How your phone betrays your location

It might not be as accurate as GPS, but as **Julian Bucknall** explains, your mobile phone transmits your location to anyone equipped to listen

"What's an azimuth?" my wife asked. For some reason, visions of naval battles came to mind. Cannons being aimed, orders shouted, shots fired, smoke and billowing sails. Other than that, I couldn't come up with much. There was nothing for it but to check Wikipedia, where I learned that it's a direction for a projection of an object (usually a star) in 3D space onto a reference plane. So, in summary, a direction.

At this point, it's best to understand what my wife does for a living. She's a Deputy District Attorney, a prosecutor, in the Fourth Judicial District in Colorado. For a long while, she's headed up the Special Victims Unit (SVU, just like on the telly) in the office, prosecuting cases where the victims are the very young, very old, or those who need special help. I get involved as a kind of home-grown expert because, like it or not, many of these cases involve computers and the internet. I get to explain technical topics in terms that she can understand and then explain to the jury in trials.

This time she had a printout of calls made to and from the defendant's mobile phone covering the period during which the alleged crime took place. The report showed the originating phone number, the called number, the duration, and then some extra information: a code for a cell tower antenna and its azimuth for both the caller and the called. But what did the azimuth number mean?

Mobile phones work by using radio waves to communicate with cell phone towers (usually known as cell sites) and from them to the normal 'wired' public switched telephone network (PSTN). The cellular network consists of strategically placed sites arranged in a regular pattern so that the carrier can provide blanket coverage over an area or region.

Cell sites

When you switch on a mobile phone, it will search for a signal from the cellular network. I should point our here that I'm in America and our system is subtly different to the one used in the UK. Ours, like yours, uses pre-defined carrier frequency bands for communication. My iPhone is on the AT&T network, and the bands are at 850MHz and 1,900MHz. These frequencies differ across the globe. Each band is divided into channels. The phone scans the bands trying to find a channel with the strongest signal, the inference being that that signal must come from the nearest cell site.

In urban areas, cell sites cover smaller areas – say half a mile in radius. In rural areas a site could cover an area up to five miles in radius. For urban areas, the reasons for the smaller coverage are two-fold. First, there are buildings that get in ▸

