Demo: Routing Overlay for Reliable Communication in Networks with Blockage *

William C. Barto, Andrea Brennen, Laura Herrera, Nathaniel M. Jones, Scott Pudlewski, Brooke Shrader, and Andrew P. Worthen
Lincoln Laboratory, Massachusetts Institute of Technology
244 Wood Street, Lexington, Massachusetts, USA

ABSTRACT

This work addresses the challenge in providing reliable communication in mobile networks with intermittent, on/off links through the use of a routing overlay designed to deal with these channel impairments. Link blockage is a predominant feature of mobile networks operating at 10+ GHz frequencies, and current techniques are ill-suited to address this problem. We present an approach comprised of multiple-path routing with end-to-end coding, queue-length-based congestion control, and a negative acknowledgement (NACK) loss-recovery scheme; these are implemented in an IP-overlay. The demonstrated scenario consists of two clusters of mobile ground vehicles operating in an "urban canyon" environment. Blockage occurs on satellite links connecting the clusters. Through the use of interactive displays, demo participants gain an understanding of routing behavior.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design - Wireless Communications

Keywords

blockage channel; multipath routing; overlay; queue-based congestion control; reliability

1. PROBLEM AND APPROACH

Mobile wireless networks can be subjected to channel blockage, where obstructions in the environment lead to on/off link behavior. Blockage is predominant in systems that operate at frequencies of tens of GHz, including mobile satellite communication (satcom) systems, as presented in this demonstration, as well as projected future commercial wireless networks, as shown by channel measurements in unlicensed bands [2, 3].

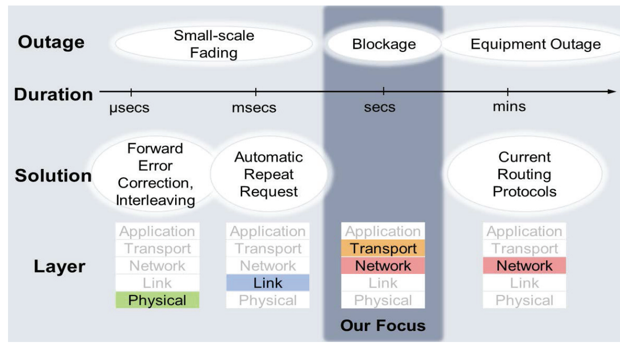

Figure 1: Comparison of durations and techniques to combat outages.

Especially challenging are recurring link blockages that last on the order of seconds. As shown in Fig. 1, solutions to combat small-scale fading, which causes outages lasting micro- or milli-seconds, are available for use in current systems, but would impose large delays when applied to longer outages. On the other hand, outages lasting minutes or longer are handled by current routing protocols. This includes routing protocols developed for the delay/disruption-tolerant networking paradigm, where a communication link between a pair of nodes only exists occasionally, as governed by the node pairs' inter-contact time. By contrast, in our scenario, communication links between pairs of nodes are persistent, but the link is sometimes blocked. Our focus is on link blockages that span seconds, as determined by the speed of the mobile and size of the obstruction. This problem is not addressed by current techniques.

We argue that the best way to combat blockages at these timescales is at and above the network layer. We demonstrate an approach consisting of three components.

- Concurrent routing over multiple paths, which provides robustness to blockages on any single link or path, coupled with end-to-end block coding, which improves efficiency by enabling the transmission of non-redundant data on different paths;
- queue-length-based congestion control and a loss recovery scheme that uses NACKs indicating the deficiency in decoding particular blocks; and

*This work is sponsored by the Office of the Assistant Secretary of Defense for Research and Engineering under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.
Contact: brooke.shrader@ll.mit.edu

MobiHoc'14, August 11–14, 2014, Philadelphia, PA, USA.
ACM 978-1-4503-2620-9/14/08.
http://dx.doi.org/10.1145/2632951.2636056.

