Alibi Routing – Public Review

Xiaowei Yang
Duke University
xwy@cs.duke.edu

It's a well-known problem that users have little control over the paths their packets take. There exist mechanisms that provide clues to where a user's packets traveled, but this work asks a new, hard, and intellectually intriguing question: how to provide an undeniable proof that a packet does not traverse certain user-defined "forbidden" regions.

This problem, coined as "provable avoidance routing" by the authors, has several useful applications. One is to avoid a geographic censor, such as China's firewall. A censor may discard or log user's packets or inject fake information into users' conversations. There exist mechanisms to detect or prevent such attacks, but avoiding the censor altogether is another orthogonal solution. Another application is to avoid the man-in-the-middle attack during a Diffie-Hellman key exchange. Two users can use provable avoidance routing to bypass untrusted geographic regions where untrusted men-in-the-middle may reside.

Conventional wisdom often regards a negative proof, i.e., proving that something does not happen, as intractable. This work provides a surprisingly simple and clever solution to this seemingly intractable problem. The basic idea is to prove that a mutually exclusive event has happened. That is, to prove that a packet does not traverse a forbidden region, we show that the packet has traversed a certain relay node, called an alibi, on its path from a source to a destination. These two events (traversing the forbidden region and relayed by the alibi) are mutually exclusive because the alibi node is so far away from the forbidden region. The shortest possible latency for the packet to reach the destination after traversing any node in the forbidden region and the alibi (or vice versa) greatly exceeds the network latency for the packet to reach the destination via the alibi node alone. A source can estimate the shortest latency to reach the forbidden region by assuming a packet travels at the speed of light. An alibi node signs the packet it receives and returns the signature to the source node as a proof of avoidance. A source can measure the latency of an alibi's proof and the latency of a destination's response. If both latencies are considerably shorter than the speed-of-light latency had the packet traversed the forbidden region, the source can ascertain that the packet has avoided the forbidden region.

A source locates an alibi using Alibi routing, a peer-to-peer overlay routing protocol. Each peer maintains a set of geographic diverse neighbors. A source sends a query that specifies a user-defined forbidden region, a target region where an alibi may reside, and a destination. Each peer forwards the query to a neighbor that is provably not in the forbidden region, until the query reaches a peer in the target region. A peer uses the same speed-of-light latency test described above to select a neighbor that is provably outside the forbidden region.

The idea of provable avoidance routing is fresh and intriguing. As the first work to address this problem, this paper leaves much room for discussion and future exploration. First, the Alibi routing protocol requires that all participating peers be trusted. Without a public key infrastructure, it is not clear how a node discovers those trusted peers and validates their signatures, unless we assume that all nodes outside a forbidden region are trusted. Future work might be able to clarify whether this assumption is necessary.

Second, it is not clear what additional value a "proof" of avoidance brings. The Alibi routing protocol itself is able to use hop-by-hop avoidance routing to forward a query to a potential alibi without traversing a forbidden region. If we modify the protocol such that a source sends a packet instead of a query using the Alibi routing protocol, and the last hop alibi node uses the same hop-by-hop avoidance routing mechanism to forward the packet to a destination, then we achieve avoidance routing without an explicit proof. Note that Table 4 shows that this "avoidance routing" protocol is likely to have low overhead, as on average less than two nodes are contacted before an alibi node is found. What do we lose or gain by providing avoidance routing without an explicit proof when all nodes outside a forbidden region are trusted? It is a thought provoking question and worth further exploration.