On the Deployment of a Hybrid Free-space Optic/Radio Frequency (FSO/RF) Mobile Ad-hoc Network

Jason Derenick†, Christopher Thorne‡ and John Spletzer†
†Computer Science and Engineering
‡Mechanical Engineering and Mechatronics
Lehigh University
Bethlehem, PA 18015
{*jcd6, cet6, josa*}*@lehigh.edu*

***Abstract*— The hybrid free-space optics/radio frequency (FSO/RF) network paradigm promises new levels of throughput for sensor and mobile ad-hoc networks [1]. However, several challenges must be addressed before such a network model can be realized. These include a means by which deployed robots can autonomously establish optical links over sufficiently long distances as well as the formulation of a mobile network architecture that can exploit high throughput FSO channels.**

In this paper, we offer solutions to these problems in the form of hierarchical link acquisition and routing protocols. The heart of our link acquisition system (LAS) is a vision-based alignment phase for locating robot link partners via high zoom camera systems. Identification is accomplished in real-time using a multi-resolution image representation and normalized intensity distribution (NID) as a similarity metric. Our routing protocol relies upon the Hierarchical State Routing (HSR) model adapted to the FSO/RF paradigm. Experimental results from a large set of link acquisition trials as well as a small scale FSO/RF deployment are provided to support our approach.

***Index Terms*— Sensor Networks, Hybrid FSO/RF Network, Free-space optics, MANET**

I. Introduction

Consider the scenario where a natural or man-made event (*e.g.*, an earthquake) destroys critical infrastructure in an urban area, including fiber links in the Metropolitan Area Network (MAN). Patching these high-throughput channels using traditional radio frequency (RF) wireless nodes would be impractical due to bandwidth and link distance requirements. Instead, a team of robots could be automatically deployed to provide carrier grade patches to the MAN using free-space optic transceivers (FSO) - all under conditions unsafe for human operators. This is only one application we envision for the hybrid FSO/RF paradigm introduced in [1].

FSO is a commercial technology used primarily in static configurations for high bandwidth, wireless communication links [2]. With FSO, an optical transceiver is placed on each side of a transmission path to establish a network link. The transmitter is typically an infrared (IR) laser or LED that emits a modulated IR signal. Link availability can be maintained under most weather conditions (heavy fog is a notable exception, but unlike RF based technologies, heavy rain has little effect).

Perhaps the greatest advantage of FSO technology is its high throughput. The need for such a channel is evident when considering bandwidth hungry applications such as surveillance and monitoring. Commercial FSO transceivers currently provide throughputs of several Gbps with link distances of a kilometer or more [3]. Terabit per second throughputs have been demonstrated under laboratory conditions. In contrast, widespread RF technologies (*e.g.*, 802.11x) are limited to link throughputs on the order of 10s of Mbps across distances of 10s of meters. Even much anticipated ultra wideband (UWB) technology - with throughputs of 100s Mbps - drops to levels lower than 802.11a at modest ranges ($r \geq 15m$) [4].

The major shortcoming of FSO technology is the requirement for optical links to maintain line-of-sight (LOS). As such, we do not suggest that FSO technology will replace RF wireless communications in mobile robotics applications. Instead, the two technologies can serve complementary roles in a hybrid FSO/RF mobile ad-hoc network (MANET).

II. Background and Related Work

Sensor networks and mobile ad-hoc networks (MANETs) have become burgeoning research areas during recent years in both the networking and robotics communities. Significant work has been done in the areas of network coverage [5], [6], [7], localization [8], network deployment [9], and navigation [10].

However, the primary motivation for our research is network bandwidth constraints. Traditional networks are bound by the provable limits in per-node throughput for radio frequency (RF) based communications. In [11], the authors proved that when n nodes are optimally distributed on the plane, the asymptotic bound for per-node throughput was $\lambda(n) = \Theta(Wn^{-\frac{1}{2}})$ where each node was capable of transmitting W bits-per-second. However, this result is based upon a non-interference protocol. To protect against the potential for interfering with neighbors' broadcasts, individual node transmissions were explicitly scheduled. As a consequence, the throughput at each node is reduced by a factor of $\frac{1}{c+1}$, where

c corresponded to the maximum number of interfering neighbors. For FSO, such a reduction is inappropriate. Links are directed and full-duplex, and as a result no such interference region exists.

We believe that this characteristic, in conjunction with the high throughput afforded by FSO, will provide an enabling technology that can expand the capabilities of wireless sensor networks.

