other hand, tools that allow for the visualization of other aspects of the communication medium can only handle small scale networks, both in terms of running-time complexity and visual complexity; [6,7,8] are popular examples of such tools.

One of the most active areas of research in visualization of networks approaches the problem from the perspective of graph theory. Problems addressed in these efforts are concerned with drawing graphs while minimizing the number of intersecting edges, pruning subgraphs based on predefined priorities, and generally maximizing the amount of visually conveyed information without sacrificing clarity and visual appeal. The popular GraphViz [9] tool applies a wide selection of cutting-edge research from this graph theoretic approach. While many graph simplification and layout algorithms are used in our tool, the focus and novelty of the visualization methods we present are in the effective visual representation of network resources commonly allocated in ad hoc networks.

3 Mobility Layer

3.1 Positions and Obstacles

The mobility layer contains the lowest level of information about the communication network. It specifies the current position, velocity, and acceleration of each node $i = 1, 2, \dots, n$ in the network at each snapshot in time. Figure 1(a) shows a snapshot of the node and obstacle positions for a large ad hoc network. Each small circle in the figure designates a node, and each large shaded polygon designates an obstacle. Each of the obstacles define the closed sections of the arena that can not be traversed by any of the nodes in the network. This visual information, in its animated form, can be used to monitor in real-time the changes in the network topology.

Figure 1(b) shows the aggregation ability of our visualization engine. A variety of spatial metrics can be gathered and plotted as a contour map, or in 3D with the z -axis containing the value of the metric. In this case, we use a contour map to show the visitation frequency of the nodes to each part of the network arena.

Light and dark areas reflect frequent and infrequent visitation, respectively. The black areas correspond to the obstacle locations. For simplicity, we used a random walk mobility model. This aggregate visual information can be used to study the movement patterns of a chosen mobility model.

A large-scale network of 10,000 nodes is shown in Figure 1(c). Specifically, the figure displays four snapshots of this network in time. The first snapshot is at the start of a simulation where the node topology is uniformly distributed over the (obstacle-free) arena. The next three snapshots capture the changing mobility graph state as the simulation progresses. The mobility model here is “follow the leader”, where ten of the nodes are chosen as leaders, and the rest of the nodes are followers. A follower constantly switches between the state of following a leader and the state of independent movement. The leaders and the followers without a leader perform a random walk. The animated version of what is shown in Figure 1(c) clearly shows that increasing the “following probability” increases the rate of convergence towards a state where small clusters of nodes move around the arena with minimal spreading. If the “following probability” is low enough, the topology dynamics are indistinguishable from one where all nodes perform a random walk.

3.2 Kinematics Graph

The kinematics graph $G_K = (V_K, E_K)$ specifies traversability information imposed by the presence of obstacles, boundaries, and kinematic restrictions in an arena. The vertices in V_K each denote a spatial location that can be occupied by a network node. The edges in E_K indicate feasible movements in the arena. Specifically, $\Lambda(i) = \{j : (i, j) \in E_K\}$ is the set of locations directly accessible from location i . Each step has an associated travel time of one time unit.

For the problem of motion planning, the kinematics graph most often takes the structured form of a lattice [10] or the unstructured form of a graph with distance-limited edges [11] generated by randomly sprinkling points V_K in the arena. The visualization of both forms is shown in Figure 1(d) and Figure 1(e), respectively. This visual information is effective for studying the chosen domain of the path planning problem for a discretized space. For example, the figures indicate that the lattice can capture the intricacies of the motion constraints in the physical space with fewer points than required by the unstructured graph, and is therefore a better choice for space discretization.

4 Physical and Medium Access Control Layers

We use the abstraction of the physical (PHY) and medium access control (MAC) layers defined in [12]. Medium access is controlled by defining a transmission schedule, consisting of a sequence of time slots. The schedule specifies a set of concurrent transmissions to occur in each time slot. The concurrent transmissions have the properties that *i*) all transmissions are successful, *ii*) adding an additional transmitter to any time slot would cause one or more transmissions

