INSTITUTE FOR DEFENSE ANALYSES

Modeling Terrain Impact on Mobile Ad Hoc Networks (MANET) Connectivity

Lance Joneckis
Corinne Kramer
David Sparrow
David Tate

May 2014

IDA Document NS D-5246
Log: H 14-000824

INSTITUTE FOR DEFENSE ANALYSES
4850 Mark Center Drive
Alexandria, Virginia 22311-1882

About This Publication

This work was conducted under IDA's independent research program (C2169). The views, opinions, and findings should not be construed as representing the official position of either the Department of Defense.

Copyright Notice

4850 Mark Center Drive, Alexandria, Virginia 22311-1882 • (703) 845-2000.

Modeling terrain impact on MANET connectivity

Lance Joneckis, Corinne Kramer, David Sparrow, and David Tate
Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, VA USA 22311-1882
ljonecki@ida.org

***Abstract*—Terrain affects connectivity in mobile ad hoc networks (MANET). Both average pairwise link closure and the rate at which the link-state changes when nodes move depend on the characteristics of the intervening terrain. However, these key parameters depend only mildly on internode distances in common real-world terrains, which suggests that network performance in real terrains can be usefully modeled using simple time- and location-invariant average link-state probability and rate of change. We use this fact to predict the number of disconnected subnets at a random instant in a given terrain, and we compare these predictions against Monte Carlo sampling of node positions.**

Keywords—mobile ad hoc networks; terrain factors; terrain mapping

I. INTRODUCTION

In previous work [1], we examined the impact that three very different terrains would have on mobile ad hoc network (MANET) performance. The terrain's most direct impact was allowing or blocking line-of-site between node pairs in the network. We used the existence of line-of-sight between node pairs as a surrogate for link-state. We found that some important metrics—in particular, the probability of link closure—are relatively insensitive to node separation, at least in the region of separation between 2 km and 8 km. This finding suggests that important features of network performance might be able to be described in terms of average values of these metrics: probability of link closure and link-state correlation length. This finding can be implemented via an Erdős-Rényi (E-R) graph treatment [2].

The three terrain sites chosen were Charlottesville, Virginia (CVA), White Sands New Mexico (WNM), and Ft. Huachuca Arizona (HAZ). WNM and HAZ were chosen because they are primary sites for the ongoing testing of military ad hoc networking radios. CVA was chosen because it represents a common rolling terrain that is quite different from that found in the desert southwest.

Fig. 1 shows topographic maps of the three areas, with north-south and east-west elevation profiles. CVA has a smaller range of elevations, but the spatial frequency of the topographic structure is greater, typical for the rolling hills. HAZ has a concave structure in the north-south direction, which will tend to increase line-of-sight probability. The chosen WNM site has a large irregular obstruction, which will often block connectivity. This obstruction is impassible by vehicles, and,consequently, nodes are prohibited from being placed at these locations.

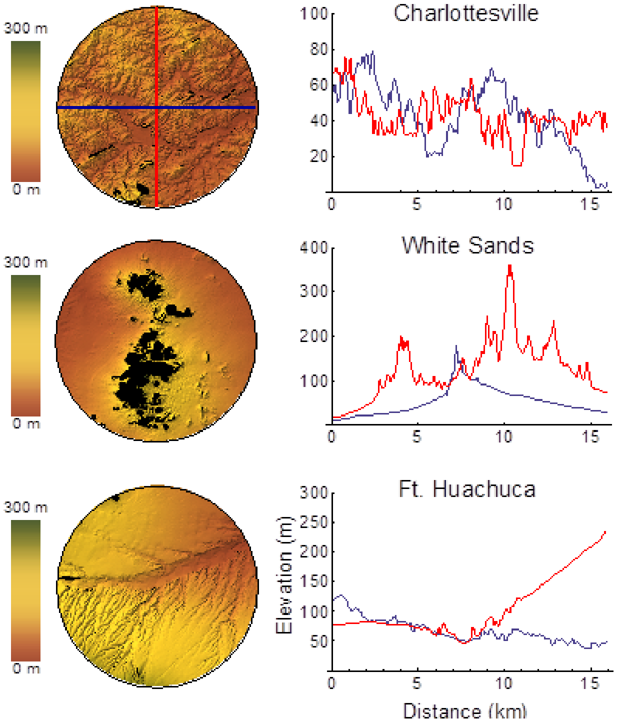

Fig. 1. Topological maps of the three areas.

