Title
Multi-Dimensional Routing

Author
Garcia-Luna-Aceves, J.J.

Publication Date
2008-08-03

Peer reviewed

Multidimensional Routing

Stephen Dabideen*
*Department of Computer Engineering
University of California, Santa Cruz
Santa Cruz, CA 95064
Email: dabideen@soe.ucsc.edu

J.J. Garcia-Luna-Aceves†*
† Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
Email: jj@soe.ucsc.edu

***Abstract*—We present a new perspective on the design and analysis of routing protocols for mobile ad-hoc networks (MANETs). Routing metrics, such as distances or link states, result in an ordering of nodes in the network with respect to the origin of the metric. The manner in which the nodes of a network are ordered can give some insight into the performance of the routing protocol. We show how the use of multiple metrics, an approach we call multidimensional routing, renders orderings among nodes that result in routing protocols that are efficient, more robust, and resilient to link failures. We explain why some routing protocols are inherently more effective than others, which can serve as guidelines for future development of routing protocols for MANETs.**

I. Introduction

Sending information from source to destination can be attained by flooding of data packets throughout the entire network. However, flooding is a very inefficient use of network resources. Routing protocols avoid the need to flood data packets by establishing an ordering among them with respect to their intended destinations. Flooding is then limited to the route discovery phase of on-demand routing protocols, or to the dissemination of distances to destinations or the state of links in proactive routing protocols. The ordering established in a routing protocol based solely on the distance to a destination can be viewed as one-dimensional, in that a single metric is used to order the nodes. Flooding, on the other hand, can be viewed as zero-dimensional, given that no routing metric is used. We define the *dimensionality* of a routing protocol as the number of metrics used to distinguish and order nodes in a network.

Clearly, there is significant improvement in efficiency moving from zero-dimensional to one-dimensional routing. Today, many routing protocols are two-dimensional, in that distance and freshness of information (sequence number) are used to order nodes with respect to destinations [1], [2], [3]. This paper explores the beneficial effects of adopting routing that uses more than two ordering dimensions.

Section III shows how an increase in the dimensionality of a routing protocol can lead to more efficient routing by maximizing the number of available routes between source and destination. We also define a quantitive means to evaluate and compare the potential efficiency of the ordering attained. We explain why some orderings are better than others and show that for every topology there is a maximal ordering and further increases in the dimensionality beyond this point yields no benefit.

As the number of ordering dimensions increases, the traditional view of destination-based shortest path routes becomes obsolete and more elaborate n-ary relations are needed to order the nodes based on their values in each dimension. Section IV briefly explores routing relations for multidimensional ordering. Section V considers the different types of metrics which can be used as dimensions.

Section VI presents the Constrained Scalable Hybrid (CaSH) routing protocol as a concrete example of the effectiveness of multidimensional routing. Section VII then compares the performance of CaSH to that of contemporary routing protocols. The results obtained from simulations clearly indicate that CaSH outperforms these protocols.

II. Related Work

The use of multiple constraints in routing is mostly associated with Quality of Service [4] routing protocols which attempt to optimize paths over several variables such as the Differentiated Service Field [5]. In this paper, we aim to solve a different problem through the use of multiple orderings. Instead of finding paths with some minimum standards, we simply aim to maximize the number of paths using a small number of variables. This is a novel approach in the field of MANETs and we show that it yields significant benefits. We use the CaSH routing protocol to show the effectiveness of multi-dimensional routing. CaSH establishes paths on demand and maintains them proactively and thus is a hybrid routing protocol. There have been several different approaches to hybrid routing in MANETs, the most well-known being the Zone Routing Protocol [6], but CaSH is the first to fully utilize multidimensional orderings.

III. Multi-dimensional Orderings

A truly resilient routing protocol is one that recovers from failures with minimal cost, in terms of both delay and network bandwidth. The resiliency of a routing protocol is a function of the number of possible paths between sources and destinations. The more paths allowed, the more resilient the protocol is.

