On Geolocation

Phil Gossett, Google Inc.

There's been a considerable amount of confusion concerning Google's position regarding geolocation (and beacons) for the TV white spaces. In particular, there have been concerns that this would somehow thwart mesh networks. As will hopefully become clear, nothing could be further from the truth. This white paper is an attempt to clarify how our proposal would actually work, particularly in the context of mesh networks.

Positive and Negative Beacons

First, a few words on terminology:

Beacons, as they've been typically used in the TV white space proceedings, mean something designed to protect wireless microphones by providing a "bubble" around a licensed wireless microphone receiver. Such a beacon would disable white space devices from using the same TV channel that the licensed wireless microphone is using. This could be called a "negative beacon", in that it disables a white space device from using that channel.

The other (less frequently used) meaning of beacon is as a signal to enable a white space device to use a particular TV channel, by sending permission to transmit to a white space device. This could be called a "positive beacon".

Google's Geolocation Proposal

What Google is proposing for geolocation is what we call a "hybrid fixed/portable network". As explained in our December 17th, 2007 ex parte:

> The current record in this proceeding creates an artificial dichotomy between fixed/access unlicensed devices on the one hand, and personal/portable unlicensed devices on the other. The implicit supposition is that a network will consist exclusively of nodes of either one type or the other. However, this assumption overlooks the fact that a “hybrid” network topology that combines elements of both models possesses significant value and practicality. In particular, fixed/access unlicensed devices can serve as base stations (or access points), while personal/portable unlicensed devices can serve as client nodes, perhaps for laptop computers or wireless in-home local area networks (LANs).
>
> In this hybrid scenario, a personal/portable device would not be allowed to transmit until it first had received a signal from a fixed/access device indicating which channels were safe to use without causing harmful interference to licensed services. As a result, this network configuration essentially marries the best aspects of both worlds. In particular, hybrid networks reduce the risk of personal/portable unlicensed devices interfering with licensed services (as would a network consisting only of fixed/access unlicensed devices), while providing the tangible benefits of mobility, low cost, and ease of installation and use (as would a network consisting only of portable unlicensed devices).

Additional information on this proposal can be found in our March 21, 2008 ex parte.

It should be obvious that a device that has both some form of geolocation, and access (via the internet) to a geolocation database, can use that information to protect known TV stations from interference by white space devices, preventing interference within the service area of those licensed TV stations. But that alone would only provide protection to licensed TV stations for white space devices that already have an internet connection (and know their location via GPS, or some other form of geolocation). Such a restriction would greatly (and unnecessarily) inhibit the usefulness of the TV white spaces.

What we're proposing is that this geolocation-based protection be extended to all devices which (directly or indirectly) are connected to at least one node with access to the geolocation database. This would allow mesh networks to provide the same level of protection to licensed TV stations as fixed networks. The permission for a white space device to use a TV channel would be passed on from the geo-located node to the other nodes in the mesh network. This can be viewed (using the above terminology) as a "positive beacon".

It may appear that each device in a mesh must be strictly geolocated for this scheme to work. Requiring each device in the mesh to have GPS (or some equivalent technology) would be overly restrictive. Happily, this apparent requirement isn't actually necessary. All we need to know is the maximum distance a mobile node in the mesh is from the known geolocated node in the mesh. We can then calculate the conservative set of permissions for a node anywhere within a circle centered on the known node, with a radius of the maximum distance the mobile node could be from that known node.

Note that the problem of calculating permissions is considerably easier than the more general problem of calculating the actual location of a mobile node. We only need an upper bound on its distance from the known geolocated node, not the actual location of the mobile node.

Network Layers

We need to be clear that there are two different things that need to be communicated. On the one hand, we need all nodes to be able to gain access to the internet-based geo-database. On the other hand, we need a mechanism to estimate the location of each device conservatively enough to ensure protection for licensed uses. Both can be easily achieved, but not at the same layer of the network protocol stack.