III. A Hierarchical Link Acquisition System

Whether employed as primarily FSO links as in a disaster relief application or as a mobile ad-hoc backbone in an RF sensor network, a fundamental challenge for the hybrid nodes is establishing and maintaining optical links. FSO technology achieves exceptional ranges by employing laser transceivers with a beam divergence on the order of 0.6 milliradians. To place this in perspective, at 100 meters this corresponds to a cone diameter of approximately 6 cm. In commercial systems, transceivers typically require an initial alignment by technicians using a telescope assembly. Instead, we propose a hierarchical link acquisition system (LAS) relying upon a "coarse to fine" approach whereby a robot pair can automatically establish the FSO links.

We assume each robot has *a priori* knowledge of its link partner's objective position. Additionally, each is able to infer its relative pose (position and orientation) to a common reference frame. With our current platforms, this can be accomplished through GPS and inclinometer sensors. Under these assumptions, the link acquisition phase is decomposed into 3 stages:

1) Coarse alignment through local sensor measurements.
2) Refinement using vision-based robot detection.
3) Final FSO alignment for link acquisition.

In the third stage, we rely upon the internal FSO tracking system. Thus, our discussions are limited to the former two.

A. Coarse Alignment Phase

Let $x_1, x_2 \in \mathbb{R}^3$ denote the objective positions for the two robots in our navigation frame $\mathcal{W}$. We assume for now that the robots are able to accurately measure their position and orientation in $\mathcal{W}$. After migrating to their objective positions, each robot will adjust its relative orientation to roughly align the optical transceivers. If the coordinate frames $\mathcal{F}_i$ (for transceiver i) and $\mathcal{W}$ are equivalent up to a translation then the desired link directions would simply be

$$x_{ij} = \frac{x_j - x_i}{\parallel x_j - x_i \parallel},\ i,j \in \{1,2\},\ i \neq j \tag{1}$$

In practice, local ground contours will only allow for a coarse transceiver alignment. Based upon our assumptions, the orientation of each robot in $\mathcal{W}$, as defined by the rotation matrices $R_1, R_2 \in SO(3)$, can be inferred from local sensor measurements. With these, we can compute the objective link direction in $\mathcal{F}_i$ as $\hat{x}_{ij} = R_i^T x_{ij}$. Converting to spherical coordinates we obtain

$$\begin{bmatrix} \theta_{ij} \\ \psi_{ij} \end{bmatrix} = \begin{bmatrix} \arctan\left[\frac{\hat{x}_{ij}(3)}{\sqrt{(\hat{x}_{ij}(1)^2 + \hat{x}_{ij}(2)^2)}}\right] \\ \arctan\left[\frac{\hat{x}_{ij}(2)}{\hat{x}_{ij}(1)}\right] \end{bmatrix} \tag{2}$$

where θ_{ij}, ψ_{ij} denote the required pan/tilt angles of the optical transceiver on robot i to establish a link with robot j, and $\hat{x}_{ij}(k)$ denotes the k^{th} element of vector $\hat{x}_{ij}$.

Under ideal conditions, Equation 2 would be sufficient to establish the link. However, our estimates for robot pose may have significant errors in both position *and* orientation. Imperfect calibration of the pan/tilt heads will introduce additional uncertainty. As a result, we rely upon a second alignment phase to refine the relative orientation estimates of the two link partners. In doing so, we exploit the fact that the ability to establish an optical link implies a clear line-of-sight (LOS). This enables the robots to employ camera systems with high zoom capabilities to assist in the alignment process.

B. Vision-based Alignment Phase

At longer link distances, the uncertainty in robot orientation will most likely be larger than the solid angle subtended by its camera's field-of-view (FOV). As a consequence, it will often be necessary for the robot to search in azimuth and elevation with its camera's pan-tilt-zoom (PTZ) system to locate its link partner. For the vision-based alignment phase, we propose a two-dimensional pattern matching approach where the coarse estimates derived in Equation 2 serve as the starting point for this search.

A major limitation of two-dimensional pattern matching approaches for tracking three dimensional targets without scale or orientation constraints is that a huge family of templates is required. Such a "brute force" approach is computationally intractable. However, our problem is not so ill posed. First, the scale of our target is approximately known since we have *a priori* knowledge of the target position (the same effect could also be achieved by fitting each node with a time-of-flight transceiver to obtain relative range measurements). Second, the relative orientation of the robots are limited by both line of sight requirements and physical constraints inherent with the pan/tilt head. As a result, the set of orientations that would permit the establishment of optical links can be quite constrained. For many applications, relying upon a small (≈ 20) set of templates in conjunction with a robust tracking algorithm would be sufficient.

