Alternate Uses for the APRS Data Stream

Using APRS mobile trackers for distributed site surveys

Darryl Smith, VK2TDS
Darryl@Radio-active.net.au

September, 2007
Sydney, Australia

Synopsis

Looking through the DCC proceedings in the last few years, there are a significant number of papers that look at how to use APRS for the benefit of the users being tracked. With the exception of PropNet looking at long distance propagation, very little has been written on how to use the combined data of the APRS network for anything other than telemetry applications such as tracking and weather reporting. This paper is designed to change that.

Introduction

In my professional life I have been working with Binary Proximity Tracking. That is, using signal presence information to determine probable locations for an unknown device based on which base stations can hear, or alternately be heard by a mobile device. The Binary term comes about because there are only two possible signal levels in such a situation – either the device can be heard, or it cannot be heard.

This has led me to think about what we other uses we can put the APRS live data stream other than for monitoring the location of vehicles. Two potential applications were devised. The first is to use the APRS data stream to plot repeater coverage maps. The second is to use the APRS data to test and develop protocols that would be used in Binary Proximity Tracking in cases, comparing the algorithmic positions with the incoming GPS locations.

Observations on the APRS data stream

There is a significant amount of information contained in the APRS data stream. Each of the mobile units that appear in its data stream identifies itself uniquely globally. The data stream also indicates which receivers have heard a particular transmission, where those receivers were locates and where the mobile unit was located at the time of the transmission.

Other information may also be available, such as antenna gains and transmitter output powers. However, as we shall see, the availability (or lack thereof) of this information does not actually affect our results.

As some APRS servers filter incoming data to remove duplicate position reports coming in via different paths, it is often better to retrieve the raw incoming data from unfiltered ports on the APRS server, or direct from the receiving station of interest. And example of this is where there are two receivers covering a particular area. On the balance of probabilities, the aggregating APRS server will indicate that 50% of packets have come from each receiver. In our case more packets mean more accurate coverage maps.

Decoding the Packet

The decoding of APRS packets is straightforward, as a publically available specification is available for APRS. The specification thankfully allows stations receiving position reports to indicate their callsign. This, in conjunction with the embedded position reports is what allows this work to proceed.

Plotting Methods

Since APRS does not use acknowledgements, we have no information about which position reports have not made it through the network. Some tracking units have a memory, and will store position reports when they are outside coverage. APRS does not work that way. Therefore we have no idea of where no coverage exists. We only know where coverage does exist. There are exceptions to this, where a RELAY, or another DIGI provide fill-in coverage, but at the moment let us assume that this information does not exist.

The lack of information actually makes plotting the coverage simpler. It may not make the plotting more accurate, but it will make it easier. I will present two methods for plotting the coverage. These two methods can be used either separately, or superimposed on one another for better results.

Scatter Graphs

Conceptually, the simplest visualization is the Scatter Graph, where a dot is placed on a map whenever a signal is heard. Since we would expect that two position reports from the same location would give the same result, we simply place a new dot at the same location. The look of the graph may be enhanced by changing the size of the dots, with smaller dots more usable where a larger number of data points are used.

Changing the size of dots depending on the number of position reports is not only not necessary, but in fact will give erroneous results, since the geographic allocation of packet transmission sites is not uniform and we have no knowledge of any transmissions going missing.

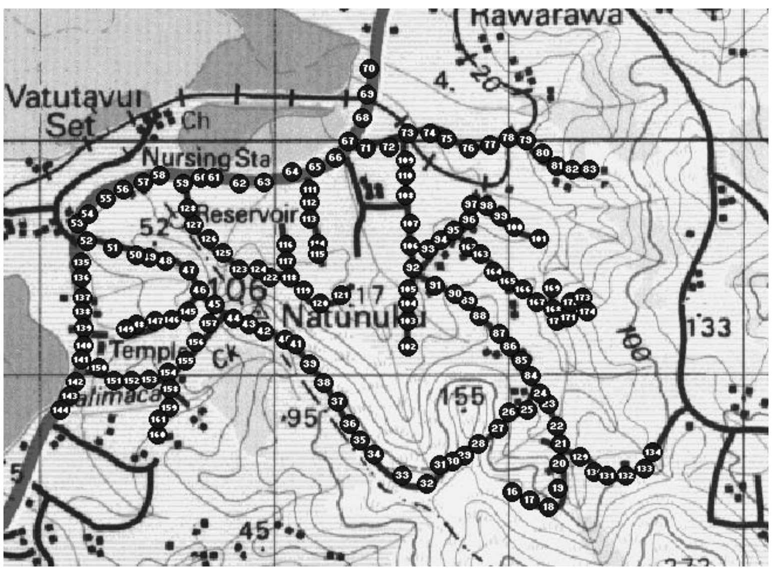

To the right is an idealized example of some radio transmission detected by a receiver in the centre of the map. It can be seen where transmissions seem to work reasonably well, and where the transmissions drop off. In this case, the transmissions quickly disappear once the transmitters get to the edge of coverage. In real life, this coverage will not be so well defined.