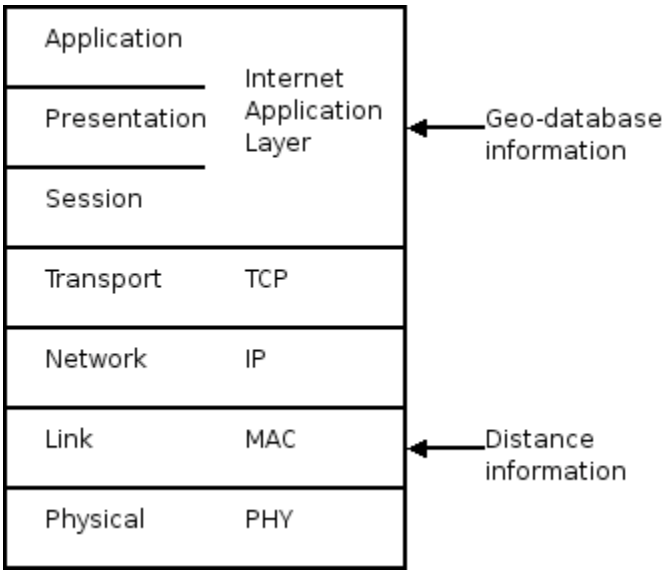

The geo-database will provide permissions as a function of distance from the known, geolocated node. This will be combined with distance information to calculate the permissions for each mobile node. Access to the geo-database is best left to a layers above TCP/IP. This allows maximum flexibility (bridging, routing and tunneling). Since it doesn't matter at all what path this information takes, there's no reason to do anything else.

For the distance calculation, things are a bit more complicated, though by no means intractable, or even particularly difficult. First, for each link from the known geolocated node to the mobile node, we need to know a reasonable upper bound on the distance of the link. This can be calculated based on maximum link range, received vs transmitted power ratios, or round-trip speed-of-light. All of these can be calculated conservatively. Assuming known antenna gains, there's a maximum link range that can be assumed for any given RF technology. For the power ratio case, reasonable upper bounds can be placed on increases in power due to "lucky" constructive interference (though some reasonable margin would be have to be included). For speed-of-light-based distances, reflection/refraction can only increase the time of flight. Combinations of these can be used to refine the distance estimates. For example, you might start with maximum link range, and then once bidirectional communication has been established, switch to round-trip speed-of-light.

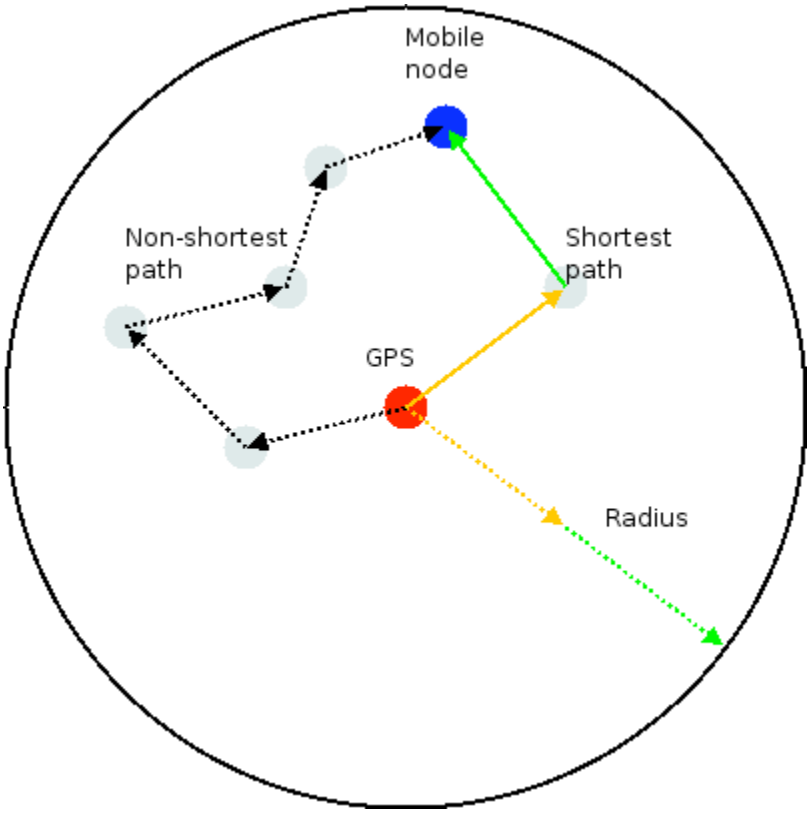

Starting from the known geolocated node, we can calculate the distances for all links between the known geolocated node and the mobile nodes. We can then calculate the upper bound for the distance between the geolocated and mobile nodes by simply summing the links in the shortest path between the nodes. If there are multiple paths between the known geolocated node and a particular mobile node, we only need to consider the shortest path. Simple geometry shows that this will always be an upper bound on the distance between those nodes. Then we can use that upper bound distance to set the radius of the circle the mobile node must be within, and calculate the permissions conservatively to be the worst case for any point within that circle.

The simple way to implement this is for each node between the known geolocated node and any mobile node to send a positive beacon packet, using the geo-database information gotten via TCP/IP, accumulating the distance of the last link into the positive beacon packet for the next link. The final mobile node will then know the upper bound of its distance from the known geolocated node, plus the permissions as a function of distance gotten from the geo-database. Each mobile node can then calculate its permissions from that.

