Probabilistic Addressing in Wireless Networks

Tobias Vaegs, Muhammad Hamad Alizai, Klaus Wehrle
Distributed Systems Group, RWTH Aachen University
{tobias.vaegs,hamad.alizai,klaus.wehrle}@rwth-aachen.de

***Abstract*— The lack of permanent network infrastructure and often unplanned deployments in many multihop wireless communication scenarios restrict nodes to determine their own addresses based on the underlying connectivity in the network. However, due to unreliable connectivity and rapidly changing link qualities in wireless networks, establishing uniform addressing and stable point-to-point routing is challenging.**

In this paper, we present Statistical Vector Routing (SVR), a virtual coordinates based addressing and routing mechanism that efficiently deals with dynamic communication links in wireless networks. It assigns stable probabilistic addresses to nodes without the need to pessimistically estimate links over longer periods of time. The routing metric predicts the current location of a node in its address distribution. Our prototype implementation over real testbeds indicates that SVR, when compared to current approaches, achieves 3-7 times more stable addressing, reduces the magnitude of change in addresses by 2-10 times, and minimizes the hop distance and transmissions in the network by 10-15%.

***Index Terms*— Wireless Communications, Routing**

I. INTRODUCTION

DEALING with unreliable and highly dynamic wireless links is a major challenge in establishing stable point-to-point routing in multihop wireless networks. The problem is further aggravated when location information is not available, a common situation in many wireless network deployments, and the nodes have to determine their own addresses reflecting their underlying connectivity. As a result, rapidly changing link conditions do not only affect packet delivery and routing topology, but also the locations and thereby the addresses of the nodes. A plethora of solutions [1]–[3] has been presented for situations, both in ad hoc and sensornets, where location information is not available at the nodes and geographic methods cannot be used for routing. The majority of the location independent addressing and routing schemes presented in the literature are based on simple tree construction primitives. Ranging from simple data collection and dissemination to complex virtual coordinate based point-to-point routing, tree-construction owns a history of hard-won successes in wireless communications. It has established itself as a building block for most location independent routing protocols.

However, most tree-construction based routing schemes put excessive focus on routing stability to maintain stable trees and addressing, while adaptability performance gets compromised to a large extent. The underlying technique is to employ a long-term link estimator [4] and restrict communication to neighbors with constantly high quality links. Although this results in a consistent routing topology and stable addressing, it circumscribes the protocol from using interesting communication opportunities provided by intermediate (short-term stable) links. Similarly, tree based addressing and routing infrastructures suffer heavily from rapid topological changes and varying link conditions.

In this paper we show how to retain the benefits of tree based routing and addressing schemes without maintaining explicit trees in the network. In contrast to existing ideas, our approach, we call *Statistical Vector Routing* (SVR), does not rely on long-term link estimation. It assigns probabilistic addresses to the nodes in the network. The basic idea is that a node learns from its past and calculates a probability distribution over its recent locations. All other nodes in the network predict the current location of a node in its distribution. A node's location is defined in terms of the probability that it exists in a certain location independently from packet loss at shorter time scales. As a result, SVR decouples addressing from routing, thereby allowing to adapt routing paths to the very recent network conditions.

II. STATISTICAL VECTOR ROUTING

In traditional virtual coordinates based routing, some number of nodes in the network would advertise themselves as landmarks. All other nodes would establish a (multihop) connection to each landmark in the form of a tree with *hop count* as a routing metric and establish an r dimensional address vector of the form $< q_1, \ldots, q_r >$, where q_i is the hop distance of a node q from the ith landmark and r is the total number of landmarks in the network.