In our approach, each robot maintains in memory a set of image templates $\mathcal{T}$, where each template $T_{\theta\phi} \in \mathcal{T}$ corresponds to a reference image of its link partner acquired at camera azimuth and elevation (θ, ϕ), respectively. All templates are acquired at a fixed reference distance $d_{\mathcal{T}}$ and camera focal length $f_{\mathcal{T}}$.

The vision system then searches in azimuth and elevation to acquire a set of images $\mathcal{I}$, where each image $I_{\theta\phi} \in \mathcal{I}$ is associated with a candidate link partner location. The search

limits are dictated by the estimated uncertainty in robot pose. These can be inferred directly through sensor models, or (more appropriately) using the second order statistics of the probability density function in Kalman/particle filter based localization implementations. By setting the focal length of the vision system to

$$f_{search} = \frac{\parallel x_j - x_i \parallel}{d_{\mathcal{T}}} f_{\mathcal{T}} \tag{3}$$

we ensure the robot scale in captured images is consistent with the reference template set $\mathcal{T}$,

To detect the robot in a given image, we use normalized intensity distribution (NID) as a similarity metric [12]. An advantage of this formulation is that it explicitly models both changes in scene brightness and contrast from the reference template image. Additionally, in a comparison to alternate approaches in model based tracking, it was less sensitive to small motion errors [13]. Our matching approach, while static, will suffer from similar error effects in the form of inconsistencies between the reference template and the viewed target (link partner) position and orientation. We shall discuss this in further detail in the sequel.

Given an image $I \in \mathcal{I}$, an $m \times n$ template T of the robot, and a block region $B \in I$ of equivalent size corresponding to a hypothetical robot position, the similarity of T and B can be defined as

$$\epsilon(T, I, B) = \sum_{u=1}^{m} \sum_{v=1}^{n} \left[\frac{T(u,v) - \mu_T}{\sigma_T} - \frac{B(u,v) - \mu_B}{\sigma_B} \right]^2 \tag{4}$$

where $T(u, v)$ denotes the grayscale value at location (u, v), μ is the mean grayscale value, and σ the standard deviation. The image region

$$B^* = \arg\min(\epsilon) \;\; \forall \, B \in I, \; \forall \, I \in \mathcal{I}, \; \forall \, T \in \mathcal{T} \tag{5}$$

would then correspond to the target robot position in the image sequence. The corresponding estimates (θ^*, ϕ^*) would be used to direct the FSO transceivers for the final alignment stage.

Even for a single template/image pair, template based pattern recognition approaches are very computationally expensive using traditional matching techniques. In our paradigm, we rely upon such an approach for a *set* of templates over an image *sequence*. To implement this in real-time, we use a multi-resolution image representation [14]. Prior to the pattern recognition phase, images are converted into Gaussian pyramid data structures. Templates are also stored in memory using a similar representation. Pattern matching is then performed at the highest level of the pyramid, with only local refinements performed at each lower level. This allows for speedups of several orders of magnitude, and real-time template based pattern recognition.

The main advantage of the vision-based approach for aligning the FSO transceivers is the reduction of the dimension of the search space. We could in fact pan/tilt the FSO transceivers, and allow the internal alignment systems to search directly. The critical difference is while the vision systems can search independently, the transceivers must *both* be aligned to a given tolerance *simultaneously* in order to establish the link. This increases the search space from $\mathbb{R}^2$ to $\mathbb{R}^4$. Exponential increases in search times would result.

IV. A General Architecture for a FSO/RF MANET

The characteristics of FSO makes the technology well suited for use as a high throughput backbone in an RF sensor network. We envision that hybrid robot nodes with both FSO and RF interfaces could be so deployed. However, to fully exploit this technology requires a routing protocol adapted specifically to the link characteristics of both mediums. To accomplish this, we propose the distributed use of Hierarchical State Routing (HSR) utilizing intelligent cluster-head election policies.

A. Architectural Overview

An FSO/RF MANET can be logically decomposed into two subcomponents. The first subcomponent is a collection of RF-based ad-hoc networks that will be required to communicate in real-time over ranges that render the use of RF communication channels impractical. It is assumed (though not required) that each of the networks in the collection will be deployed beyond the radio transmission range of any other network in the collection. Logically, each ad-hoc network can be considered a hierarchically defined domain space characterized by a dynamic topology.