This work was funded in part by the ARO under grant W911NF-05-1-0246, by the NSF under grant CNS-0435522, by DARPA through Air Force Research Laboratory (AFRL) Contract FA8750-07-C-0169, and by the Baskin Chair of Computer Engineering.

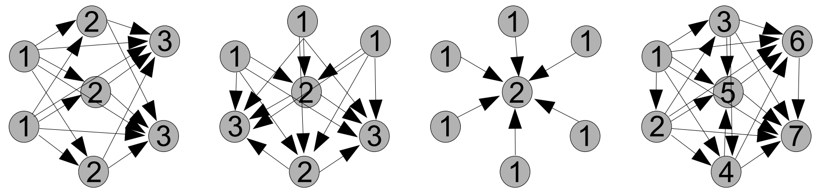

(a) Ordering A (b) Ordering B (c) Ordering C (d) Ordering D

Fig. 1. Successor graphs of three different orderings

Hence, the ordering used in a routing protocol should provide as many paths as possible between source and destination.

To evaluate the resilience of an ordering we define the *degree* of an ordering, denoted δ, to be the average size (geometric mean) of the partitions resulting from the metric used in that ordering. The degree of an ordering is indicative of the extent to which the nodes can be ordered in that dimension and thus the quality of the ordering, with a smaller degree representing a better ordering.

Definition If an ordering partitions the N nodes of a network into n disjoint sets, with the cardinality of the i^{th} partition being x_i, the *degree* of that ordering is given by $\delta = \left(\prod_{i=1}^{n} x_i\right)^{1/n}$

As an example, consider the seven-node network shown in Fig. 1 with four different orderings. The degree of ordering A in Fig. 1(a) is $\sqrt[3]{2*3*2} = 2.29$. The same is true for ordering B, and Ordering C has a degree of $\sqrt[2]{6*1} = 2.45$. Ordering D, however, has a degree of 1, which is the minimum value and therefore the best possible ordering.

Fig. 1 also shows a successor relation and thus the relationship between the degree of an ordering and the size of successor relations: for a given topology, the smaller the degree the greater the average number of successors per node. The reason for this comes from the fact that, in order to maintain loop freedom, a node cannot make an adjacent node its successor if they have the same value in the ordering. Consequently, every link between nodes with the same value can be considered a wasted link, because it cannot be used in that particular ordering. Using ordering D, the maximum number of successor relations is achieved. In a directed acyclic graph (DAG) in which each node has a different value, every link can be used in a successor/predecessor relationship. While the degree of an order does not give an exact number of paths created from the source to the destination, it is sufficiently useful for comparisons between different orderings.

A. Enhancing the Degree of an Ordering

An ordering can be made better by decreasing the average size of the partitions it creates. To minimize the degree of ordering in a single dimension, all the nodes need to have a unique value. This cannot be achieved merely by using the distance in terms of the number of hops from the destination as the value of the node. If the average connectivity in the nodes in the network is d, the number of nodes with the same number n of hops from the destination grows with $\Theta(dn)$. Hence, a more creative solution must be employed. In the Sequence-number Window Routing (SWR) protocol [7], the value of the nodes is a function of distance to the destination as well as several other parameters, which leads to an ordering with a better degree. However, in a completely distributed scenario, there is no way to guarantee that the values assigned to nodes in SWR will be unique and thus the maximal ordering can be achieved.

Another way to enhance the degree of an ordering is through multi-dimensionality. By compounding multiple orderings into a single metric, the average size of the partitions created by the multidimensional ordering can be made smaller and this can result in an ordering with a smaller degree.

B. The Limits of Multidimensionality

The improvement attained by the addition of a second dimension is a result of nodes that are identical in the first dimension being differentiated in the second dimension. Likewise, the use of a third dimension will allow for nodes identical in the first two dimensions to be distinguished in the third. By increasing the number of dimensions, the number of nodes with an identical routing value will be decreasing.

Although increasing the number of dimensions can increase the number of successors of a node, there is a limit. Once all the nodes have a unique *value* there can be no further gains in the number of paths from the source to the destination.