Implications

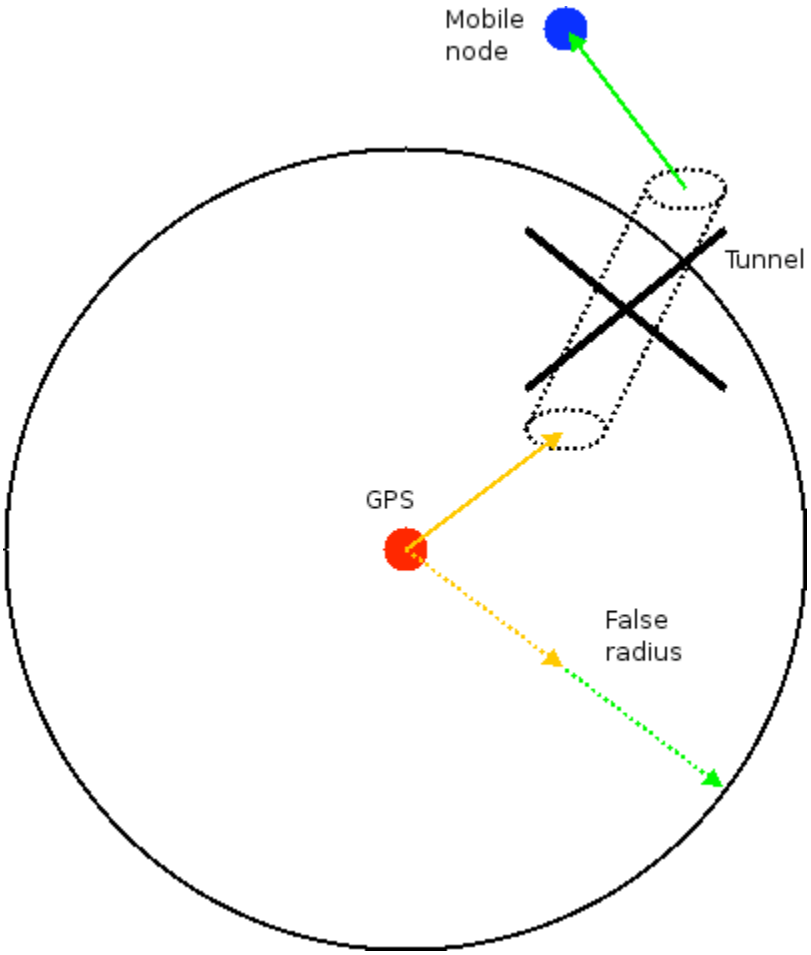

Because we can't allow tunneling to "spoof" a short distance (by forming a "wormhole" between two TV white space based sub-meshes), the distance information must be conveyed at the MAC layer. The white space devices must enforce the rule that distance information cannot be tunneled, so there's no way to "spoof" a short distance. This must be enforced by the hardware and drivers in the white space devices.