We are presenting a description of network (link-state) performance metrics in terms of terrain-averaged features of node pair link-state for uniformly random laydowns of nodes using real terrain data and comparing these data to an E-R random graph model. We expect that the E-R model, which is based on average features, will capture aggregate characteristics of network performance (those characteristics not strongly influenced by individual nodes) better than it captures rare events or network extremes. In Section 2, we describe these calculations. Section 3 presents results, and Section 4 summarizes our conclusions based on these results.

II. CALCULATONS

A. Determination of Node Pair Link-States

We selected a circular area of approximately 16 km in diameter for each of the three regions, as indicated in Fig. 1.

We used digital terrain maps [3] with horizontal spacing of approximately 30 m and vertical resolution of 1 m. We assumed that a link is closed between two locations if the line-of-sight between antennas is 3 m above ground at those locations. We determined the presence or absence of line-of-sight between all location-pairs in the region [1]. The link-state probabilities were determined for node pairs as a function of their separation and were averaged over the entire space. This average probability that two nodes separated by a distance, d, have line-of-sight varies slowly with separation, and, in the following discussion, we replace it with a constant, p, to characterize the impact of terrain. The same probability is used in the E-R model. There are multiple approaches to the choice of p. For this paper, we choose p to match the empirical average probability of line-of-sight between two nodes put down at random in the circular area.

B. *Prediction of Network Properties*

The purpose of the calculations is to assess the extent to which network properties can be described using average features of the link-state probabilities. We are particularly interested in two network metrics: (1) the probability that the network has k connected components (subnets) and (2) the expected fraction of nodes that are included in the largest k subnets. In both cases, we assume a uniform random laydown of nodes in the terrain. We are especially interested in the behavior of these metrics as a function of the number of nodes in the network.

These features are estimated in two ways, which we then compare. First, we generate random laydowns in the actual terrain and describe the network connectivity for each. Second, we generate E-R random graphs with edge probability equal to the empirical average node pair link closure probability for the terrain and describe the connectivity for each. Averaging over many laydowns gives two different predictions of expected network topology. Comparing the properties of the terrain scenarios to the E-R scenario illustrates the extent to which actual terrain impacts network performance and can be captured by a simple random graph model using average features of the terrain.

C. *Precision of Simulation*

We derived a recursion relationship that allows us to compute exact values for the probability of forming a single network in terms of the link-state probability. This relationship was used to check the variance of the E-R averages. Discrepancies were typically 5% or less. The pairwise link probability (or the probability of forming a single network with two nodes) constrains the percentage of nodes in the largest network for N = 2 and N = 3 nodes. We used this fact to check the accuracy of the terrain sampling, and, again, we found discrepancies typically 5% or less. Thus, we conclude the numerical comparisons are sound.

III. RESULTS

In this section, we compare network performance metrics, as a function of the number of nodes, N, as determined using real terrain, and as predicted using an E-R model with matching probability of link closure (probability of an edge). The probabilities of link closure are given in Table 1 for the three sites.

TABLE I. PROBABIITY OF LINK CLOSURE FOR THE THREE SITES

Location	Link Closure Probability
CVA	0.052
WNM	0.19
HAZ	0.41

The probability for a fraction of the total nodes being within the largest k subnets are presented in Fig. 2 for the three sites under consideration and their E-R approximations.

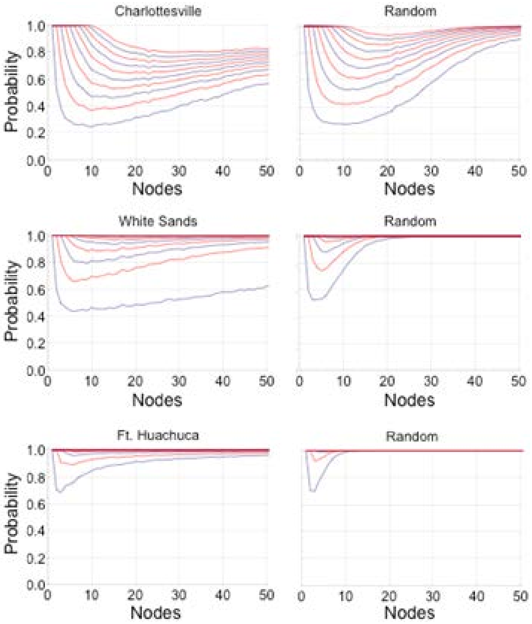

Fig. 2. The cumulative average fraction of nodes in the largest k subnets. Each curve in a plot is for the next subnet or component ordered by the number of nodes in the subnet. The lowest curve is for the largest single subnet; the curve above that is for the largest two subnets, and so forth.