The correlation between orderings and the degree of each order will affect the magnitude of the gain achieved by adding another dimension. The more similar orderings are, the higher we say their correlation is. The more correlated orderings are, the smaller the gains achieved by compounding them into a multidimensional ordering. An extreme example would be using the same ordering as both the first and second order of a two-dimensional ordering. Given that both dimensions are the same, the degree of the two-dimensional order would be the same as that of the one-dimensional order. Ordering A and ordering B of Fig. 1 are not very correlated as can be seen by the directions of the successor relations and there is a significant gain achieved by the two-dimensional ordering achieved my combining these two orderings.

In addition, the closer the degree of the ordering to the lower limit of 1, the less the gain is likely to be. An ordering with a low degree has very few nodes not related by a successor/predecessor relationship and unless the new dimensions can specifically create such relations between these nodes, there will be no gain. As long as the orderings are not fully correlated, the maximum gain decreases with the number of dimensions. For the third dimension, the number of new successors will grow as $O(\delta_{1,2} * N)$, which is the average number of nodes identical in both the first and second dimensions, and $\delta_{1,2} < \delta_1$.

IV. Multi-dimensional Relations

Multi-dimensional routing requires a total order relation on sets of two or more elements depending on the dimensionality of the protocol. For one-dimensional routing protocols the "greater than or equal to" relation is sufficient to achieve the desired total ordering of the nodes. As the number of dimensions increases more complex n-ary relations must be used to fully exploit the benefits of multidimensionality and the properties of transitivity and antisymmetry of such relations would guarantee loop freedom.

In general, a weighted sum of the value in each dimension can be used in the determination of successor-predecessor relationships. If w_i is the weight assigned to the ith dimension and v_i^x is the value of node x in the ith dimension, an n-dimensional value, $V(x)$, can be assigned to each node x, which is given by $V(x) = \sum_{i=0}^{n} w_i * v_i^x$. The choice of the weights assigned to each dimension greatly influences routing decisions in the protocol. A simple loop-free condition would be that a node x can choose a node y as its successor if $V(x) < V(y)$.

Lexicographic ordering is one case in which $w_i >> w_{i+1}$, and can be used as an n-ary relation to create an n-dimensional total ordering of the nodes in the network. A multi-dimensional DAG is formed by adding edges, one dimension at a time, and with each successive dimension having a lower priority than the one which came before it. Once the direction of an edge is set by a dimension, it cannot be altered by any other dimension.

V. Possible Dimensions

Now that the value of multi-dimensional routing has been established, we consider the possible metrics that can be used for the dimensions of the routing protocol. A broad distinction can be made between arbitrary, meaningless metrics and those whose optimization can positively affect the performance of the routing protocol.

A. Optimization Metrics

The most intuitive choices for dimensions are metrics to be optimized. One used in many routing protocols is the shortest distance to the destination. The shortest-path route is often the most efficient in terms of network resources and using shortest paths to destinations as the first dimension results in a routing protocol that prioritizes an efficient use of resources. Shortest path distance can be used as any dimension depending on which metric the protocol should optimize. Distance from the source can also be used as a metric, with nodes further from the source having priority over those closer the source. While this metric does not guarantee progress to the destination, it can, however, work effectively as a secondary metric. If shortest path distance to the destination is used as the first dimension and shortest path distance from the source is used as the second dimension, then if a node no longer has a neighbor closer to the destination than itself, it may choose as a successor a node further from the source that can lead to a path to the destination.

Various Quality of Service (QoS) metrics can also be employed as secondary dimensions in a routing protocol. It is important that the metrics be weighted in order of importance. The available bandwidth along a path and the end-to-end delay are typical QoS metrics that can be used.

MANET-specific variables such as mobility can also be used as a dimension. The more mobile a node is, the least preferable it becomes when creating a path, because the use of a highly mobile node will lead to link failure as the node moves out of range of its neighbors. If each node can determine its degree of mobility, absolute or relative, then this can be used as a dimension to favor paths consisting of nodes with lower mobility.

