YOU'RE CROSSING THE LINE: LOCALIZING BORDER CROSSINGS USING WIRELESS RF LINKS

[1]*Peter Hillyard,* [1,3]*Neal Patwari,* [2]*Samira Daruki,* [2]*Suresh Venkatasubramanian*

[1]Electrical and Computer Engineering, University of Utah
[2]School of Computing, University of Utah
[3]Xandem Technology

ABSTRACT

Detecting and localizing a person crossing a line segment, i.e., border, is valuable information in security systems and human context awareness. To that end, we propose a border crossing localization system that uses the changes in measured received signal strength (RSS) on links between transceivers deployed linearly along the border. Any single link has a low signal-to-noise ratio because its RSS also varies due to environmental change, (e.g., branches swaying in wind), and sometimes does not change significantly when a person crosses it. The redundant, overlapping nature of the links between many possible pairs of nodes in the network provides an opportunity to mitigate errors. We propose new classifiers to use the redundancy to estimate where a person crosses the border. Specifically, the solution of these classifiers indicates which pair of neighboring nodes the person crosses between. We demonstrate that in many cases, these classifiers provide more robust border crossing localization compared to a classifier that excludes these noisy, redundant measurements.

1. INTRODUCTION

Knowing when and where people leave one region and enter another is an important piece of information in an age of increasing security and human context-aware computing systems. A person illegally crossing a national border, a driver passing through an intersection, or a shopper entering an aisle in a store are examples of "border crossings," i.e., people moving from one region to another by crossing the line segment that separates them.

In this paper, we present methods for localizing a person crossing a border by measuring overlapping line segment crossings. A related idea is shown in movies like *Entrapment*, *Ocean's Twelve*, and *The Return of the Pink Panther*, in which a seemingly impossible-to-bypass mesh of laser beams serves as a security system, and blocking any one laser triggers an alarm. Instead of lasers, we propose using received signal strength (RSS) measurements from a radio frequency (RF) link to detect line segment crossings. As in [1, 2, 3, 4, 5], we conceptualize a radio link as a *link line*, i.e., the line segment between a transmitter and receiver. Compared to a laser beam, an RF link experiences changes in RSS in a relatively large area around the link line. With some RF links, that area around the link line is small. The RSS on these links experience large changes when a person is on the link line; but when a person is far away from the link line, the RSS shows very little change. However, other RF links can experience large changes even when a person is far away from the link line; moreover, the link can experience very little change even when the person is on the link line (see Fig. 1). For our purposes it suffices to say that these "noisy" links have low signal and high noise and thus are less reliable as link line crossing detectors.

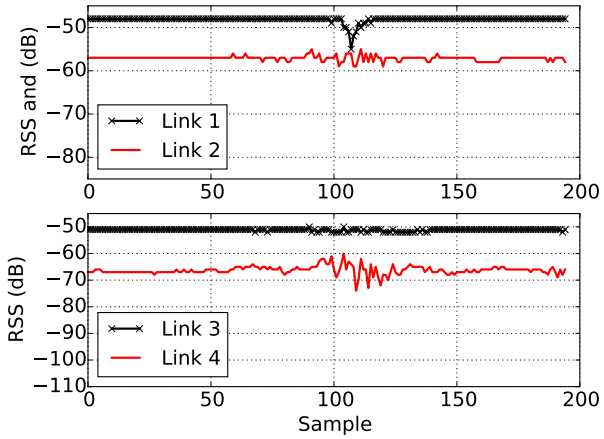

Fig. 1. (top) The drop in Link 1 RSS at sample 108 correctly indicates that a person is crossing the link line. Link 2 sees small RSS changes and fails to detect the crossing. (bottom) No person is near the link line. Link 3 shows little change, but link 4 shows a significant drop in RSS at sample 108, which would incorrectly be detected as a crossing.