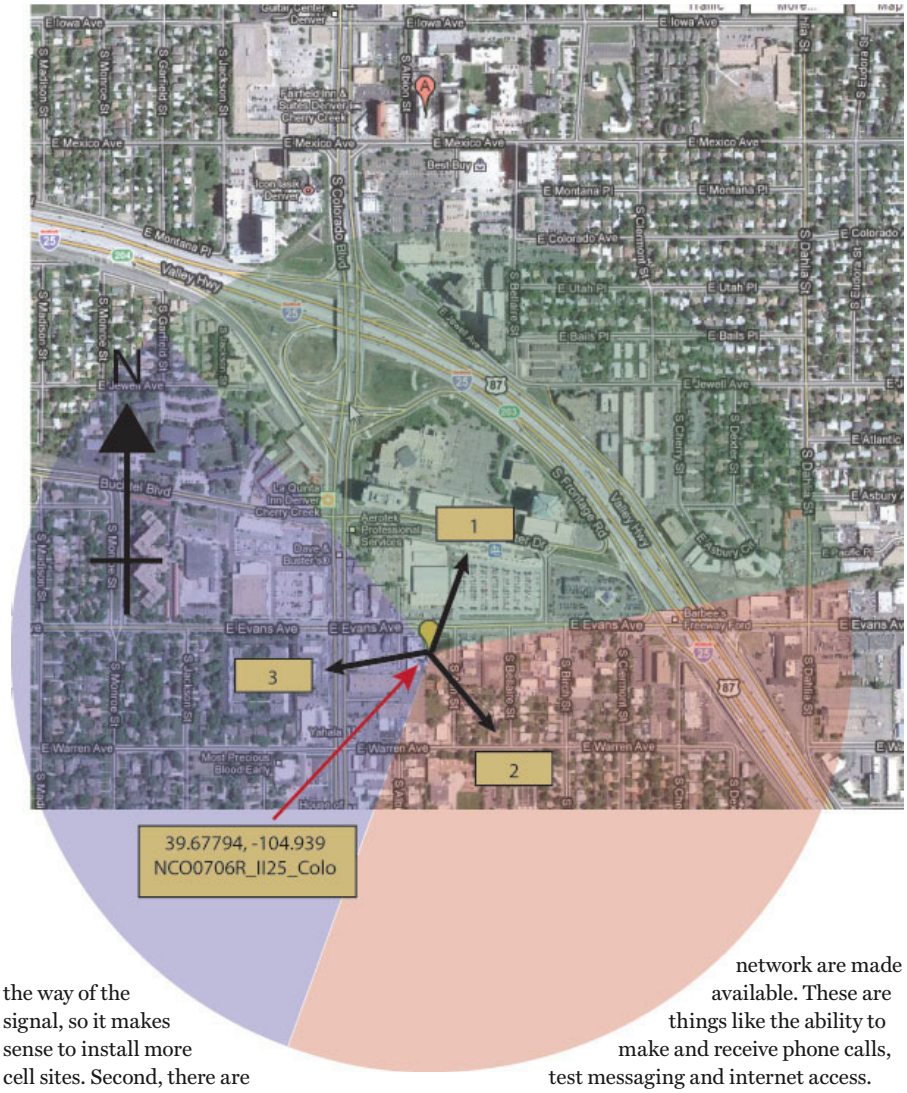

◄ Figure 2: The antenna azimuths and coverage for the given tower marked on a Google Map.

► the way of the signal, so it makes sense to install more cell sites. Second, there are more subscribers in urban areas, so in order to avoid a site's channels getting swamped with connections, a carrier will build more towers. In practice, a cell tower will be leased to several carriers, which is why a tower in an urban area fairly bristles with antennas.

Once the strongest signal has been identified, the phone will negotiate using a standard protocol to log into the cell site. During this process, the phone will transmit a couple of numbers to the cellular network for identification purposes. The first is the International Mobile Equipment Identity, or IMEI. This is a unique number, and it's big – my phone's is 15 digits. If you have a GSM phone, you can type ***#06#** on the keypad to see if that identifies the device. If your phone is stolen, you can ask for the IMEI to be blocked, rendering the phone useless on the network. The second is the International Mobile Subscriber Identity or the IMSI number. This uniquely identifies you, the subscriber, and is encoded on the SIM card. Again, it's a 15-digit number. The combination of these two fields identifies the mobile's phone number, though sometimes the IMEI number is used by itself.

After the protocol negotiation is complete, the phone's location is registered with the cellular network and the features of the network are made available. These are things like the ability to make and receive phone calls, test messaging and internet access.

Since the whole point of mobile phones is that they are mobile, the cellular network has protocols in place to make sure the phone is connected using the strongest signal possible as it's moved around. As you travel, the phone and network are cooperating and handing over your phone from cell site to cell site. As part of this handoff, your phone may switch from one channel to another. The network and phone are designed so that this switch can happen as a call is taking place. To you, it's as if the phone is connected to the network permanently and is the only device on that network.

Staying connected

That's the theory. In practice things can be rather different. The internet is full of stories of dropped calls and bad reception. Carriers craft their advertising to target a competitor's perceived bad network, but all carriers actually suffer. The more phones are within range of a cell site, the more calls are possible at any one time, which can swamp the possible channels.

The cellular network (and hence the carrier) knows where the phone is at all times when it's switched on – not necessarily being used for a phone call or for data communications. This isn't a GPS-style position to within a few tens of meters, but a location in terms of the area covered by the cell site. If you think about it, the carrier has to know where the phone is, otherwise you wouldn't be able to receive any calls or get any push notifications.

Going back to my wife's phone call list, that explains the codes for the cell site for the originating phone and the receiving phone (for a landline the cell site code was blank), but where does azimuth come in?

Azimuth explained

If you look at a picture for a cell tower, you'll notice that the antennas are in a triangular formation. The antennas are directional: each points in one direction. Since the tower uses a triangular form, each antenna provides about 120 degrees of coverage, with the three giving the full 360 degrees. As it happens, each antenna actually provides about 130 degrees of coverage. This is to ensure a little overlap so a phone doesn't 'fall off the radar' in between two antennas and get handed over to a more distant site. The direction each antenna of a cell site points in is known as the azimuth.