However, in SVR, a node's virtual coordinates are expressed in the form of distributions. If a node knows that it can reach a landmark in the network over multiple paths, it will not derive its coordinate component for that landmark by selecting the best path in terms of the offered quality and the number of hops. Rather, it will represent its coordinate component in the form of a probability function that expresses all the paths and the relative frequencies at which they are available. Hence, the notion of path quality is automatically embedded in SVR's coordinates. These distributions have to be published completely, if no suitable smaller representation for them can be found. E.g. assuming the hop distances for each node approximately follow a normal distribution a formal representation for the components q_i of the virtual coordinates could be as follows:

$$q_i(c) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-(c-\mu_i)^2/2\sigma_i^2} \tag{1}$$

Where $q_i(c)$ is the normal distribution of the different hop counts c to the ith landmark. The basic parameters that would only have to be published in this case are mean (μ_i) and variance (σ_i):

$$\mu_i = \frac{1}{N} \sum_{n \in N} h_{q_i,n} \tag{2}$$

$$\sigma_i = \frac{1}{N} \sum_{n \in N} (h_{q_i,n} - \bar{q_i})^2 \tag{3}$$

where $h_{q_i,n}$ is the nth place in the coordinate history for q_i and N is the size of the gathered statistics, i.e. the *history size*.

These coordinate distributions remain independent of the data loss over a path. Although they are derived from regular beacon messages exchanged among nodes, in the long run, a node's coordinate distribution is supposed to stabilize eventually and not to be affected directly by instantaneous changes in the link conditions anymore. SVR's basic design goal is to decouple addressing from routing and to provide a consistent routing topology even in the existence of links with highly variable qualities. Moreover, SVR can take advantage of intermediate links by embedding the information regarding all possible node locations in its address distribution. Thus, it provides the opportunity to utilize these links that are very important for routing [5], rather than limiting forwarding to long-term stable links.

A. Algorithm

After presenting SVR's philosophy, we now discuss how to achieve such an addressing algorithmically within a routing infrastructure. We need two ingredients: (1) *Beacon Exchange* among neighbors to share addressing information periodically, and (2) *Coordinate Calculations*, performed by each neighbor locally at the end of each beacon interval.

1) Beacon Exchange: Beacon exchange among neighbors is a key requirement to establish a scalable tree-based routing infrastructure. Our goal is not to present a new beacon exchange mechanism, but to provide the details about the information that is exchanged via these beacons in SVR. Each node broadcasts a beacon message every interval t with the following information.

- **Sender ID:** The unique ID of the packet source.
- **Sequence Number:** A unique sequence number assigned by the source to each beacon packet.
- **Current Coordinates:** A vector of the minimum hop distance to each landmark in the last beacon interval. To reemphasize, there is no link quality information used to calculate this.
- **Traces**: For each landmark, a trace of the last five nodes on the path from which the current coordinates are derived are included. This information is important to avoid *routing loops* and *count to infinity* phenomena.
- **Neighbors**: This is a list of nodes from which the source node received a beacon packet in the last beacon interval. This is used to identify neighbors with symmetric links.

The use of *sender ID* and *sequence number* is trivial, i.e. to identify the source of beacons and uniquely identify each beacon from a particular source, respectively. The size of the beacon messages depends upon the number of landmark nodes in the network and the number of neighbors.

2) Coordinate Calculations: At the end of each beacon interval t, i.e. when a node is about to send its beacon message for the next interval, it performs the following operations.

- **Update Current Coordinates:** A node derives new *current coordinates* simply by selecting the minimum hop distances to each landmark in the network. The only piece of intuition here is that for calculating these minima a node will only include neighbors with which it shares a symmetric link.

- **Update Statistical Vector (SVR coordinate):** A node updates its SVR coordinate, calculated from the history of the last m current coordinates[1], after including the coordinate vector calculated during the current beacon interval and removing the oldest one from the history.

It is worth mentioning that there exists an error value for the calculations of the SVR coordinate, which determines the threshold to decide whether the difference between the newly calculated coordinate distribution and the previous one is significant and hence requires an update in the global coordinate database. If the difference is negligible, no overhead is caused.

The exclusive use of minimum hop distances in calculating the current coordinates is an algorithmic choice, as the coordinate distribution can also be derived by selecting all available hop distances in each beacon interval. However, we avoid this for two reasons: (1) We want our virtual coordinate distribution to be dominated by the smallest paths, and (2) we want to avoid loops and unnecessarily long and useless paths. After calculating the new SVR coordinate, a node will only trigger an update in the address database if the previous and the newly calculated SVR coordinate deviate by more than a certain threshold e.