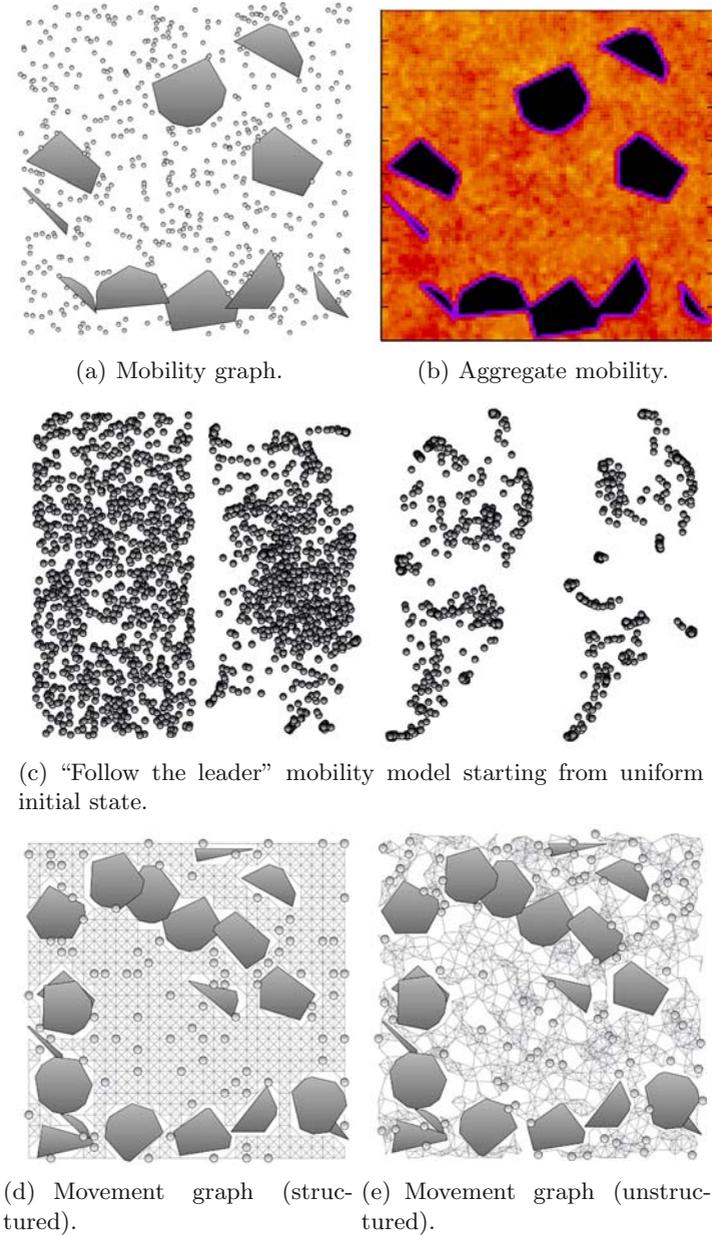


Fig. 1. Mobility visualization

to fail, and *iii*) the number of time slots is sufficient to ensure that each node is allowed to transmit in at least one time slot per round. The random packing procedure in [12] is employed to identify a random schedule satisfying these properties. The left part of Figure 2 shows the set of concurrent transmissions

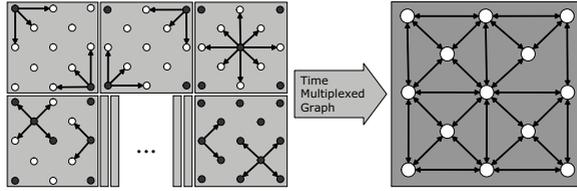


Fig. 2. The process of forming a capacitated graph from a set of schedules

at some of the time slots; the right part of the figure shows the *flow graph* $G_F = (V_F, E_F)$ formed by multiplexing the concurrent transmissions together.

The result of the packing procedure is a collection of graphs

$$\mathbf{G}_P = \{G_P^1, \dots, G_P^m\}, \quad G_P^z = (V_P^z, E_P^z), \quad z = 1, \dots, m, \quad (1)$$

where the graphs indicates concurrent transmitter receiver pairs at each time slot in the schedule. Granting preference to nodes that have not yet had the opportunity to transmit in a previous time slot ensures $m \leq n$.