It should also be noted that this type of graph tends to highlight the routes taken by the drivers. It can be used to assist with the creation of public domain street maps for an area.

Radar Graphs

A Radar Graph is one where the maximum coverage of the repeater is plotted, producing a graph that looks very much like a star. In this case, reports closer to the centre will be cancelled out by reports further away. This has the effect of hiding poor coverage areas within the coverage area that might be caused by a gully or similar shadowing.

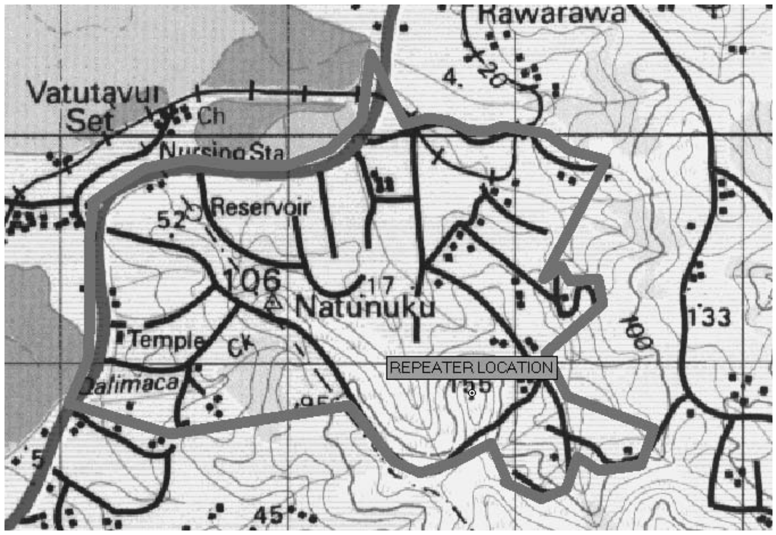

A great amount of information is lost with this type of graph. That does make it simpler to determine the maximum range, at the expense of accuracy.

Combined Graphs

The graph below shows the Radar and Scatter Graphs superimposed on each other. It clearly highlights the advantages and disadvantages of both graphs, showing the likely extents of coverage, and the known coverage areas.

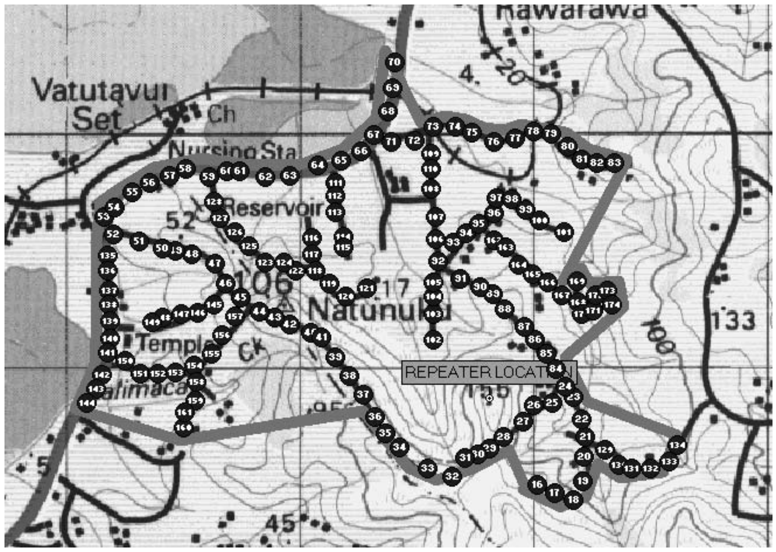

Power and Antenna Effects

Of course, this only works if we are only interested in using one vehicle to plot the coverage maps, or if all the vehicles have identical transmission powers and antenna gains. Whist transmission power and antenna gains are included in the APRS specification, very few mobile users actually send correct information out on these parameters.