In this paper, we propose using multiple, overlapping links of varying signal-to-noise ratio (SNR) in methods that not only answer the question of when a border is crossed, but where the border is crossed. To accomplish these goals,

we deploy N nodes linearly along a border such that links form between pairs of nodes (see Fig. 2). Nodes k and l form a link (k, l). Each link has an associated line segment between the positions of nodes k and l, which we call the link line. We denote $\mathcal{L}$ as the set of all links $(k, l), k \neq l$, which are measured in the network. A special subset of $\mathcal{L}$ is what we call the set of "short segments", the links $(j, j+1)$ for $j = 1, \ldots, N-1$. Short segments have neighboring nodes as endpoints. As shorthand notation, we call the link line on $(j, j+1)$ "short segment j".

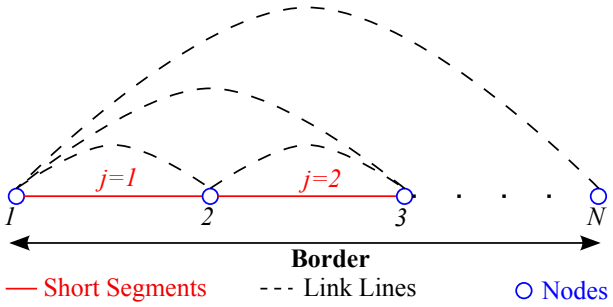

Fig. 2. Border crossing system with N nodes aligned linearly. Short segments are created between neighboring pairs of nodes. Link lines are curved in this figure for visual purposes only. In practice, the link lines are viewed as line segments with the nodes as their endpoints.

First, consider a network in which we measure RSS only on the short segments j, for $j = 1, \ldots, N-1$, and have as a goal to estimate which one of these, if any, were crossed by a person. Any single false alarm or missed detection among the $N-1$ links would cause the system to be unable to determine the crossed short segment.

Next, consider a network in which we measure links (k, l) for many overlapping links $\mathcal{L}$. When a person crosses through a particular short segment j, they also must simultaneously cross some other longer link lines associated with links (k, l) for which $k \leq j < j+1 \leq l$ (see Figure 3). In this paper, we

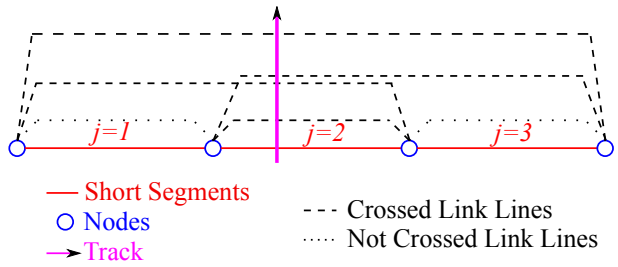

Fig. 3. A certain set of links are crossed when a person passes over short segment s_2.

propose that a system that exploits the measurements of the redundant links can be more robust to single link errors. We compare, via experimentation, the redundant links network to the short segment-only network in terms of classification accuracy.

With our goal of analyzing the robustness of both systems, we compare the accuracy of three crossing segment classifiers we propose in this paper. The closest codeword classifier (CCC) and the maximum a posteriori classifier (MAPC) leverage the measurements of the redundant link; the simple classifier (SC), excludes these measurements. Although our focus is on the performance of these classifiers, we must evaluate them by implementing link line crossing detectors (LLCD) [1, 2, 3, 4, 5]. Algorithms in a LLCD use link RSS measurements to produce link line crossing measurements x_i for each link $i \in \mathcal{L}$. The output x_i is a binary 1 when a crossing is detected and 0 otherwise. A vector $\mathbf{x} = [x_1, \ldots, x_L]^T$ from $L = |\mathcal{L}|$ links is then fed into a crossing segment classifier (see Fig. 4). In our paper, a LLCD outputs one value for x_i for RSS measurements collected during time interval T_{max}, where r_i is an RSS measurement for link i.

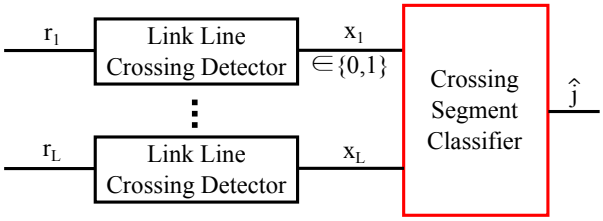

Fig. 4. Block diagram of a border crossing localization system.