The parameters of the PHY layer are *i*) the fixed transmission power P_i on each node i , and *ii*) the minimum acceptable signal-to-interference-plus-noise (SINR) ratio, γ . Given these values, the random packing procedure adds transmitters to a time slot incrementally such that the SINR requirement, γ , is not violated for any of the links already selected for the time slot. The SINR from transmitter i to receiver j in time slot z is given by:

$$\text{SINR}_{ij}(z) = \frac{P_i A_{ij}}{\sum_{k \in V_P^z \setminus i} P_k A_{kj} + \eta}, \quad (i, j) \in E_P^z, \quad z = 1, \dots, m. \quad (2)$$

Here, η is the channel noise power. The channel attenuation between transmitter i and receiver j is:

$$A_{ij} = \left(\frac{d_{\text{ref}}}{d_{ij}} \right)^\alpha, \quad (3)$$

where d_{ij} is the distance between nodes i and j , d_{ref} is a channel reference distance, and α is the attenuation exponent. All edges (i, j) in each graph G_P^z satisfy the requirement that $\text{SINR}_{ij}(z) > \gamma$. The capacity on each edge (i, j) at each time slot, $C_{ij}(z)$, may be assigned to be unity, or may be selected using the Shannon capacity:

$$C_{ij}(z) = B_{ij} \log_2 (1 + \text{SINR}_{ij}(z)), \quad (i, j) \in E_P^z, \quad z = 1, \dots, m. \quad (4)$$

The multiplexing procedure is additive for capacity, that is, the capacity of an edge (i, j) in the flow graph G_F is computed from \mathbf{G}_P by:

$$C_{ij} = \sum_{z=1}^m C_{ij}(z). \quad (5)$$

The result of the multiplexing procedure is a capacitated flow graph $G_F = (V_F, E_F)$ that has an edge $(i, j) \in E_F$ with capacity C_{ij} if and only if $(i, j) \in E_P^z$ for one or more $z = 1, \dots, m$.

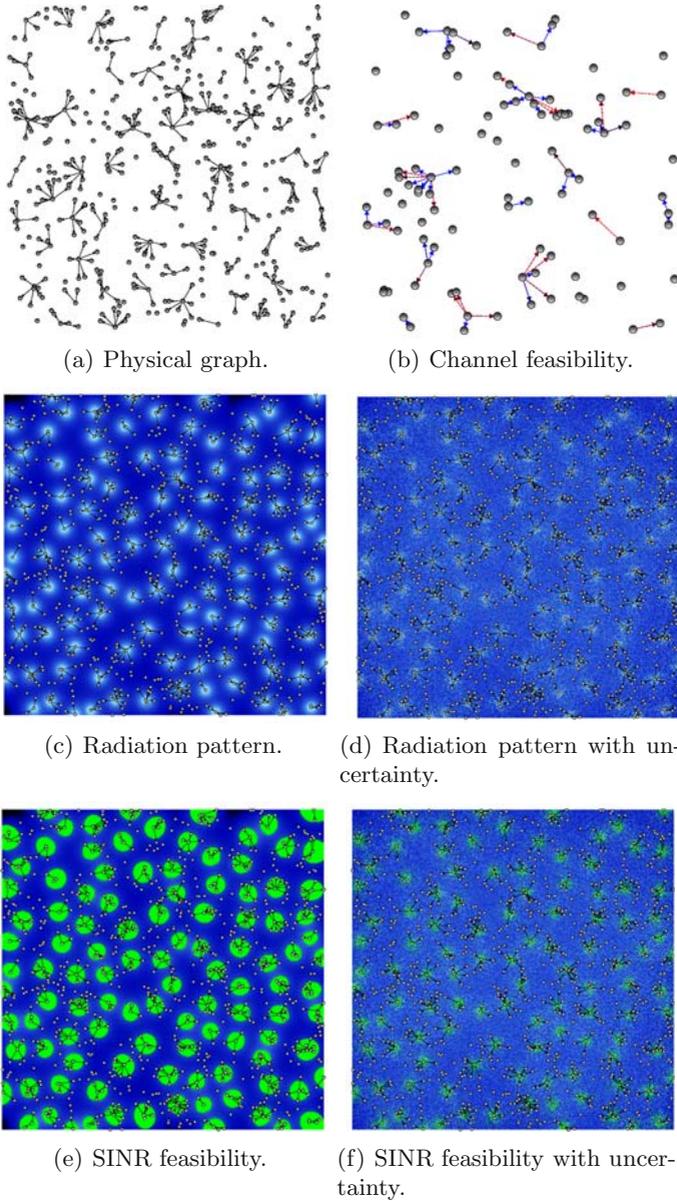


Fig. 3. Physical layer visualization

The visualization of a physical layer snapshot is done by drawing arrows outgoing from each active transmitter to each of its desired receivers. An example is shown in Figure 3(a). The capacity of each edge, $C_{ij}(z)$ can effectively be visualized with color, line thickness, and line dash patterns. Figure 3(b) demonstrates the use of color and line dashes to distinguish between feasible and infeasible edges. Specifically, the solid lines indicate feasible edges, and dashed lines

indicate infeasible (but nearly feasible) edges. This visualization technique is extremely effective at identifying nodes in the network that need to adjust their power either by increasing (in order to be better heard) or completely shutting itself off (to reduce total interference of that physical state).