The upper left-hand plot shows the (cumulative) expected probability that a fraction of the nodes will be included in the largest k subnets, as a function of the total number of nodes, N, in the CVA area. The lowest curve is for the largest single subnet, the curve above that is for the largest two subnets, and so forth. These curves must hit unity for k = 1 and for k = N.

For a 10-node laydown, we see that, on average, at least a quarter of the nodes will be in the largest single subnet. CVA has the lowest pairwise connection probability, which is not surprising when p is less than 1/N. This fraction is greater than 80% for favorable terrains, such as HAZ, with more than 10 nodes. As we consider the largest two, three, … subnets, the probability of a node being included increases by about 0.10 per subnet for the largest six subnets. For a 20-node laydown, the probability of nodes being in the largest subnet is higher, but this probability increases more slowly and is lower than the corresponding probability for the 10-node case when k = 4.

The chart on the upper right of Fig. 2 shows the results from the E-R graph model using the CVA value of p. The families of curves in the two cases are qualitatively similar. However, the E-R is universally more optimistic about network connectivity. The optimism is modest for 10 nodes but increases with network size until the curves converge at unity.

This pattern is largely replicated in the results for WNM and HAZ, where p is 0.19 and 0.42, respectively. The E-R model results are more optimistic to start with and converge to unity faster as the network size, N, increases. In addition, the performance under either approach is better in terrain that has higher values of p. The one obvious outlier to these trends is the probability of nodes being in the largest single subnet at WNM, which is notably low in the terrain-based calculation. The failure to converge to a single subnet as N grows is caused, in part, by the irregular obstruction, which readily separates the nodes into two or more subnets.

Fig. 3 presents the results for the probability of forming k subnets, again using both approaches for the three locations.

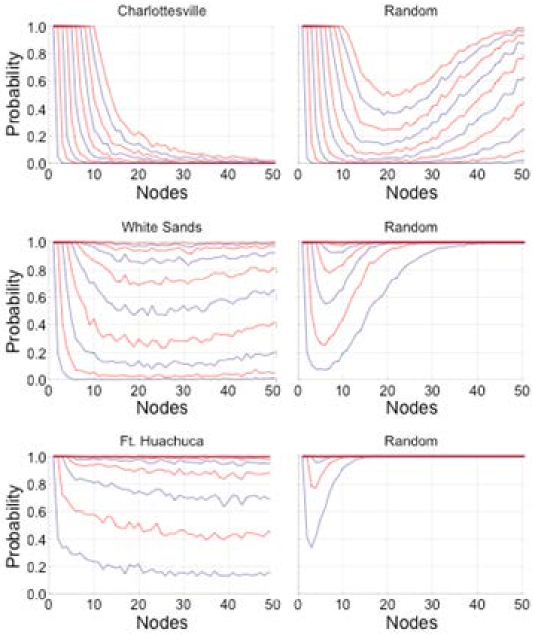

Fig. 3. The probability of forming k or fewer subnets. The curves in each plot are for sucessive values of k with the lowest curve having k = 1.

The upper left-hand plot of Fig. 3 shows the expected probability that the network will form k or fewer subnets as a function of the total number of nodes, N, in the CVA area. The curves must hit unity for N = k, and for N = infinity. However, at N = 50, all of the curves for k ≤ 10 subnets are still decreasing, with near certainty that there will be more than 10 subnets. This persistence of large numbers of subnets for large values of N is in dramatic contrast to the E-R results, where the largest few subnets contain almost all nodes for N > 10 at HAZ and for N > 50 in all three locations.

The results for the expected number of connected subnets based on real terrain and on the E-R model using the average value of p are qualitatively different for all three locations. There is increasingly rapid convergence to a single subnet for the E-R model but no observed convergence in real terrain for networks of up to 50 nodes. For WNM, the E-R results have converged before the terrain-based results have even hit their minimum values. For HAZ, which has a p that is high enough and where we should expect routine formation of fully connected networks, we see that the probability of a single network falls below 20% around N = 10 for the real terrain. In this case, the E-R model predicts nearly always forming a single subnet.

(We note, in passing, that the simulation results using real terrain have significant implications for the use of an elevated node to tie the network together. In most concepts, a "nearest neighbor" prohibition of some sort is used to prevent the elevated node from becoming an advantaged node that is then overwhelmed with traffic and becomes a bottleneck and a single point of failure for the network. A large number of subnets preclude this approach and impose an impractical burden on an air tier that tries to knit the subnets together. The probability of five or more subnets is 99% in CVA, 70 % in WNM, and 10% in HAZ. Only in HAZ would a single air node be able to fully connect the network without creating a bottleneck and/or single point of failure.)