We discuss the details of each classifier in Section 2, compare classifier accuracies through two experiments in Section 3, discuss the results, and conclude.

2. METHODS

In this section, we propose two crossing segment classifiers that use the noisy, redundant binary measurement vector $\mathbf{x}$ to classify which short segment j a person crossed while passing over a border, or if no crossing took place. In addition, we describe a third simple classification method that excludes the measurements from the redundant links in its decision. We first describe a LLCD that feeds binary vector measurements into the classifiers.

2.1. Detecting Link Line Crossings

In order to evaluate our classifiers, we first implement a LLCD. Although any LLCD could be used, we choose to implement the moving average based detection method from Section 4.3.1 in [2] because of its straightforward implementation and its accuracy in detecting link line crossings. In this method, RSS measurements made on a link i feed into a LLCD as shown in Figure 5. The moving average based detector adds RSS measurements to a short and long term

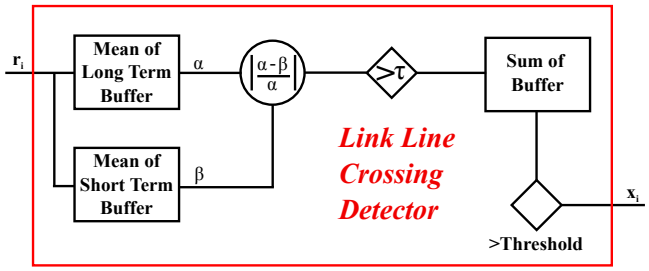

Fig. 5. Moving average based detection [2]: Detector is 1 if the difference between short-term and long-term average exceeds τ multiple times during a time interval.

buffer. The long term buffer stores the static behavior of the link while the short term buffer stores the current behavior. Upon adding a new RSS measurement to the buffers, the detector computes the relative difference between the means of the two buffers. When the relative difference exceeds a threshold τ, an event is detected. These events are stored in a buffer that holds T_{max} amount of samples. When full, the buffer is summed to count the number of events that occurred during T_{max}. If the number of events exceeds another threshold, then $x_i = 1$. The buffer is then cleared to receive the next T_{max} events. A LLCD is created for each link, and we store the binary measurements in $\mathbf{x}$.

If the LLCDs measured link line crossings without error, then a person crossing a particular short segment of the border would result in a unique $\mathbf{x}$. We call these unique binary measurement vectors *codewords* where the codeword for a person crossing short segment j is $\mathbf{w}(j) = [w_1(j), w_2(j), \ldots, w_L(j)]^T$ where $w_i(j) = 1$ if by crossing j, the person also crosses link i, and 0 otherwise. Formally,

$$w_i(j) = \begin{cases} 1, & \text{if } (l_i \leq j) \text{ and } (j+1 \leq k_i) \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where k_i and l_i are the endpoints of link i, and $k_i < l_i$, without loss of generality. We denote the "off" codeword $\mathbf{w}(0)$ whose elements are all 0. We define short segment 0 as the class that no person crossed the border. As a result, an N node system will have N unique codewords, which we denote as set $\mathcal{W} = \{\mathbf{w}(j)\}_{j=0}^{N-1}$.

2.2. Closest Codeword Classifier

The first classifier we propose, the CCC, finds the codeword that is closest to $\mathbf{x}$ in terms of Hamming distance. The CCC classifies which short segment j is crossed using

$$\widehat{j} = \arg\min_j \|\mathbf{x} - \mathbf{w}(j)\|^2 \quad (2)$$

where $\|\cdot\|^2$ is the ℓ_2 norm. Note that if there are no errors in any of the L link line crossing detectors, then $\mathbf{x} = \mathbf{w}(j)$ for the correct short segment j which was crossed. If $\widehat{j} = 0$ then the CCC classifies that there was no crossing.

In some instances, more than one codeword will minimize (2), thus the CCC does not definitively classify which of the short line segments were crossed. Under these conditions, the CCC randomly chooses one of the short line segments whose codeword satisfies (2). Finding the closest match between $\mathbf{x}$ and $\mathbf{w}(j)$ can be viewed as error correction where $\mathbf{x}$ differs from the true codeword $\mathbf{w}(j)$ in γ places. The CCC corrects γ errors when the detector chooses the correct crossed segment.