Figure 3(c) and Figure 3(d) show the signal radiation patterns from transmitting nodes under two different models of the channel. The former illustrates a channel state that is assumed to be fully known with no fading, shadowing, or noise in the computation of the attenuation. The latter illustrates a channel state that is highly unstable, modeling uncertainty in the attenuation exponent, shadowing, fading, node positions, and noise. These two visual presentations of the physical state help demonstrate that it is much more difficult to make resource allocation decisions under uncertainty.

Figure 3(e) and Figure 3(f) show bright green regions where the SINR is acceptably high for a feasible channel to exist. That is, if a transmitter were to move into one of those regions, it would be able to transmit a signal to the receiving node associated with that region. Once again, the presence of uncertainty makes it difficult to design node movement and scheduling such that the receiver nodes can remain in the feasible regions of the transmitter with whom they seek to communicate.

The formation of a multiplexed MAC graph from a set of physical layer graphs is a powerful aggregate abstraction of resource allocation. A lot of visual complexity is removed in performing this abstraction while not losing much of the critical network state information. Figure 4 shows a subset of the physical layer schedules generated by a random packing scheduling algorithm from [12] on the left. On the right is a multiplexed MAC graph which is formed from

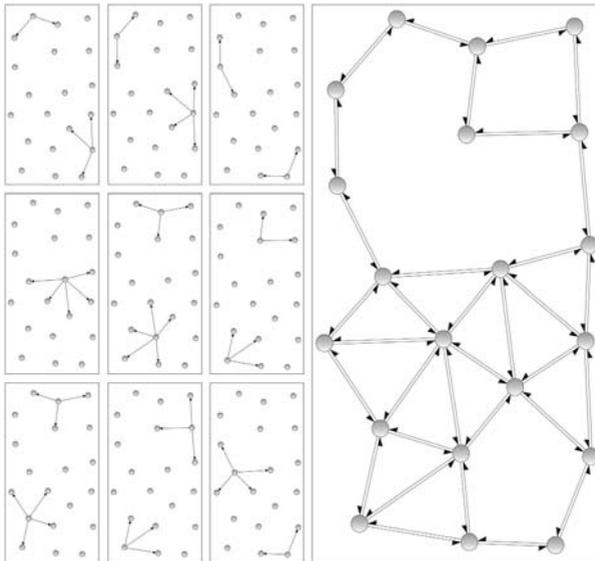


Fig. 4. Multiplexed MAC graph

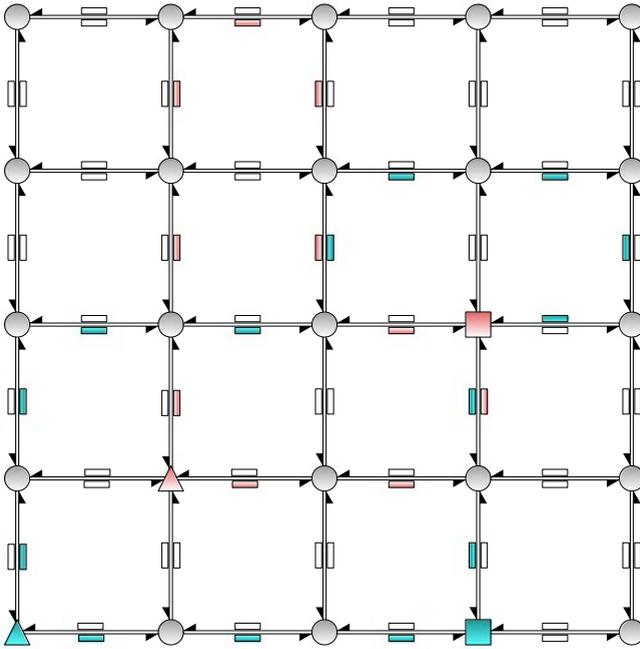


Fig. 5. Flow graph

superimposing all of the physical graphs output by this scheduling method. The graph illustrates global and local connectivity.