B. Arbitrary Dimensions

An arbitrary dimension is a possibly unique, meaningless value assigned to each node in the network. Such dimensions can be a node ID or some other pre-set value and would most suitably function as a non-primary dimension. A simple but effective protocol could be one which uses shortest path distance as the first dimension and node ID as the second. Shortest path route would have precedence, but in case of link or node failures, a node can choose as its new successor another node with the same shortest path distance as itself but with a lower node ID, for example. This would lead to loop-free routing with more allowed paths than a simple one-dimensional shortest path routing protocol.

C. Non-shortest Path Routing

The concepts of routing and shortest path routing have become synonymous, with most routing protocols delivering shortest path routes. In a wireless environment, shortest path routes may not be the best, but finding efficient non-shortest path routes has also been a challenge.

Using a multidimensional approach to routing, it is possible to easily develop efficient non-shortest path routing protocols. If several destination-based orderings are compounded in a multidimensional fashion, then the importance of the distance-based ordering can be reduced in a way that the accuracy of the distances can be relaxed. If we were to use an approximate distance, for example, if $\lfloor \frac{D}{5} \rfloor$ is used as the ordering in the first dimension where D is the actual distance, then many nodes relatively close to each other can be viewed as equidistant from the destination, even though they are not. In this case, routing decisions will depend on the orderings used in the remaining dimensions. For protocols to work using non-shortest path metrics, all orderings must lead to the destination and therefore cannot all be arbitrary.

VI. The Constrained Scalable Hybrid (CaSH) Routing Protocol

The CaSH routing protocol was developed to take advantage of multidimensional routing as it uses a four-dimensional ordering to maximize the number of successor-predecessor relationships between nodes and thus minimizing the number of unusable links.

A. Signaling

The signaling in CaSH is hybrid in that paths are established on-demand and are maintained proactively to achieve routing which incurs both low overhead and low latency.

Route Request packets (RREQs) are used in a similar manner to that of AODV in setting up routes to destination nodes. Routes are established on demand by the flooding of RREQs which contain the number of hops from the source and a sequence number which is incremented every time the source initiates a RREQ.

Once the destination receives the RREQ, it initiates a route reply message (RREP) but, unlike AODV, this RREP is flooded among a subset of nodes between the source and destination. The RREP will contain a sequence number set by the destination so that it can be uniquely identified and the distance from the destination. Upon receiving a RREP, a node will increase the hop count and retransmit the RREP if it is in the Active Routing Area as defined in Equation 1.

Definition Let D_A(X) denote the distance of node X from node A, let S_{AB} be the shortest path distance from A to B and let M_{AB} denote the maximum allowed path from nodes A to B. Then the *Active Routing Area (ARA)* of source A and destination B is defined to be the set of nodes which satisfies the constraints of inequality (1).

$$S_{AB} \leq D_A(X) + D_B(X) \leq M_{AB} \tag{1}$$

If a node no longer has a route to the destination it will initiate a local Route Error packet (RERR). Its neighbors, upon receiving this RERR, will update their routing table. If they have a route to the destination then nothing further is done, if they do not have a route to the destination then they issue a RERR. If RERRs propagates to the source, then there is no known path from the source to the destination and the source increments its sequence number and sends a new RREQ to the destination.

After receiving the initial RREQ, the destination node will proactively flood the ARA with periodic RREPs and the source will do the same. This results in nodes maintaining up-to-date distance information to both source and destination as long as the path is needed. Once there is no longer need for a the link between the source-destination pair the proactive updates will cease.