2.3. Maximum a Posteriori Classifier

Here, we propose the MAPC, which builds upon the CCC by adding in prior knowledge of the accuracy of each LLCD. We capture the accuracy of the detector for link i by its probability of detection, P_i^d, and its probability of false alarm, P_i^{fa}. These probabilities help us make more informed decisions about which short line segment a person crossed based on $\mathbf{x}$.

To learn the prior probabilities, as a proof of concept, we assume that the network is trained by a person crossing each short segment multiple times. The value of P_i^d and P_i^{fa} are then estimated from the training data. To prevent the MAPC from assuming any particular measurement provides certainty, we limit P_i^d and P_i^{fa} to be in the range $\epsilon < P_i^d, P_i^{fa} < 1 - \epsilon$ for some small $\epsilon > 0$. In our experiments, we use $\epsilon = 0.0001$. Future work could explore methods to predict performance via statistics that can be measured without training, e.g., RSS mean or variance.

To formulate the MAPC, we first formulate the likelihood of measurement $\mathbf{x}$ given that a person crossed short segment j. Under the assumption of conditional independence of links, this is $p(\mathbf{x} \mid j) = \prod_{i=1}^{L} p(x_i \mid j)$, where

$$p(x_i \mid j) = \left[(P_i^d)^{x_i}(1-P_i^d)^{1-x_i}\right]^{w_i(j)} \cdot \left[(P_i^{fa})^{x_i}(1-P_i^{fa})^{1-x_i}\right]^{1-w_i(j)}. \quad (3)$$

Next we consider the prior probabilities of j. If the N nodes are equally spaced along the border, this might make crossing each short segment j equally probable. If, however, N nodes are unevenly distributed along the border, we may want to assume that crossing a longer short segment is more probable than crossing a shorter short segment. We may also want to modify how likely crossing j is given environmental factors, e.g. thick vegetation, hills, etc. To account for these conditions, we impose a prior, $p_J(j)$, on each j. With a prior on j, we form the more general MAPC, formally defined as

$$\widehat{j} = \arg\max_j p_J(j)\, p(\mathbf{x} \mid j). \quad (4)$$

Note again that if multiple j have equal joint probability, we randomly choose one of them.

2.4. Simple Classifier

A more basic method to classify which short segment was crossed would be to eliminate all the redundant links and measure only the short segments. The third classifier, i.e. SC, in contrast to the CCC and MAPC, only measures link line crossings on $(j, j+1)$ for $j = 1, \ldots, N-1$. If a LLCD for short link i measures a 1 for x_i, then the system would classify short segment i was crossed. In the event that more than one LLCD outputs a 1, SC randomly selects from the candidate short segments for its solution. In like manner, if none of the LLCDs for the short link lines report a crossing, then SC selects the 0 class, for no crossing, as its solution. We show in the next section how each classifier compares in accuracy.

3. EXPERIMENTAL VERIFICATION

In this section we show the accuracy of the CCC and MAPC compared to the SC using two experimental campaigns.

3.1. Equipment and Procedures

The wireless nodes used in the following experiments are Texas Instruments CC2531 USB dongles, each of which transmits at 2.4 GHz with $+4.5$ dBm. Following a TDMA protocol, each node takes a turn transmitting a packet while the others receive the packet. An additional listen node is connected to a laptop to overhear the wireless traffic and to log RSS measurements for each link. Each node is placed on a tripod that stands 0.91 meters high.

Two experimental sites are chosen to show the accuracy of each classifier in different environments. The first site is in a natural area with many trees. Nine nodes are deployed such that they are 4 meters apart (see Figure 6). The second site is inside the Union cafeteria at the University of Utah by lunch tables and large pillars. The indoor environment induces more complex fading patterns and is similar to indoor border crossing scenarios (e.g. the aisles of a store). Nine nodes are again equally spaced along the 29.26-meter border (see Figure 7).

To test with fewer than nine nodes, we keep the deployment along the same total border length, and four nodes are removed such that five evenly spaced nodes remain. Next, two of the remaining five nodes are removed, leaving three evenly spaced nodes. Under these conditions, the data from the nine node configuration can be used for the five and three node case by removing the appropriate links' RSS measurements.