On the whole APRS mobile units transmit between 25W and 50W. Very few transmit less than this power, and the number that transmits more than 50W on UHF or VHF could be considered a statistical anomaly. The mobile units that transmit a lower power will provide information on coverage for areas close to the receive site.

Likewise, most mobile units will operate with a ¼ λ antenna providing 3 dBi on-axis gain. Higher gain antennas will tend to not operate as one would numerically have expected due to mobile signal dispersion effects. Of course, all bets are off if the mobile user is operating with a high gain directional antenna such as a Yagi. Thankfully these are rare in a mobile environment for APRS tracking.

Some users will operate using antennas with a lower gain for the selected frequency due to size considerations on their vehicles. The author of this paper is an example of such a user, who operates with a ¼ λ 70cm antenna on the 2m band. This has the same effect as reducing the output power.

Ignoring summertime 'Sporadic-E' ducting events and space operations, VHF and UHF FM operations show a relatively consistent range over time. The coverage of a repeater will not vary dramatically over the course of a 24 hour period, nor will it change significantly between seasons.

Determination of Coverage for Non-Digipeater Sites

Using the collected data to determine the coverage of a digipeater site is useful for determining the coverage of an existing site, but in reality once the coverage has been plotted there is not a great need to do it again for some time. The real power comes from using the APRS data to compute coverage maps for new sites using actual radio transmission.

It would be possible to place a TNC at a proposed repeater site monitoring APRS transmissions building coverage maps. Since APRS can operate without a radio making any transmissions, it should be legal for a commercial entity to use the received APRS data to produce coverage maps.

Binary Proximity Tracking Algorithm Development

During the last year, I have been working on developing some Indoor Tracking products. That is, devices and systems designed to track people and objects as they move within a building, without access to GPS. A large amount of work has been done in this field, but there is still a lot of work to be done. Algorithms are available, but the best ones tend to be proprietary and unavailable.

Developing algorithms tends to be a catch 22 with regard to the software. How do you determine where a person or device is within a building so that you can compare this with the predicted location? This might sound easy, but we really want to be able to monitor the location once per second as they move through the area, and this is not really feasible to do accurately with human interaction. At least not for a large number of tracked people objects.

The area of Indoor Tracking I am interested in is 'Proximity Based' tracking, where the mere presence of a radio transmission indicates you are within an area. Or the lack of transmission indicates you are outside the coverage area. Since the transmission is either detected or not detected, this is commonly called Binary Proximity Tracking.

If you were to remove the GPS data from the APRS data stream, the information received is essentially identical to that from sensors in Binary Proximity Tracking. The difference is the timescale and physical sizes involved. Instead of dealing with floors and buildings, we are talking towns, cities and states. Instead of seconds, we are talking at the very least minutes to move between coverage areas. Therefore it is possible to take the APRS data stream, and use it to develop algorithms for Binary Proximity Tracking. Then the GPS location of a vehicle can be compared with the projected location for protocol verification and tuning.

Thus, the same algorithm is usable with adjustment of two tuning parameters – the size of coverage circles and the time between position reports. Adjusting these two parameters should allow realistic testing to take place. Improved results will also come from monitoring received signal reports from each internet connected APRS station in an area, even if they have chosen not to generally IGATE position reports.

Conclusion

Through this paper, I have shown some non-traditional uses for the data contained in the APRS data streams worldwide. The uses provide a benefit to all uses of the radio spectrum, highlighting that Ham Radio really can be the radio equivalent of a National Park.

Appendix 1: Range and Area Effects of Power Doubling

Doubling power from 25W to 50W does not double the range. The range is only increased by about 41%. The math relies on the surface area of a sphere, which is four times the area of a similarly sized circle.

$$A = 4\pi r^2 \therefore r = \sqrt{\frac{A}{4\pi}} \; or \; r = \sqrt{A}\sqrt{\frac{1}{4\pi}}$$

Simplifying this

$r \propto \sqrt{A}$

If A is actually the transmission power, we can see what happens to the radius increasing from 25W to 50W. In this case, the equation becomes

$$Range\,Ratio\,Increase = \frac{r_{50}}{r_{25}} = \frac{\sqrt{50}}{\sqrt{25}} = 1.41$$

$$Range\,Area\,Increase = \frac{A_{50}}{A_{25}} = \frac{\pi \cdot 1.41^2}{\pi \cdot 1.00^2} = 2.00$$