A feedback mechanism is used to control the rate at which proactive updates are flooded in the ARA. Destination nodes use the reception of RREQs, or lack thereof, as feedback to adjust the update frequency. If a node is already a destination and it receives a RREQ, then paths are being invalidated faster than they are established so the node should increase the update frequency. If an active destination does not receive a RREQ between periodic updates, then it might be possible to increase the update period, and consequently reduce routing overhead without increasing the latency of the route. Eventually, the update period should reach minimal fluctuations around the optimal period. The update interval at any instant, t, is the sum of a base and an offset, given by equation (2).

$$f(t) = \beta + \delta(t) \tag{2}$$

where β, the base, is a fixed value which is the minimum update interval and $\delta(t)$, the offset, is determined dynamically and defined recursively as:

$$\delta(t) = \begin{cases} \delta(t) = 0 & \text{for t = 0} \\ \delta(t-1) + \epsilon & \text{for t} > 0\text{, if no RREQ} \\ \alpha * \delta(t-1) & \text{for t} > 0\text{, if RREQ.} \end{cases} \tag{3}$$

If the node receives a RREQ, the offset is multiplicatively decreased. If, on the other hand, no RREQs are received between updates, then the node may safely increase the update interval by a constant ϵ. Whenever a RREQ is sent, the source is holding a packet and there is a penalty in terms of latency associated with this. In most situations it would be best to have $\alpha < 0.5$ since this would reduce the update interval rapidly to prevent this penalty from being incurred in successive intervals. Also, the offset is slowly increased in order to prevent too much overshooting of the optimal update interval since this too will incur latency penalties. In the experiments described in Section VII, α was set to 0.2, ϵ was set to 2 seconds and β was set to 1 second, as this configuration delivered the most promising results.

Nodes within the ARA are labelled with the quadruplet (SN_D, D_D, SN_S, D_S) called a Route Label (RL). When routing packets, nodes select a successor with a RL lexicographically smaller than itself, except larger values of D_S are favored. Intuitively, if a node cannot send a packet closer to the destination, then it should send it to a node at the same distance to the destination but further away from the source. As links get broken packets take alternate routes. If there is no remaining routes from the source to the destination which are allowed by the existing RLs, the source initiates a RREQ. When there are no more packets to be sent, the destination stops sending proactive updates and the ARA is dissolved.

VII. Performance

In order to gauge the performance of the CaSH routing protocol it was implemented in the Qualnet 3.9.5 simulator and its performance was compared to that of several well known routing protocols, representing each of the main approaches to routing in MANETs. DSR is representative of source routing schemes, OLSR of proactive routing and AODV of on-demand routing.

A. Simulation Environment

Two scenarios were used and the parameters are summarized in Tables II. The first of these was designed to rigorously

TABLE I
SIMULATION RESULTS

	Scenario A			Secnario B		
	Delivery Ratio	Latency	Net Load	Delivery Ratio	Latency	Net Load
AODV	0.600±0.102	0.086±0.037	14.4±5.3	0.901 ± 0.031	0.072 ± 0.015	5.04 ± 1.31
DSR	0.142±0.099	18.47±15.91	5.02±1.24	0.1435 ± 0.04	42.69 ± 12.89	2.65 ± 0.34
OLSR	0.300±0.081	0.072±0.015	67.46±1.2	0.714 ± 0.044	0.104 ± 0.021	17.21 ± 0.16
CaSH	0.782±0.103	0.147±0.104	7.9 ±2.7	0.975 ± 0.031	0.067 ± 0.047	1.92 ± 0.20

test the performance of the protocols in a dynamic environment with volatile links. This choice of parameters satisfies the minimum standards for rigorous MANET protocol evaluation as prescribed in [8] as it results in an *average shortest path hop count* [8] of 4.03 and *average network partitioning* [8] of 3.9%. This would ensure that packets should travel several hops from source to the destination and thus would test the robustness of the protocols.

The second scenario uses a greater radio range, 200m, to add more stability to the links and create more multi-path opportunities, for which the CaSH routing protocol was designed to exploit. Consequently the average network partitioning as well as the average shortest path hop count would be reduced.

Each experiment lasted for 900s and for each protocol the experiment was repeated 250 times with random node placement and mobility. In each experiment, there were 10 CBR sources, which started generating packets at a random time to a randomly chosen destination. Each CBR source generated 800 packets at a rate of 4 packets per second.