In each experiment, a border crosser indicates into a voice recorder the start of the experiment. RSS values for all of the links are then logged in a computer as the person crosses the border. The border crosser then indicates into the voice recorder the end of the T_{max}-long crossing experiment and the computer stops logging the RSS values. This process is repeated thirty times for each short segment j, for $j = 1, \ldots, N-1$. In addition to crossing each short segment, the border crosser then walks back and forth 2.4 meters away from and parallel to the border, so that we can evaluate what happens when a person walks near but never crosses the border. The above procedures are carried out twice, first for a training data set, and second for a testing data set.

Fig. 6. Experiment in outdoor environment

Fig. 7. Experiment near cafeteria tables

In this paper, the LLCDs use the same best parameters recorded for the long term and short term buffers in the last of Section 3.4.1 of [2] in both the indoor and outdoor experiments. However, we vary the threshold τ to achieve varying levels of detection and false alarm accuracy for the LLCDs. We compare the accuracy of the classifiers using the following metrics: the probability of correct classification (P_{cc}), meaning $\widehat{j}$ matches the true short segment crossed; and the probability of border crossing false alarm (P_{fb}), meaning $\widehat{j} \neq 0$ when the border was not crossed. We envision that a border crossing localization system should have a low P_{fb} (at the expense of a lower P_{cc}) in some applications, because a false alarm could result in wasted time and resources devoted to investigate a falsely reported border crossing.

3.2. Experiment 1 (Outdoor) Results

In this section, we compare the probability of P_{cc} and P_{fb} as a function of τ using the MAPC, CCC, and SC in the outdoor environment experiment. We show the relationship between τ and these probabilities in Figure 8 and 9 where five nodes were deployed. Figure 8 shows that as τ approaches 0, P_{fb} increases. A low τ causes an LLCD to frequently output a binary 1 even when the person does not cross the link. When many of the LLCDs output a one, the classifiers tend to decide that multiple short segments were crossed. Thus, even when the border is not crossed, the classifiers yield a high false alarm classification probability. Moreover, when the border is crossed, the classifier must randomly choose between the classified crossed segments, thus lowering P_{cc} (see Figure 9).

In contrast, when τ is large, few, if any, link lines will be detected as crossed. When many of the LLCDs output a zero, the classifiers tend to decide that no short segments were crossed, i.e. $\widehat{j} = 0$. This reduces P_{fb}, but we also observe a

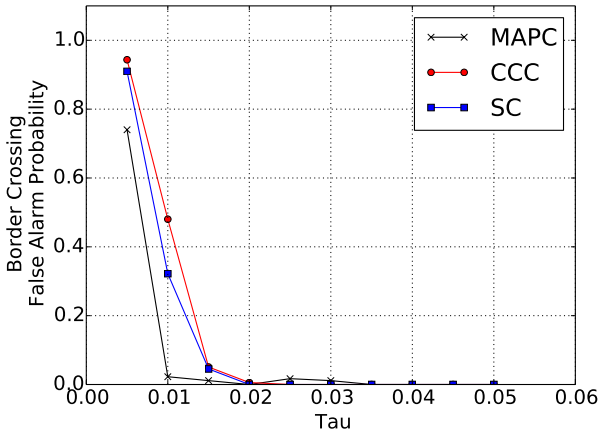

Fig. 8. Probability of border crossing false alarm as a function of τ for $N = 5$.

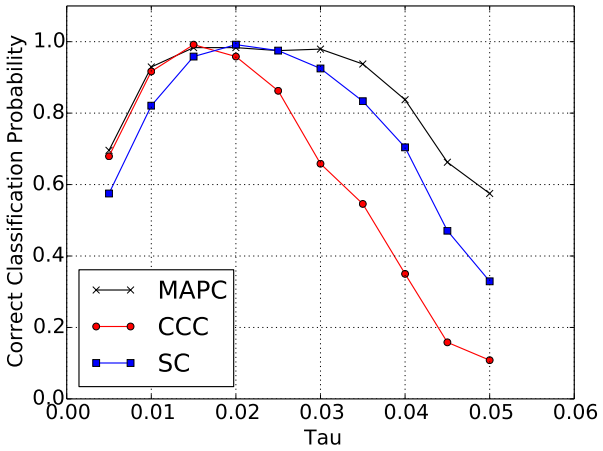

Fig. 9. Probability of correct classification as a function of τ for $N = 5$.

decrease in P_{cc} since we are unlikely to classify the correct crossed short segment.