The difference between these two sets of calculations is driven by having more isolated or nearly isolated nodes than predicted by the E-R model. These nodes will have a generally small effect on the percentage of nodes included in the largest subnets, whereas they can dramatically affect the total number of subnets formed. The systematic overprediction of connectivity by the E-R approach indicates the presence of nodes in particularly disadvantaged locations. We can quantify this observation using by examining the terrain data.

The nodes are either isolated, contained in the largest subnet, or neither. These fractions are shown for N = 10 and N = 20 for our three locations in Table II. Except for WNM, with the large obstruction, most of the nodes not in the largest subnet are isolated.

TABLE II. DISTRIBUTION OF NODES AMONG THE SUBNETS

Location and Network Size	**Percent of Isolated Nodes**	**Percent of Nodes in Largest Subnet**	**Percent of Nodes Neither in Largest Subnet, nor Isolated**
CVA 10 nodes	70.52	24.4	5.08
CVA 20 nodes	53.60	31.2	15.2
WNM 10 nodes	27.30	46.5	26.2
WNM 20 nodes	14.10	48.5	37.4
HAZ 10 nodes	14.00	84.7	1.3
HAZ 20 nodes	8.00	91.5	0.5

If we combine the results from Fig. 1 and Table II for HAZ, we see that most of the nodes are in the largest subnet and that the remainder is almost entirely isolated single nodes—despite the fact that the probability of forming a single network is fairly low. The low probability of the network being comprised of only one or two subnets is caused by a significant density of small, disadvantaged locations that lead to isolated nodes.

For CVA, most of the nodes in a laydown are isolated until the number of nodes, N, exceeds 20. There are relatively few nodes that are not in the largest subnet or isolated. However, the fact that the fraction of nodes that are neither in the largest subnet nor isolated increases with N suggests that for CVA, the disadvantaged locations are both more common and larger—large enough to occasionally capture more than one node.

For WNM, there is substantial probability of non-isolated nodes forming additional subnets rather than coalescing into the largest subnet. This probability of additional subnets is a consequence of the irregular obstruction, which effectively partitions the region into independent areas that will have their own large subnet until N becomes large enough to ensure connecting paths around the obstructions.

To summarize, the E-R model qualitatively describes the fraction of nodes captured in the k largest subnets, but it does not accurately estimate the probability of forming a single subnet. In fact, it is qualitatively incorrect with respect to the behavior of formation for any number of subnets. In the terrain-based simulations of the well-known random graph result, there is no evidence of a phase change to a single connected network.

We explore this idea further by plotting the probability of forming a single network in a E-R graph, as a function of the edge probability, p, (or the link-state closure probability) (see Fig. 4). This result we compare with the data from our three sites, using their particular values of p. Note that for N = 2, both approaches yield the fraction p by construction.

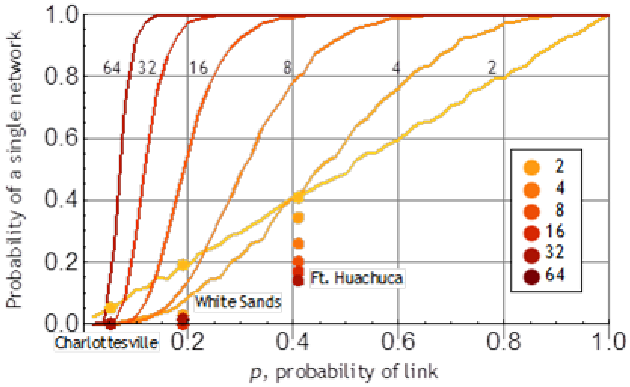

Fig. 4. The probability of forming a single (fully connected) network for various network sizes, as a function of link-state closure probability. The curves are for the E-R model, while the points show simulations using actual terrain. Ideally, the simulation points should fall near the corresponding E-R curve (i.e., the curve of the same color as the point). It is evident that this relationship is not present and that the trend is actually in the opposite direction. The E-R curves clearly show the presence of the well-known phase transistion to a fully connected network for E-R random graphs. This phase transistion is not present in any of the three cases using real terrain.

We see that as the number of nodes increases, the E-R model's probability of forming a single subnet that includes all of the nodes converges to certainty at lower and lower values of p. However, for the probability based on real terrain, there is no such limiting behavior. In fact, in most cases, the probability decreases as nodes are added. Forming a single network from a random laydown of nodes in real terrain is much less likely than the E-R model would predict.