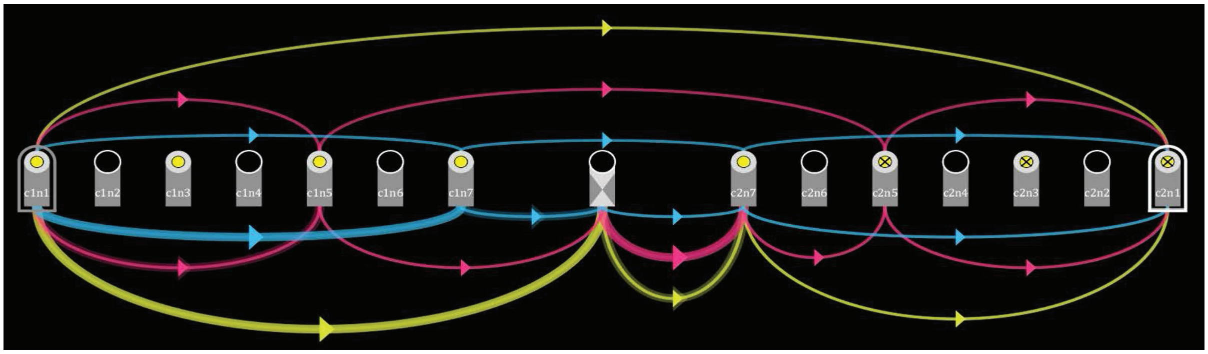

Figure 3: Snapshot of two clusters of mobile nodes connected by satellite (center node). Filled circles indicate satcom-enabled nodes, and an $\times$ indicates blockage. Lines above nodes show logical overlay routes, where each hop of the route initiates and terminates at an overlay node. Lines below nodes show the corresponding underlay routes, as chosen by the underlay routing protocol; line thickness indicates relative traffic intensity.

Figure 2: Overhead view of Cambridge, Mass., with satellite link conditions overlaid on roads [1]. Yellow: unblocked. Red: blocked.

- an IP-overlay implementation, which provides compatibility with current systems and the option of deployment to a subset of network nodes.

We refer to this collection of techniques as the "routing overlay" approach.

2. DEMONSTRATION

We demonstrate the problem and the routing overlay approach in a scenario with urban satcom blockage. The network consists of two clusters of ground mobile vehicles; clusters are connected by satellite. Within each cluster, nodes communicate over a fully-connected wireless network with error-free, fixed-rate links. Four of the seven nodes in each cluster are satcom-enabled and their satcom links can experience blockage. The blockage realizations use data from field measurements in Massachusetts for EHF satcom [1].

This scenario is setup and executed in a network simulator, and the routing overlay approach operates as follows.

- There are three static overlay paths. Random linear coding over $\mathbb{GF}(256)$ with blocks of twelve packets (1400 bytes/packet) is performed as described in [4]. Initially four coded packets are sent on each path.

- The congestion controller operates by tracking and constraining the largest queue among all overlay nodes on all paths. A delay-based algorithm is used to estimate capacity on overlay edges, i.e., between pairs of overlay nodes. The error recovery scheme sends additional coded packets, as needed, equally among all paths.

- All of the above is implemented in a standalone engine that interfaces with IP and UDP. Only satcom-capable mobile nodes are overlay nodes. A single-path routing protocol (OSPFv3) runs in the underlay network to provide routes between overlay nodes.

The demonstration provides multiple displays, synchronized in time, and showing the environment and satellite link conditions (e.g., Fig. 2), the operation of a baseline set of routing and transport protocols, the operation of the routing overlay approach (e.g., Fig. 3), and file transfer application performance. It is interactive: participants can enable/disable views of different traffic flows and can pause, fast-forward and rewind to watch blockage state changes and the corresponding behavior of routing. This demonstration provides a means of understanding and visualizing the operation of routing protocols, which is especially challenging in large networks and networks with multiple routing layers.

3. REFERENCES

[1] W. M. Smith, "Channel Characterization for EHF Satellite Communications On the Move," *MIT Lincoln Laboratory Technical Report,* TR-1109, 12 July 2006.

[2] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter Wave Wireless Networks: Potentials and Challenges," *Proceedings of the IEEE,* March 2014.

[3] J. G. Andrews, S. Buzzi, W. Choi, S. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, "What Will 5G Be?" *IEEE Journal on Selected Areas in Communications,* September 2014, to appear.

[4] B. Shrader, A. Babikyan, N. M. Jones, T. H. Shake, and A. P. Worthen, "Rate Control for Network-Coded Multipath Relaying with Time-varying Connectivity," *IEEE Journal on Selected Areas in Communications,* May 2011.