The second component is a high throughput backbone FSO network. This network serves as the cornerstone of the FSO/RF MANET architecture since its presence enables the dissemination of information among otherwise disjoint ad-hoc domains. It can be constructed dynamically by the hybrid node(s) that have been assigned as members of each respective ad-hoc network. The hybrid nodes serve to bridge the gap between the RF domain and the FSO backbone domain space. As such, the routing protocol employed by each RF-based network must account for this constraint so that all ad-hoc RF nodes have knowledge of where to forward packets destined for a non-local address.

In order to solve the routing problem in RF ad-hoc domains, we propose the use of Hierarchical State Routing (HSR) [15], [16]. Among the more appealing features of this protocol is its support for "logical" network partitions and differentiated services. The latter is enabled by the architecture's definition over a collection of physical network clusters. As such, its underlying theory lends itself nicely to the task of propagating time-critical information. It is also highly scalable, facilitating the support of a large number of wireless nodes.

The routing protocol used among the hybrid peers depends upon the number of hybrid nodes in the system. For an FSO backbone consisting of only a few hybrid nodes, flat link-state routing can be used. When the number of nodes increases, HSR can also be applied among these nodes, providing an additional level of hierarchy.

B. Hierarchical State Routing (HSR)

HSR extends the capabilities of traditional link-state routing by introducing the concept of hierarchy into the network topology through the use of physical clusters and logical subnetworks [15], [16]. Initially, the nodes that comprise the network organize themselves into a collection of clusters. Generally, these subgroups are formed based on the spatial locality of nodes (*e.g.* a collection of neighboring vehicles). Once the clusters are formed, each elects a cluster-head to manage the group's communications and to essentially act as a representative for all of its participating members. The cluster-heads define a new level in the hierarchy and accordingly, they partition themselves into clusters and elect appropriate cluster-heads to represent them at the next higher level. This process continues iteratively until the desired hierarchy is formed. At each level, link-state routing information is shared among respective cluster members, and it is propagated to the lower levels of the hierarchy. Intuitively, HSR defines a virtual topology of nested clusters, abstracting over the underlying physical clusters. For details about message propagation, the reader is referred to [15]. Figure 1 illustrates a simple HSR deployment.

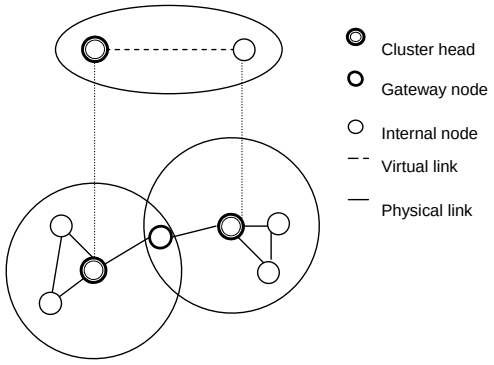

Fig. 1. Simple HSR Topology

Each cluster within the hierarchy features three types of nodes: cluster-heads, gateway nodes, and internal nodes. The role of each node is related to its proximity to selected cluster leaders. Gateway nodes are located within range of two or more communicating cluster-heads, while internal nodes only require connection to a single cluster-head. All of the nodes must be accessible from any point in the hierarchy. Given this requirement of the network topology, a new addressing scheme is adopted to utilize the hierarchical view that each node locally possesses. The original protocol defines hierarchical addresses as being the sequence of MAC layer addresses that define a node's location with respect to the hierarchy's top layer. This method is practical because MAC addresses are guaranteed to be unique and thus, the assignment of unique network-wide, hierarchical identifiers can be easily realized.

One of HSR's more appealing qualities is that it supports further abstractions over the organized clusters. Particularly, it supports logical network partitioning that can span several clusters within the hierarchy. Each of these partitions can be considered subnetworks, and each is identified by a unique network number. In a similar manner to the Internet Protocol (IP), a host in each network is assigned a unique host-id within its partition. Additionally, the introduction of subnetworks in tandem with the use of home-agents allows for source nodes to forward packets without knowledge of the precise cluster with which the destination belongs.

HSR also supports mobility by assigning each logical subnetwork at least one home-agent that is responsible for maintaining a list of hosts belonging to its subnetwork. Each node is responsible for ensuring that its home-agent has knowledge of its current hierarchical id at all times. As such, every node must know the hierarchical id of its respective home-agent. It notifies the agent with both periodic and triggered broadcasts to inform it of its current location within the network. Each node at the top layer is responsible for maintaining a mapping between all network home-agents and their respective hierarchical IDs.