Parameter	Value
Simuation time	900s
Number of Nodes	100
Simulation Area	1000m x 1000m
Node Placement	Uniform
Mobility Model	Random Waypoint
Min-Max Speed	1-10m/s
Pause time	30s
Propagation model	Two-ray
Physical layer	802.11
Antenna model	Omnidirectional
MAC Protocol	802.11 DCF
Data Source	CBR
Number of packets per flow	800
Packet rate	4 packets per second
Node density	0.001 nodes/m^2

TABLE II
EXPLICIT SIMULATION PARAMETERS

Three metrics were used to evaluate and compare the performance of the protocols. Delivery ratio is the fraction of packets that arrive at the corresponding destination by the end of the simulation. Latency is the average end-to-end delay experienced by the data packets. Net load is the number of overhead packets (RREQs, RREPs, RERRs, Hellos, etc.) which were initiated or forwarded, divided by the number of data packets sent. This takes into account packets that were sent into the network and were dropped or did not make it to the destination for any reason. This last metric gives an indication of the number of overhead packets needed to send a packet from the source to the destination. The simulation results for the four routed protocols tested are summarized in Table I, where the mean and a 95 percent confidence interval are given.

B. Delivery Ratio

In both scenarios, the CaSH routing protocol delivers a higher number of packets than the rest. The poor delivery ratio of DSR and OLSR in large networks shows that these protocols are not scalable. AODV sets up only a single route while CaSH sets up multiple routes thus allowing greater resilience to link failure and this is reflected in scenario A. There is a considerable improvement in performance between scenarios A and B, because links are more stable and fewer packets are dropped due to lack of a route to the destination. Also, CaSH sets up even more paths to the destination and achieves a near perfect delivery ratio.

C. Latency

There is a significant difference in the performance of AODV and CaSH versus that of DSR and OLSR. For CaSH, there are two interesting contributors to the value of the latency experienced. When an alternate path exists to a destination, the time taken to perform local route repair is negligible. On the other hand, if such a path does not exist, there will be some backtracking during which time nodes may discover broken links as they try different neighbors. When there is no longer a known valid path from the source to the destination, it takes a significantly longer time for the source to discover this than in AODV with local repair. Consequently, in a volatile environment as in Scenario A, CaSH experiences a slightly higher latency, while in a more stable environment, Scenario B, CaSH enjoys a performance advantage over AODV.

D. Network Overhead

The most interesting advantage of CaSH over the other routing protocols is in terms of network overhead. This is the number of non-data transmissions (including forwarding RREQ, REEPs and hello messages) divided by the number of data packets transmitted by the source. Flooding, as used in AODV adds significantly to the overhead as a single RREQ is retransmitted by almost every node in the network. Even though CaSH employs some proactive signaling, it incurs considerably lower overhead than the fully reactive AODV and DSR routing protocols. This might be influenced by the lifetime of valid links. If the expected lifetime of a link is small, then reactive routing protocols would have to frequently flood the network as they search for a path and CaSH

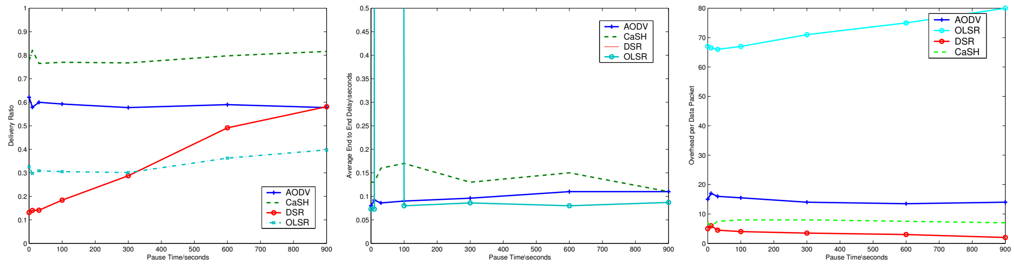

(a) Delivery Ratio (b) End to End Delay (c) Network Load

Fig. 2. Performance variation with Pause Time

demonstrates that in such scenarios it would be a better idea to introduce some proactive path maintenance. Also, the fact that the proactive updates in CaSH are restricted to a subset of nodes contributes to its relatively low overhead. Furthermore, the feedback mechanism employed in CaSH would allow the nodes to find an optimal period for proactive updates. It is not surprising that OLSR experienced the greatest overhead as it is a fully proactive routing protocol.