B. Routing

The first step in routing is to elucidate that SVR's coordinate distributions are meaningful addresses that can be used to derive routing decisions. The simplest way to use them is to calculate the $mean$ for each distribution corresponding to a node's address vector component. We need a distance function to find a neighboring node as a next hop that minimizes the remaining virtual distance to the destination. The key to this selection is to find a neighbor whose coordinates are most similar to the destination node. We propose the *sum distance* $\delta_k^s(\bar{p}, \bar{d})$ metric to calculate the remaining distance. The $\delta_k^s(\bar{p}, \bar{d})$ between $\bar{p}$ and $\bar{d}$ is calculated by deriving the number of hops from $\bar{p}$ to a landmark and from there to $\bar{d}$ and averaging this over all landmarks.

$$\delta_k^s(\bar{p}, \bar{d}) = \frac{1}{|C_k(d)|} \sum_{i \in C_k(d)} (\bar{p}_i + \bar{d}_i) \qquad (4)$$

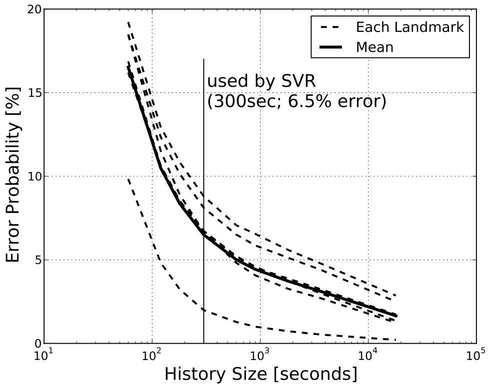

Fig. 1. Results of Pearson's χ^2-Test when calculating the SVR coordinate for different history sizes m and 10 seconds length of beacon intervals.

where $\bar{p}_i$ represents the mean distance of neighbor p to landmark i and $C_k(d)$ is the set of the k closest landmarks to d [1].

III. EVALUATION

The current implementation of SVR is in TinyOS 2.x for the IEEE 802.15.4 based Tmote Sky platform. Our evaluation focuses on two aspects: (1) Experimentally deriving the history size m for calculating SVR's coordinate distributions (see Section II-A.2) and (2) thoroughly comparing SVR with an existing virtual coordinates based addressing approache (regarding address stability and routing performance) to observe potential benefits and drawbacks of our approach.

A. Pearson's χ^2-Test

In order to derive the coordinate history m, we ran SVR with six landmarks on the TWIST testbed, a standard sensornet testbed deployment at the TU Berlin, with a transmission power level of $-15dBm$. We use Pearson's χ^2-Test, a test of goodness of fit, which shows how much two distributions differ from one another. Our goal is to calculate a $p-value$ that measures the probability of the deviations between two distributions to be caused by chance or by real substantial differences.

Figure 1 shows the average $p-value$ for different history sizes. It points out an initially rapid decrease in the error probability when we increase the history size. However, later increases do not substantially impact the error probability. For example, when increasing the history size from 60 to 300 seconds (i.e. from 6 to 30 beacon intervals), the $p - value$ decreases from 17% to 6.5%. Thereafter, increasing

---

[1] The history size m corresponds to multiple beacon intervals. In our experiments these intervals had a length of 10 seconds, so the history size can also be expressed by $10m$ seconds.

the history size from 300 seconds to 1,000 seconds (30 to 100 intervals) only results in a 2% decrease while significantly dampening the adaptability of coordinates and increasing the memory overhead for computing the SVR coordinates. Our cutoff is therefore at an error probability of 6.5%, which gives us the history size m of 300 seconds, i.e. 30 beacon intervals, for our experiments. For the remaining evaluation in this paper, we calculate the SVR coordinate from a history of 30 beacon packets.