With home-agent support, the process of delivering packets to mobile nodes within the network is greatly simplified through the use of indirection. When a node is sending a packet, it obtains the subnetwork number from the logical address associated with the destination node. It then sends a request to the top of the hierarchy, which will provide to the source node the hierarchical id for the home-agent associated with the destination subnetwork. Receiving this information, the source node forwards the packet accordingly to the destination's home-agent. Upon receipt, the home-agent obtains the hierarchical id for the destination and forwards the packet to that node completing the transmission.

C. Hybrid FSO/RF HSR

HSR possesses the desired qualities of a routing protocol that can be successfully applied to the FSO/RF MANET architecture. It is envisioned that the FSO/RF MANET will be deployed in scenarios that require the immediate assembly of a high-throughput network that joins a collection of groups while still supporting real-time communications. HSR satisfies this requirement by ensuring the dissemination of time-sensitive data through its underlying use of link-state routing which can propagate QoS information along with link-state status. Additionally, each cluster can employ any of an array of MAC layer protocols that further facilitate QoS. It is believed that many of the nodes in the network will feature group-dynamics. With its use of logical subnetworks and their associated home-agents, HSR is well suited to provide scalability and support for such group-based networks [15].

In order to properly incorporate HSR within the FSO/RF MANET architecture, a simplifying assumption about the network topology is made. Particularly, the network architecture assumes symmetrical physical FSO links. This restriction is imposed in order to ensure that data transmissions between the hybrid nodes can be made reliable through the use of direct acknowledgments. Given an asymmetric FSO link composition, the use of indirect acknowledgments may

provide a viable option; however, we have not fully explored how this functionality can be properly integrated. HSR can be successfully applied to the FSO/RF MANET architecture through the incorporation of the following cluster-head election policies with respect to each spatially disjoint ad-hoc network.

1) If a hybrid node is a member of a cluster (at any level of the hierarchy), it is elected the head of that cluster.
2) If a hybrid node is not a member of a cluster, then the election of cluster-head can proceed using whatever policy the chosen clustering algorithm employs.

Once a hierarchy is defined around each hybrid node, subnetwork reachability information is then shared among all of the hybrid peers, giving each a view of the network. Both periodic and triggered updates can be used to ensure that a proper view is maintained by each of the hybrid nodes. Additionally, this view can be flooded to the lower levels of the hierarchy. When an arbitrary node wishes to send information to a destination within one of the disjoint subnetworks, it will forward the packet to its respective hybrid node. The hybrid node will then check which of its peers handles the destination subnetwork, and it will proceed to forward the packet accordingly. In a configuration consisting of many hybrid nodes, with some being separated by multiple hops, routing can be accomplished by using HSR at this level as well by properly clustering hybrid nodes and using appropriate cluster-head election polices.

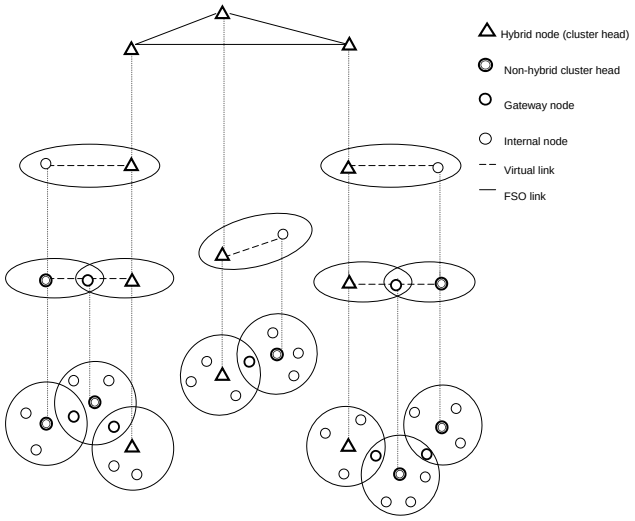

Fig. 2. FSO/RF MANET HSR Deployment.

Figure 2 illustrates the distributed application of HSR to a symmetrically linked FSO/RF MANET architecture. The architecture features three physically disjoint radio-based networks assumed to be separated over a sufficiently long distance. To avoid unnecessary complexity, RF links have been omitted from the diagram at the lowest level and the FSO links are only specified at the highest level. This diagram illustrates a simple configuration where all hybrid nodes are directly connected.

V. Experimental Results

Our experiments employed two hybrid nodes (Balrog and Gimli) based upon the pioneer P3-AT platform. Each was equipped with a Laserwire 100 Mbps FSO transceiver mounted on a pan/tilt head, 802.11g RF interface, and a Sony EVI-D70 18X PTZ system that allows for the automatic control of camera azimuth, elevation, and focal length settings (Figure 6). Due to local weather constraints, experiments were limited to indoor trials at this time.