5 Network and Transport Layers

Let there be K commodities, each commodity specified by a source node $\sigma_k \in V_F$, and a destination node $\delta_k \in V_F$. A commodity specifies a stream of unique data. A *throughput vector* $\mathbf{f} = \{f_1, \dots, f_K\}$ is *feasible* if it respects both network capacity constraints and conservation of flow constraints. The objective is to maximize the sum throughput summed over the K commodities. Each commodity may be split across multiple paths; we define a *flow* as the set $\{x_e^k, e \in E_F, k \in [K]\}$, so that x_e^k is the flow for commodity k on edge e . The max multicommodity flow problem is:

$$\begin{aligned}
 \max_{\mathbf{x}} \quad & F(\mathbf{f}) = \sum_{k=1}^K f_k \\
 \text{s.t.} \quad & \sum_{i \in V_M} x_{i, \delta_k}^k = f_k, \quad k \in [K] \\
 & \sum_{k \in [K]} x_{ij}^k \leq c_{ij}, \quad i \in V_F, j \in V_F \\
 & \sum_{j \in V_M} x_{ji}^k = \sum_{j \in V_M} x_{ij}^k, \quad k \in [K], i \in V_F
 \end{aligned} \tag{6}$$

The first line states the objective is to maximize the sum commodity throughput. The first constraint defines the throughput for each commodity as the sum of the flow over all edges which terminate in commodity k 's terminal node δ_k . The second constraint is the capacity constraint; the third is the conservation of flow constraint.

Figure 5 presents a visual representation of a two-commodity flow graph. To demonstrate the amount of flow on each edge, a rectangular “container” is drawn to the right of the edge. If the container is white, this means that no data is flowing on that edge. If the container is filled, data is flowing on the edge, and the color of the fill indicates which commodity the data belongs to. The triangle shape is the source node of the data, while the rectangular shape is the node which is the desired destination for the data.

6 Conclusion

The visualization of the mobility graph informs the observer of the evolution of global patterns in the network topology. The visualization of the physical layer displays the interaction between transmitting and receiving nodes. The visualization of the MAC layer gives a set of “capacity pipes” through which information can travel as represented by the flow graph. Its visualization provides an overview of the movement patterns of data. All of this visual information is critical for making network control decisions.

References

1. Dhoutaut, D., Vo, Q., Lassous, I.G.: Global visualization of experiments in ad hoc networks (2003)
2. Bettstetter, C.: On the connectivity of ad hoc networks. *COMPJ: The Computer Journal* 47 (2004)
3. van Ham, F., van Wijk, J.J.: Interactive visualization of small world graphs. In: *INFOVIS*, pp. 199–206. IEEE Computer Society Press, Los Alamitos (2004)
4. Becker, R., Eick, S., Wilks, A.: Visualizing network data. *IEEE Transactions on Visualization and Computer Graphics* 1, 16–28 (1995)
5. Herman, I., Melançon, G., Marshall, M.S.: Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics* 6, 24–43 (2000)
6. Fitzek, F.H.P., Seeling, P., Reisslein, M., Zorzi, M.: ViTAN - visualisation tool for ad hoc networks. *IEEE Network* 17, 5 (2003)
7. Matthes, M., Biehl, H., Lauer, M., Drobnik, O.: MASSIVE: An emulation environment for mobile ad-hoc networks. In: *WONS*, pp. 54–59. IEEE Computer Society Press, Los Alamitos (2005)
8. Estrin, D., Handley, M., Heidemann, J., McCanne, S., Xu, Y., Yu, H.: Network visualization with nam, the VINT network animator. *Computer* 33, 63–68 (2000)
9. Ellson, Gansner, Koutsofios, North, Woodhull: Graphviz – open source graph drawing tools. In: *GDRAWING: Conference on Graph Drawing (GD)* (2001)

10. Chatila, R.: Mobile robot navigation: space modelling and decisional processes. In: Giralt, F. (ed.) *Robotics Research 3*, pp. 373–378. MIT Press, Cambridge (1986)
11. Canny, J.: *The Complexity of Robot Motion Planning*. MIT Press, Cambridge (1988)
12. Wu, Y., Jain, P., Kung, K.: Network planning in wireless ad hoc networks: a cross-layer approach. *IEEE Journal on Selected Areas in Communications* 23, 136–150 (2005)