B. Comparison with Beacon Vector Routing (BVR)

Now we thoroughly compare SVR with the *Beacon Vector Routing (BVR)* protocol. This is a state-of-the-art point-to-point routing protocol also using virtual coordinates. We base our evaluation on the following factors:

- **Coordinate Change Rate:** This is the rate of changes in the coordinates in SVR and BVR. It is our key evaluation aspect to show the stability of the coordinates over time.
- **Hope Distance:** This is the nodes' average hop distances from the landmarks.
- **Coordinate Range:** This is the difference between the maximum and minimum hop distance over time, i.e. the magnitude of change. For each component of the address vector the range of change from each landmark is analyzed.

Figure 2 shows the coordinate stability results from the TWIST testbed[2]. SVR achieves three times more stable coordinates than BVR. It reduces the hop distance from landmarks by 10-15%. The range of coordinates is reduced by 2-10 times depending on the testbed. Concluding our comparative evaluation, we have seen that SVR makes significant strides in enhancing the efficiency of tree-construction based virtual addressing in wireless networks. It shows that stable addressing across the network can be achieved without compromising the adaptability of virtual coordinate based routing, which has been the trade of existing routing approaches for a long time.

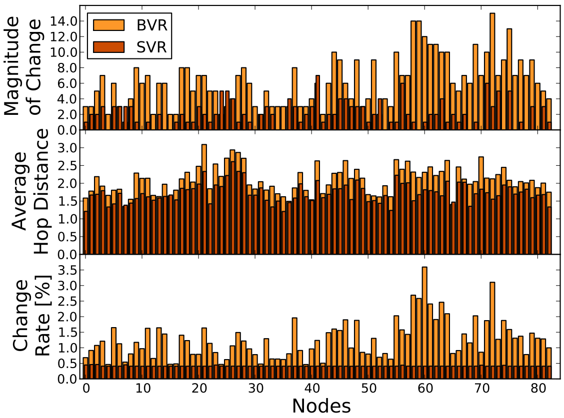

Fig. 2. Summary of the coordinate stability evaluation results from the TWIST testbed.

As our prototype implementation is for sensor networks our two main performance metrics from a routing perspective are the number of hops and transmissions required by a packet to reach its destination. Our results from real deployments indicate a 10-15% reduction in the transmission costs by adapting the routing paths for each packet. These improvements are also accompanied by a reduced packet loss and are achieved inherently due to SVR's design and the associated routing strategy. The results prove the feasibility of SVR in principle, however, we strongly believe SVR's transmission costs to benefit from more sophisticated routing metrics.

IV. FUTURE WORK

This paper presents the basic design, implementation, and evaluation of Statistical Vector Routing (SVR). We are still in the active development phase of our work. Thorough evaluation of SVR's routing performance in realistic simulation and testbed environments is mandatory to further measure the effectiveness of the concepts presented here. Our future work will mainly be to look into different routing metrics that can exploit the path quality information embedded in SVR's coordinate distributions.

REFERENCES

[1] Rodrigo Fonseca, Sylvia Ratnasamy, Jerry Zhao, Cheng T. Ee, David Culler, Scott Shenker, and Ion Stoica, "Beacon vector routing: Scalable point-to-point routing in wireless sensornets," in *NSDI*, May 2005.

[2] Qing Cao and Tarek Abdelzaher, "Scalable logical coordinates framework for routing in wireless sensor networks," *ACM Trans. Sen. Netw.*, vol. 2, no. 4, pp. 557–593, 2006.

[3] Ananth Rao, Sylvia Ratnasamy, Christos Papadimitriou, Scott Shenker, and Ion Stoica, "Geographic routing without location information," in *MobiCom '03*, 2003.

[4] Rodrigo Fonseca, Omprakash Gnawali, Kyle Jamieson, and Philip Levis, "Four-bit wireless link estimation," in *HotNets*, 2007.

[5] Muhammad Hamad Alizai, Olaf Landsiedel, J Agila Bitsch Link, Stefan Goetz, and Klaus Wehrle, "Bursty traffic over bursty links," in *SenSys'09*, Nov. 2009.

---

[2] We ran our experiments on the MoteLab (Harvard University) and Indriya (National University of Singapore) testbeds as well and gained similar results.