Fig. 3. Hybrid robot nodes employed in the experiments. Each is equipped with a 100 Mbps pan/tilt FSO transceiver, 802.11g RF link, and 18X PTZ vision system.

A. Vision-based Alignment Validation

In an attempt to characterize the performance of our vision approach under range and orientation errors, two sets of trials were conducted. In the first, the robots were initially placed at a distance of 40 meters and the camera focal length was adjusted appropriately. The target robot was set at an orientation approximately equal to one of the templates used by the tracking robot. The camera then scanned over a 30° arc composed of 15 images, and the tracking robot attempted to detect its link partner within the scan. The distance between the robots was then decreased *without adjusting the focal length*, and detection was again attempted. The procedure iterated until the tracking robot consistently failed to detect its link partner. The entire process was repeated a minimum of 5 times for each test range. The test procedure for the second set of trials was identical to the first, except the target robot was placed at an orientation roughly equispaced between two of the tracking robot's template images. This corresponds to an orientation error of approximately 11°.

A total of 80 detection runs were conducted. In each, 16 template images were used by the vision system for detecting the link partner. These were taken from roughly equispaced robot orientations (every $\pi/8$ radians) on the plane. Although only 3 templates were actually required given the transceiver pan/tilt constraints, all 16 were used for each trial to better approximate the computational resource requirements for outdoor operations. Template size was typically 128x64 pixels, and a 3-level Gaussian pyramid representation was used. Under this configuration, a single 640x480 pixel image could be searched

with all 16 templates in 180ms using a 1.7 GHz Centrino notebook computer. This included the time for image pyramid generation. A SIMD implementation of the NID would further reduce these times. Ambient illumination of the target robot varied from 20-200 lumens.

The performance of the robot detection algorithm was very consistent. It correctly identified the target robot *every time* for each orientation and range errors up to 25%. However, at range errors of 30%, it always failed to detect the robot in the image sequence. To provide a little more insight into detection robustness, we define a signal-to-noise ratio (SNR) metric as

$$SNR = \frac{\min[NID(\mathcal{I} \setminus robot)]}{NID(robot)} \tag{6}$$

where $\min[NID(\mathcal{I} \setminus robot)]$ was the minimum normalized intensity distribution value obtained across the image scan that was not associated with the target robot, and $NID(robot)$ was the target robot's NID value. This is an appropriate definition, as anything else in the image sequence that is not a robot can be categorized as background noise. A plot of the SNR values vs. range for both sets of trials is at Figure 4.

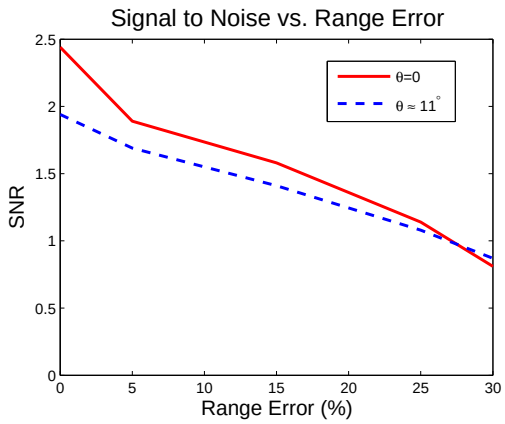

Fig. 4. SNR versus range error for two target robot orientations. As expected, the SNR drops as the pose error increases. With the exception of when range errors reached 30% (where all trials failed), the target robot was successfully detected every time.

As expected, the SNR value decays as range and orientation errors are introduced. However, the performance is still quite good at range errors up to 15%. These initial results are promising, and indicate that the approach will exhibit certain levels of robustness to robot localization errors.

B. Network Deployment

Our deployment experiments of an FSO/RF MANET included two hybrid and two RF nodes (notebook computers). Once physically linked, the two hybrid nodes defined a symmetric, mono-link FSO backbone that spanned the length of a corridor in Packard Laboratory ($\approx$ 45 meters in length) as shown in Figure 5. Each RF node was associated with a single hybrid node in the context of an 802.11g ad-hoc network. Logically, this configuration consisted of two spatially disjoint subnetworks connected via a high-throughput channel. As such, each subnetwork was associated with its own IP address domain. The RF network formed by Balrog's RF interface and Notebook 2 was assigned a subnetwork number of 169.254.0.0 with a network mask of 255.255.128.0. Similarly, the RF network defined by Gimli and Notebook 1 was assigned a subnetwork number 169.254.128.0 with a mask of 255.255.128.0. The interfaces associated with the FSO backbone were assigned an identifier in the subnetwork address space 192.168.0.0, which had an associated mask of 255.255.255.0.