For the remainder of this paper, we let τ be the value that maximizes P_{cc} while keeping $P_{fb} \leq 0.01$ for each classifier. In Figure 10, we show P_{cc} for each classifier when three, five, and nine nodes are deployed.

Both the MAPC and the CCC improve in P_{cc} as the number of nodes increases, achieving almost perfect classification for $N = 5$ and 9. Adding more nodes to the system creates greater distance between codewords and $\mathbf{x}$ and therefore allows us to make more correct classifications. We also observe that the MAPC outperforms the CCC for $N = 3$. By adding in probabilistic conditions in the MAPC, we improve the accuracy compared to the CCC, which does not take into account these probabilities. When we compare the SC to the MAPC and CCC, we observe that the SC achieves a comparable P_{cc} to the MAPC for $N = 3$ and 5, but suffers when

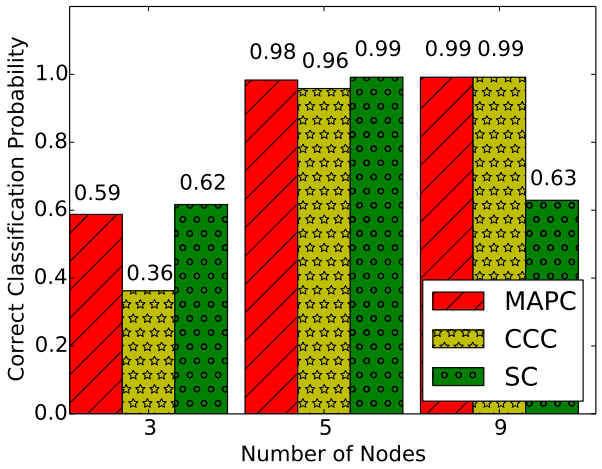

Fig. 10. Probability of correct classification in the outdoor environment using the MAPC, CCC, and SC with a 0.01 probability of border crossing false alarm.

$N = 9$. So whereas MAPC and CCC strictly increase in P_{cc} with more nodes, the SC can perform worse. The SC relies on the short link lines for classification, and when one or more short link lines are poor line segment crossing detectors, the SC can suffer in accuracy.

3.3. Experiment 2 (Indoor) Results

In this section, we again compare P_{cc} of the three classifiers by using the τ that maximizes P_{cc} while keeping $P_{fb} \leq 0.01$ (see Figure 11). We observe dramatically reduced P_{cc} for al-

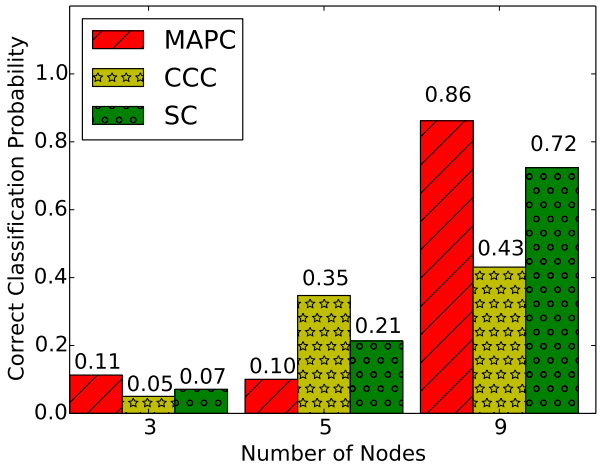

Fig. 11. Probability of correct classification in the indoor environment using the MAPC, CCC, and SC with a 0.01 probability of border crossing false alarm.

most all node configurations and classifiers compared to the outdoor case. One explanation for this reduction is that an

indoor setting introduces more multipath fading than an outdoor environment, which in turn causes more links to be poor line segment crossing detectors.