Graphs showing how these statistics vary with the pause time of the node are shown in Figure 2 for Scenario A. It can be seen that the CaSH routing protocol is not very sensitive to mobility and maintains a high performance,with respect to the tested metrics, even in very mobile networks. The delay of DSR is much greater than the others and does not appear in the range of the graph in Figure 3.

E. A Dimensional Analysis

AODV, DSR and OLSR use a sequence number to gauge the freshness of the routing information. If routing messages are properly disseminated across the network then all the nodes should have the same sequence numbers. However, due to the uncertainty of the wireless media some nodes might have an older sequence number. In general, the degree of the ordering achieved by the sequence number across the nodes would be relatively large compared to the ordering achieved by the distance to the destination metric. In fact the multidimensional gains of such sequence numbering as an ordering metric can be negligible.

CaSH on the other hand is four-dimensional, albeit two of its dimensions are sequence numbers with large degrees. While there is some correlation between the remaining two dimensions: distance to the destination and distance to the source, there is certainly an improvement in the ordering among the nodes in the network. If the nodes were to be grouped into equivalence classes of these two metrics, it would appear as concentric circles centered at the source and the destination. Consequently there will be nodes identical in one dimension and different in the other, which leads to the improvement in the ordering especially as the distance between the source and destination increases.

The performance of CaSH is heavily dependent on performing local re-routing of packets after a link failure with minimal overhead. This in turn is dependent on the degree of the multidimensional ordering. Although there are other factors which contribute to the success of the CaSH routing protocol, the increase in the number of paths established with each proactive update decreases the need for proactive updates since more link failures can be tolerated before an update is absolutely necessary and CaSH exploits this by adjusting the update period so as to minimize bandwidth consumption without incurring significant delay.

VIII. Conclusion

We introduced the concept of multidimensional routing and have shown how it can lead to a better ordering of nodes in a network. The ordering of nodes in a network can be improved by compounding several ordering into one. We presented tools to analyze the efficiency of orderings and discussed the use of n-ary relations in such routing environments. As a concrete example of this framework, we presented the CaSH routing protocol. Simulation experiments were used to illustrate that CaSH, outperforms several commonly used routing protocols. The CaSH routing protocol uses multidimensional ordering, and employs several unique mechanisms to improve its performance. It uses interest-based dissemination of information which dictates the nature of the routing information distributed and the extent to which it is distributed in the network.

References

[1] C. E. Perkins and E. M. Royer, "Ad hoc On-Demand Distance Vector Routing," *IETF Internet draft, draft-perkins-manet-aodvbis-00.txt*, October 2003.

[2] D. Johnson, D. Maltz, and Y. Hu, "Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR)," *IETF Internet draft, draft-ietf-manet-dsr-09.txt*, April 2003.

[3] C. Perkins and P. Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers," *Computer Communication Review*, 1994.

[4] G. Huston, "Next Steps for the IP QoS Architecture," *IETF RFC 2990*, 2000.

[5] K. Nichols, S. Blake, F. Baker, and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," *IETF RFC 2474*, 1998.

[6] Z. Haas and M. Pearlman, "The Zone Routing Protocol (ZRP) for Ad Hoc Networks," *IETF Internet draft, draft-ietf-manet-zone-zrp-02.txt*, 2002.

[7] J. J. Garcia-Luna-Aceves and H. Rangarajan, "A New Framework for Loop-Free On-Demand Routing Using Destination Sequence Numbers," *Proc. IEEE MASS 2004: The 1st IEEE International Conference on Mobile Ad-hoc and Sensor Systems*, 2004.

[8] S. Kurkowski, T. Camp, and W. Navidi, "Minimal Standards for Rigorous MANET Routing Protocol Evaluation," *Technical Report MCS 06-02, Colorado School of Mines*, 2006.