The hybrid nodes were assigned the responsibility of routing packets between the RF address spaces. In order to endow them with this functionality, each was configured as an IP router. Each node was outfitted with a manually configured routing table in order to expedite the process of deploying the network.

A total of 5 deployments were conducted. In each the hybrid nodes relied upon the hierarchical LAS, and were successful in establishing the hybrid network and routing remote video data in real-time to the RF nodes. The mean optical link acquisition time for these trials was 65 seconds, with min and max values of 60 and 73, respectively. Acquisition time was dominated by 2 components: the final FSO alignment stage, and camera motion constraints. In our trials, the internal FSO alignment system can require 30 seconds or longer to establish the optical link - even when the transceivers are well aligned. Furthermore, at each pan angle we required the camera system to remain static for 2 seconds prior to capturing an image. This was a consequence of the large changes in scene depth, which would initiate the camera's auto-focus feature. This delay ensured the auto-focus system had stabilized. For a 15 image sequence, this added 30 seconds of acquisition time. The image processing portion was by far the smallest component, amounting to only several seconds of link acquisition time.

To put these times into perspective, consider instead only relying upon the internal FSO alignment system. In this case, link acquisition would be possible for only 1 of 225 possible transceiver orientations, for a worst case link acquisition time of 112.5 minutes (using a 30 second timeout for failure). This two orders-of-magnitude difference demonstrates the efficacy of our hierarchical link acquisition paradigm. Images from a sample network deployment are shown in Figure 6.

VI. Discussion and Conclusion

In this paper, we offer link acquisition and routing protocols to realize the deployment of a hybrid FSO/RF MANET. Our primary motivation is the tremendous throughput provided by FSO over long link distances. However, the technology also offers additional advantages such as lower per-bit power consumption and a secure/robust transmission medium.

To demonstrate the feasibility of the link acquisition approach, a series of experiments was conducted whereby a FSO/RF MANET was deployed and FSO links established dynamically using a hierarchical vision/FSO based link acquisition system (LAS). To date, our LAS experiments have been limited to distances of ≈ 45 meters. From extrapolating these results, we believe that extending this to 150-200 meters is

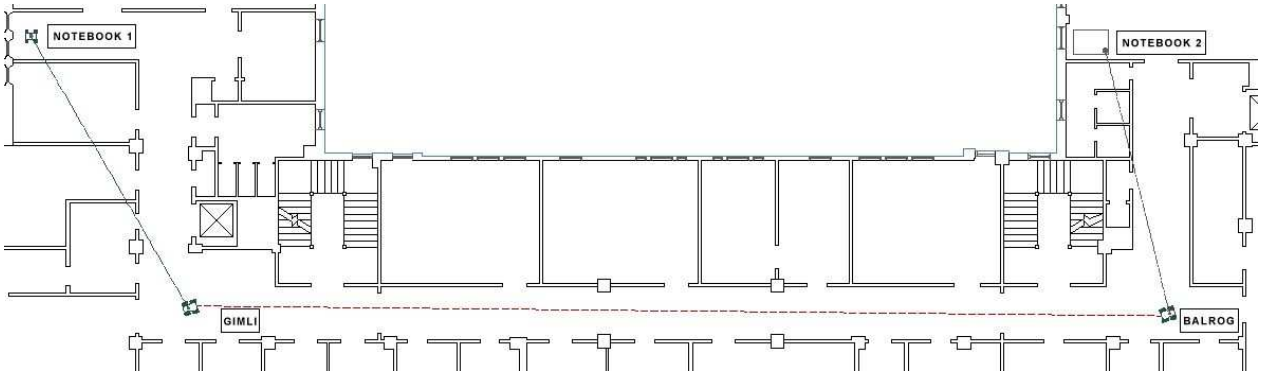

Fig. 5. Deployment configuration for the hybrid FSO/RF network.