IV. CONCLUSIONS

We have computed network properties using the probability (p) of antenna line-of-sight as a surrogate for the probability of link closure in three different real-world locations. These results are compared against an Erdős-Rényi (E-R) random graph model, using the same p.

We showed that for an aggregate measure, such as the number of nodes that will be found in the largest k subnets, there is qualitative agreement between the E-R model and simulation results when the actual terrain is used. However, the E-R approach systematically overestimates formation of large subnets. Thus, some aggregate measures of network performance can be usefully approximated by E-R models in terms of readily determined average properties of the terrain. This approximation permits both analysis and planning using generic terrain characterized by a few parameters rather than by detailed map data tied to a specific location. However, more work in characterizing which aggregate measures are robust under E-R modeling is needed before this approximation can be reliably applied.

The terrain-based and E-R methods make qualitatively different predictions for the distribution of the number of connected subnets in the network. These differences are driven primarily by a larger than predicted number of isolated nodes, except for the WNM location. At WNM, the large irregular obstruction leads to multiple subnets with more than one node. For all three locations, the large number of isolated nodes has little impact on the nature and connectivity of the largest subnets but completely changes the character of the distribution of the overall number of subnets.

Proposed concepts of operations (CONOPS) for military MANET tend to assume that units will maneuver without explicit consideration of the impact that their locations and movements will have on network connectivity. The fragility of network connectivity and the frequency of node isolation under random configurations show that this assumption is not warranted in typical terrain. Except in the most favorable terrain, "terrain-aware" approaches to network management will probably be required.

There are also implications of this work for the use of an air tier to handle terrain. The unexpectedly large number of subnets for modest values of N imposes a burden on the air tier to knit the subnets together. Only under the most favorable terrain circumstances would a single air node not be forced to handle the bulk of network traffic. This concept will be explored in future work.

It is not surprising that a description in terms of average terrain features better predicts more aggregate network metrics. In this case, the connectivity shown in real terrain, even using aggregate measures, is systematically worse than the E-R model predictions. For the probabilities of forming a particular number of subnets as the network size, N, increases, the behavior predicted by the E-R approach is qualitatively more favorable and nonpredictive of real-world behavior. There is no evidence in the terrain-based calculations that the network will coalesce into a single subnet, even for operationally large numbers of nodes and high probabilities of pairwise connection. The assumption that one can route around terrain obstacles in a practical MANET is questionable in at least some common terrain types.

ACKNOWLEDGMENT

The authors would like to thank John Fregeau for enlightening discussions.

REFERENCES

[1] L. Joneckis, J Fregeau, C. Kramer, D. Sparrow, and D. Tate, “Impact of Terrain Features for Tactical Network Connectivity,” IDA Document NS D-5026, submitted to Ad Hoc Networks.

[2] P. Erdős and A. Rényi, “On random graphs,” Publ. Math., vol. 6, pp. 290–297, 1959.

[3] USGS National Elevation Data Set, http://ned.usgs.gov.

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE May 2014	2. REPORT TYPE Final	3. DATES COVERED *(From–To)* January 2014 – May 2014

4. TITLE AND SUBTITLE

Modeling Terrain Impact on Mobile Ad Hoc Networks (MANET) Connectivity

5a. CONTRACT NUMBER
HQ0034-14-D-0001

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Lance Joneckis
Corinne Kramer
David Sparrow
David Tate

5d. PROJECT NUMBER

5e. TASK NUMBER
AE-2-2169 (C2169)

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, VA 22311-1882

8. PERFORMING ORGANIZATION REPORT NUMBER

IDA Document NS D-5246
Log: H14-000824

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, VA 22311-1882

10. SPONSOR/MONITOR’S ACRONYM(S)

IDA

11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited (2 May 2014).

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Terrain affects connectivity in mobile ad hoc networks (MANET). Both average pairwise link closure and the rate at which the link-state changes when nodes move depend on the characteristics of the intervening terrain. However, these key parameters depend only mildly on internode distances in common real-world terrains, which suggests that network performance in real terrains can be usefully modeled using simple time and location-invariant average link-state probability and rate of change. We use this fact to predict the number of disconnected subnets at a random instant in a given terrain, and we compare these predictions against Monte Carlo sampling of node positions.

15. SUBJECT TERMS

mobile ad hoc networks, terrain factors, terrain mapping

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Mr. Philip Major
a. REPORT Uncl.	b. ABSTRACT Uncl.	c. THIS PAGE Uncl.	SAR	8	19b. TELEPHONE NUMBER *(include area code)* (703) 845-2201

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18