In spite of the significant drop in P_{cc}, we observe the same general improvement in classification as N increases for the MAPC and CCC. We also observe that the SC shares this same improving behavior, which we did not observe in the outdoor case. It is probable that all short link lines were sufficiently reliable line segment crossing detectors in this setting and SC improved as a result with increasing N.

A surprising result is that the CCC and SC perform twice as well as the MAPC when $N = 5$. The MAPC uses probabilities estimated from training data, and we would therefore anticipate that the MAPC would always perform at least as well as the CCC and SC. However, the probability estimates may be inaccurate. The inaccuracies are influenced by the places a person walks with relation to the border during training. Ideally, we would want a person to train the MAPC by visiting several locations near and on the border and then record the link accuracies. But this would take a significant amount of time to train. When the estimates are close to the true probability, the performance gain of the MAPC can be as great as two times, as in the $N = 9$ case.

In a practical outdoor border crossing localization system, we might conclude that using the CCD would be a ideal choice since it achieves a high accuracy and does not need to be trained. However, if the number of nodes were limited, the SC may be the better choice for its accuracy and plug-and-play nature. However, in the more complicated indoor environment, the MAPC would achieve a higher P_{cc} using a large number of nodes. Although the MAPC must be trained, this may only have to be done occasionally.

4. CONCLUSION

In this paper, we proposed two new classifiers that provide robust border crossing localization using RSS measurements on redundant RF links. Each classifier localized a border crossing by deciding which short segment was crossed. The CCC obtained near perfect classification at a 0.01 false alarm rate in the outdoor border for five and nine nodes and did not have to be trained. The MAPC matched or exceeded the CCC in accuracy but required a training period. The SC achieved the highest accuracy with three nodes, but can degrade as the number of nodes increases. In the indoor environment, however, nine nodes were required for any classifier to reach adequate accuracy; the MAPC reached a 0.86 probability of correct classification while the SC reached 0.72. The CCC was not a reliable classifier in this environment in any of the node configurations. Future experiments will test performance in other environments and with higher numbers of nodes.

We found that the number of nodes and the border environment were critical components in determining which of these three classifiers would be best. We found that leveraging redundancy improved accuracy when the number of nodes was high. However, when fewer nodes were deployed and thus little redundancy was present, methods that exploit redundancy performed as well or worse than a method that excluded redundant measurements. For an indoor border localization system, none of the classifiers worked well when three and five nodes were deployed. However, when we used nine nodes, using the redundant measurements provided more robustness despite noisy links.

In summary, we have demonstrated that when many nodes are employed in either indoor or outdoor environments, it is in our best interest to use a classifier that leverages redundancy.

5. ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant Nos. #1407949 and #1035565.

6. REFERENCES

[1] Ossi Kaltiokallio and Maurizio Bocca, "Real-time intrusion detection and tracking in indoor environment through distributed rssi processing," in *2011 IEEE 17th Intl. Conf. Embedded and Real-Time Computing Systems and Applications (RTCSA)*, Aug. 2011, vol. 1, pp. 61 –70.

[2] Moustafa Youssef, Matthew Mah, and Ashok Agrawala, "Challenges: device-free passive localization for wireless environments," in *MobiCom '07: ACM Int'l Conf. Mobile Computing and Networking*, 2007, pp. 222–229.

[3] Yang Zhao, Neal Patwari, Jeff M. Phillips, and Suresh Venkatasubramanian, "Radio tomographic imaging and tracking of stationary and moving people via kernel distance," in *Proceedings of the 12th international conference on Information processing in sensor networks*, New York, NY, USA, 2013, IPSN '13, pp. 229–240, ACM.

[4] J. Wilson and N. Patwari, "A fade-level skew-laplace signal strength model for device-free localization with wireless networks," *Mobile Computing, IEEE Transactions on*, vol. 11, no. 6, pp. 947–958, 2012.

[5] Yi Zheng and Aidong Men, "Through-wall tracking with radio tomography networks using foreground detection," in *Wireless Communications and Networking Conference (WCNC), 2012 IEEE*. IEEE, 2012, pp. 3278–3283.