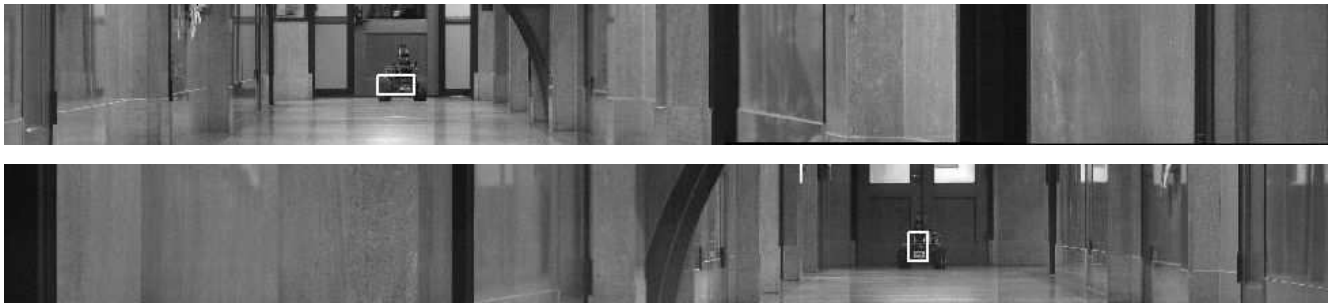

Fig. 6. Partial mosaic from the images acquired during a link acquisition trial. Using an 18X optical zoom, each robot was able to accurately localize its link partner at a range of ≈ 45 meters to enable fast FSO link acquisition.

within the capabilities of our current camera system. Extending this further would require an increase in image resolution.

Although contrived, these experiments demonstrated the feasibility and utility of a hierarchical hybrid network scheme. The next logical step is to implement a simulation featuring the single deployment of agents utilizing our routing scheme in the context of a large-scale hybrid FSO/RF MANET. Using empirical data as input, insight could be obtained in regards to the affects certain parameters have on hybrid network performance. For instance, the effects of link acquisition time on backbone routing convergence could be explored. Also, the affects of link failure and link reacquisition on per-node throughput could be gauged.

ACKNOWLEDGMENTS

Special thanks to Liang Cheng for providing valuable insight in response to our questions regarding HSR. This project was financed in part by a grant from the Commonwealth of Pennsylvania, Department of Community and Economic Development.

REFERENCES

[1] J. Derenick, C. Thorne, and J. Spletzer, *Accepted to Multi-Robot Systems*, chapter Hybrid Free-space Optics/Radio Frequency (FSO/RF) Networks for Mobile Robot Teams, Kluwer, 2005.

[2] H. Willebrand and B. Ghuman, *Free Space Optics: Enabling Optical Connectivity in Today's Networks*, Sams Publishing, 2002.

[3] System Support Solutions, *Free Space Optics*, http://www.systemsupportsolutions.com/.

[4] J. Wilson, "Ultra-wideband / a disruptive rf technology?," Tech. Rep., Intel Research and Development, 2002.

[5] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, 2002.

[6] M. Mataricand G. Sukhatme A. Howard, *Distributed Autonomous Robotic Systems*, chapter Mobile Sensor Network Deployment using Potential Fields: A Distributed, Scalable Solution to the Area Coverage Problem, Springer, 2002.

[7] S. Poduri and G. Sukhatme, "Constrained coverage for mobile sensor networks," in *IEEE ICRA*, Mar 2004.

[8] K. Langendoen and N. Reijers, *Computer Networks*, chapter Distributed Localization in Wirelees Sensor Networks: A quantitative Comparison, Elsevier, 2003.

[9] P. Corke, R. Peterson, and D. Rus, "Networked robots: Flying robot navigation with a sensor net," in *International Symposium on Robotics Research*, 2003.

[10] Q. Li, M. deRosa, and D. Rus, "Distributed algorithms for guiding navigation across a sensor net," in *Mobicom*, 2003.

[11] P. Gupta and P.R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, 2000.

[12] A. Fusiello, E. Trucco, T. Tommasini, and V. Roberto, "Improving feature tracking with robust statistics," vol. 2, pp. 312–320, 1999.

[13] Ch. Grl, T. Ziner, and H. Niemann, *Pattern Recognition*, chapter Illumination insensitive template matching with hyperplanes, pp. 273–280, Springer, 2003.

[14] E. Adelson, C. Anderson, J. Bergen, P. Burt, and J. Ogden, "Pyramid methods in image processing," vol. 29-6, 1984.

[15] A. Iwata, C. Chiang, G. Pei, M. Gerla, and T. Chen, "Scalable routing strategies for ad-hoc wireless networks," *IEEE Journal of Select. Areas Communication*, vol. 17, no. 8, pp. 1369–1379, Aug 1999.

[16] X. Hong, K. Xu, and M. Gerla, "Scalable routing protocols for mobile ad hoc networks," *IEEE Network Magazine*, vol. 16, no. 4, 2002.