The report from the phone company therefore let us track where the defendant was every time he made or received a phone call on his mobile phone. My wife's case needed to place the defendant at a particular place in Denver, but the defense attorney's counter to that was that the defendant was driving along I-25 at the time. If he was too close to another site, the phone would register there. But how close were the cell sites for this carrier? How many sites were in the general area? More evidence was needed to rebut that explanation.

I told my wife that she needed to subpoena the carrier for a complete list of cell sites in the Denver area. Once we had that, we could more easily determine where the defendant's mobile phone was located. A week or so later, we had the list as an Excel spreadsheet. I was totally surprised: there are hundreds of cell sites in

On the server

Recently there have been numerous newspaper articles about location-based services, especially in regard to social networks like Facebook and Twitter, culminating in the brouhaha about the location database on Apple's iPhone.

These location-based services use a combination of hardware and software on the device itself to calculate the position of the device (and therefore the user). The hardware uses a combination of a built-in GPS system (accurate to approximately 20m) and knowledge of nearby cell sites (a database sliced and downloaded when needed to improve the accuracy). The cell sites aren't used actively (the network isn't used for location) – instead, knowledge about the position of the nearest cell tower helps improve the accuracy of the location calculation.

These services attach latitude and longitude information to tweets and posts, or can be used interactively together with maps and, say, local restaurant information to aid the user. ■

Acronyms

GSM (Global System for Mobile Communications) is a set of standards for digital cellular networks. It was the second such standard, and is therefore also known as 2G. Later it was enhanced to provide packet data through GPRS (General Packet Radio Service), and still later to the faster EDGE (Enhanced Data rate for GSM Evolution). There's now a third generation standard for digital cellular networks called UTMS (Universal Mobile Telecommunications System), or 3G. A new one, sometimes called 3.5G, adds HSDPA (High Speed Downlink Packet Access) to speed up data downloads. CDMA (Code Division Multiple Access) defines how channels are shared among users and is used in 3G networks; 2G networks use TDMA (Time Division Multiple Access). Confused yet? ■

Spotlight on… Accurate location

The main article described an investigation into phone and cell site data some months after the events happened. In this case, the only information available was what had been logged by the carrier. This tends to be call information or data usage information – just enough to bill the subscriber accurately.

In real-time, the situation is a little different. A phone broadcasts radio waves, and although the system is designed to latch the phone onto the nearest cell site for the best service, the handset does not then direct its signals to only that cell tower. There is no directional broadcast system in a phone: the radio waves propagate equally in all directions.

It's therefore possible to listen in on the signal from a mobile phone from several antennas and triangulate the position of the device. This is much more accurate than the method we used below. This triangulation can only happen in real-time of course.

This method's main drawback is that it requires the carriers to add extra software and hardware to their network infrastructure, which means it's decidedly only within the purview of governments, not individuals. ■

Denver. This was looking like a real problem. The data included the name of each cell site (the one we were interested in is NCO0706R_II25_Colo), and its latitude and longitude.

Limiting the data

I wanted to limit the data I had to a particular square – say five miles on a side. Twenty-five square miles seemed like a good start.

My first problem was calculating how many miles there are in a degree of latitude and a degree of longitude. I could then sort the cell site data to limit it to sites around the cell site in question. The Earth is about 4,000 miles in radius. Taking latitudes first (on the surface of the Earth, the distance subtended by a degree of latitude is constant), the number of miles for a degree of latitude is 4,000 x sin(1), or about 70 miles. So I needed to limit my data to +/-0.035 degrees of latitude from NCO0706R.

Longitude is harder: the lines of longitude converge at the poles, so the number of miles for a degree of longitude is at a maximum at the equator (70 miles) and at a minimum at the pole (essentially, 0 miles). In other words, the number of miles for a degree of longitude varies according to the latitude of the place in question (it varies according to the cosine of the latitude). Denver's latitude is roughly 39.6 degrees, so the number of miles for a degree of longitude is 70 x cos(39.6), or 60 miles. This meant I had to limit my data to +/- 0.042 degrees of longitude around the cell site. I extracted data whose latitude and longitude fell into these ranges and found 17 cell sites in the 25 square miles centred on this cell tower.