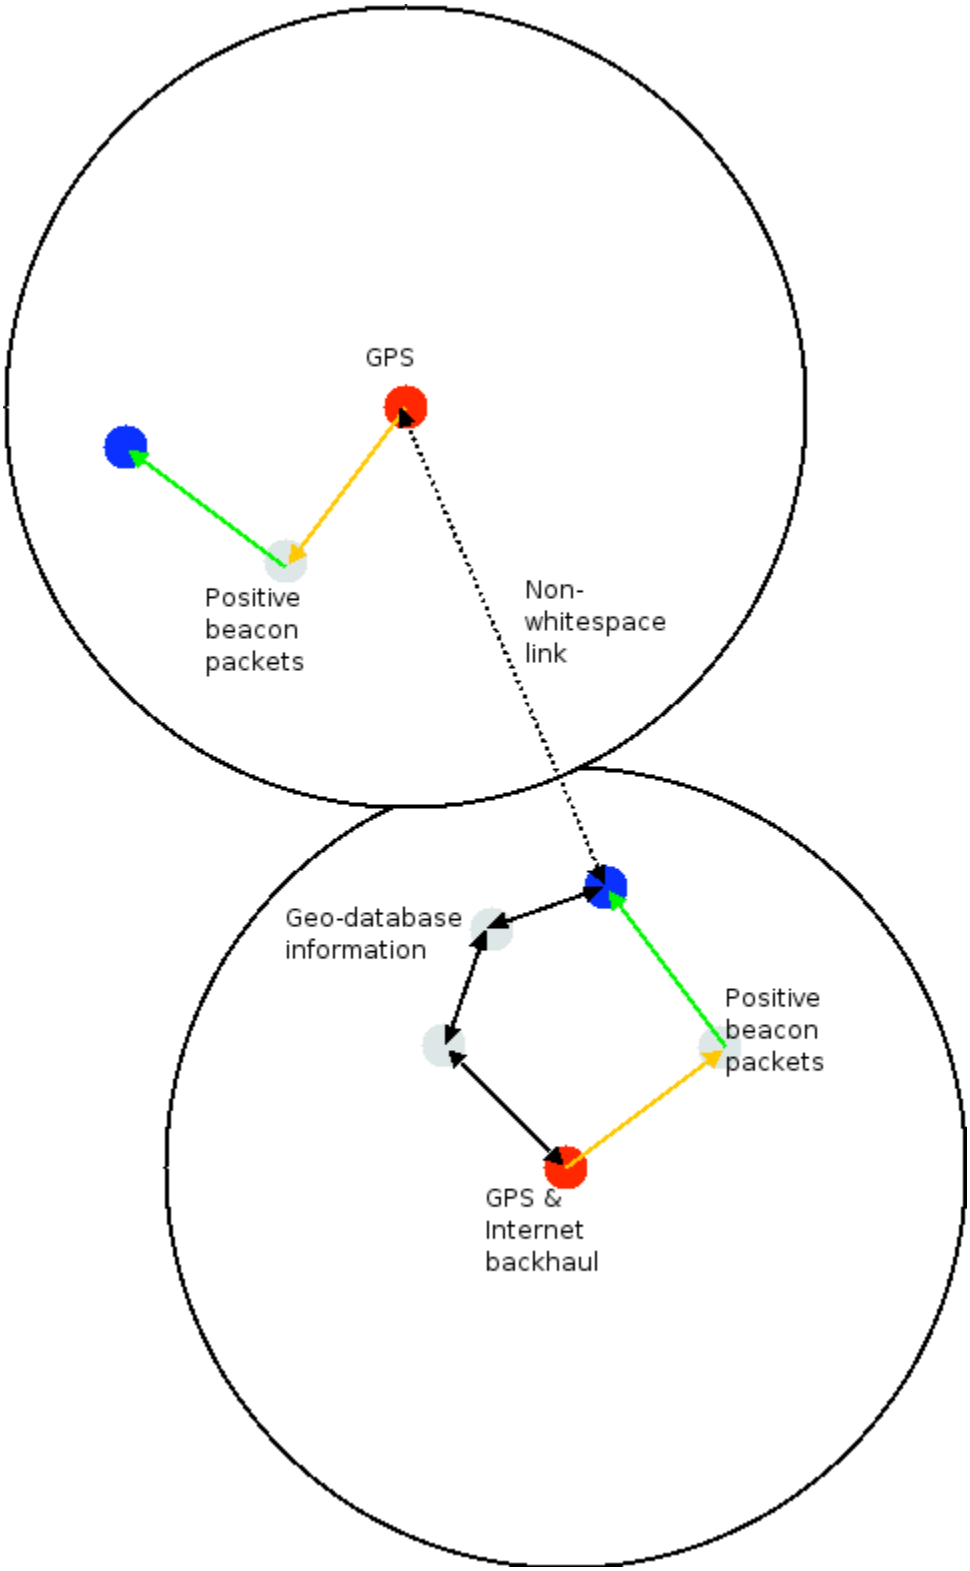

Even in a heterogeneous mesh network (using TV white space as well as other RF links), it is still the case that only one node in the overall mesh needs access (by whatever

mechanism) to the geo-database. This can be carried via TCP/IP, with no unnatural restrictions. However, at least one node in the sub-mesh formed entirely by white space devices will need to have geolocation (by GPS, or some equivalent technology), since the positive beacon packets must be at the MAC layer to prevent tunneling.

While this is bit less flexible than what is required for just the geo-database information, it still seems an entirely acceptable constraint.

Disaster Recovery

One important application of the TV white spaces is for use during natural disasters and other emergency situations to facilitate communication between members of the public safety community. To provide for this important case, special positive beacons could be provided at very low cost, under tightly regulated conditions to prevent abuse. These would send the same permissions that a geolocated white space device would send under normal situations, even if access to the interned geo-database were interrupted.

Summary

We believe that (properly implemented), our geolocation proposal will provide both the protection of a fixed geolocated device, and the flexibility of a mobile mesh device. While the technical details must be carefully attended, none of the problems are insurmountable, or even particularly challenging.