There was no way my wife could present an Excel spreadsheet of 17 rows of cell site data to the jury and ask them to visualise it. I decided to use Google Maps for this, but since I didn't want to enter 17 pairs of numbers to drop markers, I had to do it programmatically.

It turns out there's a complete JavaScript API to Google Maps. Using a specially crafted HTML5 page and a little bit of JavaScript, you can display a map centred on a point of interest, and then display one or more markers on the map. All I had to do was write a bit of code that converted the cell site data from Excel into a JavaScript array of objects and import that into the page's script tags. It it would have been quicker to enter all the lat/long data into the Google Maps site, but I'm a programmer at heart. Besides, I could use it for another set of data. Yes, that's the reason.

Figure 1 shows the result. The marker in the middle is our friend NCO0706R_II25_Colo. As you can see, cell sites are placed pretty close together in this part of south Denver. The next cell sites north and south are very close (about half a mile), whereas the sites to the east and west are about a mile away. In other words, the defendant was less than half a mile from the cell tower when he was making or receiving calls. Could I place him any better than that?

Back to the original data. The cell tower had three antennas, and their azimuths were quoted as 10, 130, and 250 degrees. In other words, the first antenna was pointing 10 degrees from north (north is taken to be 0, and we count clockwise). Since an antenna has a 120-degree span (we'll ignore the overlap for now), that means it picks up signals from 310 to 70 degrees. The second spans from 70 to 190 degrees and points southeast. The third spans from 190 to 310 degrees and points almost due west. Figure 2 shows the situation.

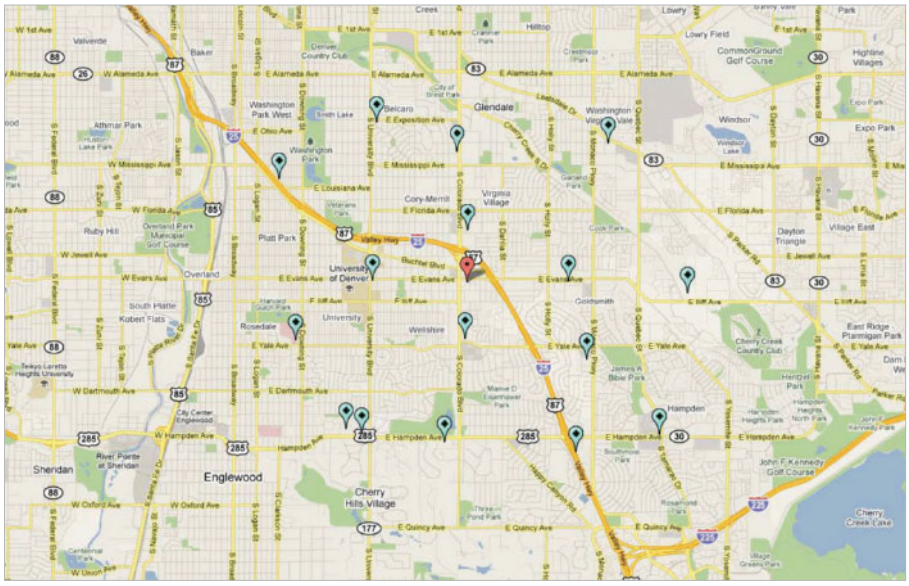

▲ **Figure 1:** The positions of the cell sites around the given tower.

Closing in

Now the fun stuff: there were three calls in quick succession. The first was picked up by antenna 2, the second a couple of minutes later by antenna 1, the third a few minutes later by antenna 2. The remainder of the calls used antenna 2. This couldn't be explained by driving up and then down the Interstate – there wasn't time to make a U-turn. You can't park on the Interstate either. What's going on?

It's simple to explain if we assume that the defendant was stationary, on the dividing line between using antenna 1 and antenna 2. Since there is an overlap, it's entirely feasible that his phone switched antennas at the nearest cell site for one call. And my wife had other corroborating evidence that he was in a building on that dividing line. Case closed. PCP

Julian M Bucknall has worked for companies ranging from TurboPower to Microsoft and is now CTO for Developer Express.
feedback@pcplus.co.uk