Brought to you by

Forensic Analytics

TIMING ADVANCE

PRECISION & RELIABILITY TESTING

April 2025

v1.7

Contents

1. Executive Summary

Timing Advance (TA) is a type of digital forensics evidence provided by cellular networks.

It is generally regarded as providing evidence of a target phone's distance from the serving cell tower when connections were made, with some digital forensic practitioners using terms such as 'pinpoint' and 'accurate' in relation to TA's ability to estimate the target phone's location.

A common defence response to the use of TA-based evidence is to state that there seems to have been no published empirical testing of the precision or reliability of TA evidence, specifically of the 'distance from tower' data provided by TA.

This document is intended to provide some initial testing of this type of evidence, to assess whether it can be regarded as precise or reliable and to set some recommendations for the interpretation of TA-based evidence.

This will be achieved by reviewing the technical basis for TA, to help readers understand how it works and what precisely TA values are telling us. It also looks at some potential complexities related to TA data, which can lead to TA values being less precise that some digital forensics practitioners may realise.

We have also undertaken some basic testing of TA – in both the UK and the US – and we present the initial results of that testing below.

In the past, some on the prosecution side have argued that TA is totally reliable and is able to accurately 'pinpoint' the location of a target phone within a sector. Others, on the defence side, have attempted to portray TA evidence as worthless, stating that it can only be said that the target phone was somewhere within the cell it used, but cannot be located any more precisely than that.

The initial results of our testing are, predictably, somewhere between these two extremes. We submit that while TA cannot be regarded as being precise enough to 'pinpoint' a phone's location, it can be used to assign an approximate location for the phone within an area that extends for some 200-250m (650-820ft or three TA bands) around the stated TA distance. Whilst this is less specific than saying 'the phone was within the stated TA band area' it is more precise than saying 'the phone was somewhere within the used cell' and offers, we think, a degree of reassurance that location evidence obtained from TA data is valid and viable and can safely be used in criminal cases, with the correct degree of caution.

1.1 Contributors

This document was written by and with the help of:

- Joe Hoy – Forensic Analytics Ltd (UK)
- Martin Griffiths – Forensic Analytics Ltd (UK)
- Cory Brodzinski - Denver District Attorney's Office (US)
- Nick Falcicchio – ABM Intel (US)

2. Timing Advance

Each cell in a cellular network is capable of serving multiple devices simultaneously by employing a variety of 'multiple access' techniques.

Whilst this is highly desirable in the sense that it allows the network to maximise the number of devices each cell can serve, it creates problems if the active devices are operating at different distances from the base station.

Radio signals travel at the speed of light – roughly 300,000km/s – whilst this is very fast, it is still a finite speed. It takes a finite amount of time for a radio signal to travel a given distance.

For example, if a device was 300,000km away from the base station (leaving aside the fact that a cell could never be that big!), it would take 1 second for a signal to travel from the device to the tower; if it was 150,000km away the signal would take 0.5s, and so on.

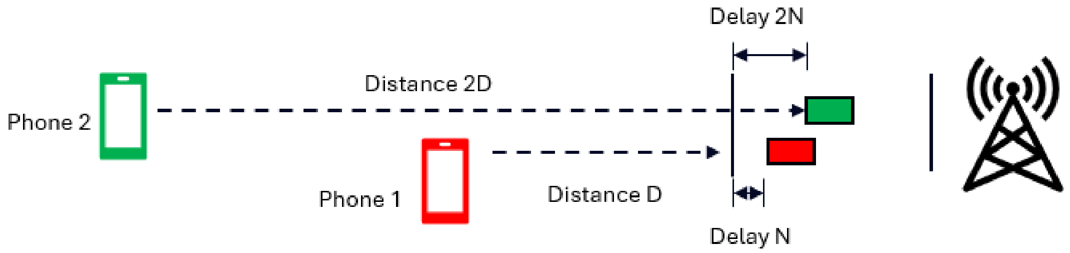

At the real distances that cellular signals are required to travel (usually up to a few km), the differences in travel times at different distances would be in the microseconds, but even time differences of that level can have unwanted effects. Most cellular multiple access techniques, used in different network generations, employ some form of 'time division', meaning that signals from different devices need to arrive at specific points in time to avoid overlapping with signals transmitted by other devices. It becomes important to be able to factor a device's distance from the tower (and the resulting delay due to the speed of light) into the decision about when each device begins to transmit its 'bursts' of uplink data.

This activity is managed using the Timing Advance process.

Timing Advance is concerned with synchronising each individual device's uplink transmissions to compensate for the delay associated with the device's distance from the serving tower so that each uplink transmission arrives at the base station within the allocated time period.

Timing Advance is required on the uplink because transmissions come from multiple sources that are located at multiple distances from the tower. Timing Advance is not required on the downlink, as all transmissions in a cell come from a single source, the base station.

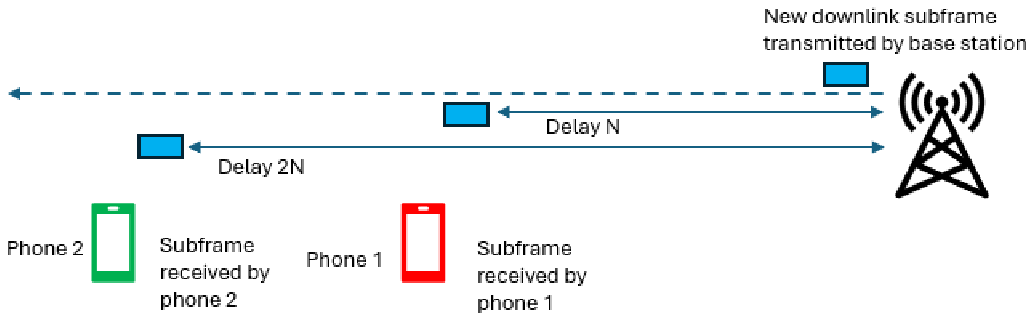
New downlink subframe transmitted by base station
Delay N
Delay 2N
Phone 2
Subframe received by phone 2
Phone 1
Subframe received by phone 1

Each device synchronises itself to the downlink signal from the base station, but this synchronisation is relative rather than absolute – devices at different distances from the tower will experience the start of a new downlink subframe at different times from each other; devices that are closer to the tower will react to the start of the subframe earlier than devices that are further away, as the downlink signal carrying that new subframe will take longer to propagate out to more distant devices.

The base station has the opportunity to evaluate each device's initial Timing Advance requirements during the random access procedure.

A device that requires a connection waits for the beginning of the downlink subframe that contains the next random access period and transmits a Random Access Channel (RACH) Request on the uplink to the base station. A random access period typically lasts for several milliseconds and the device transmits its RACH burst at what it determines (based on relative downlink synchronisation) is the start of the period. Depending on the device's distance from the tower, the point at which it determines that the RACH period is starting may be some microseconds after the base station actually began transmitting the subframe in which the RACH opportunity exists – this is because the corresponding downlink frame, transmitted by the base station, may take a short amount of time to propagate to the device's location. The device will only start to transmit its RACH burst when it sees the corresponding subframe starting, so this adds to the delay associated with the RACH process.

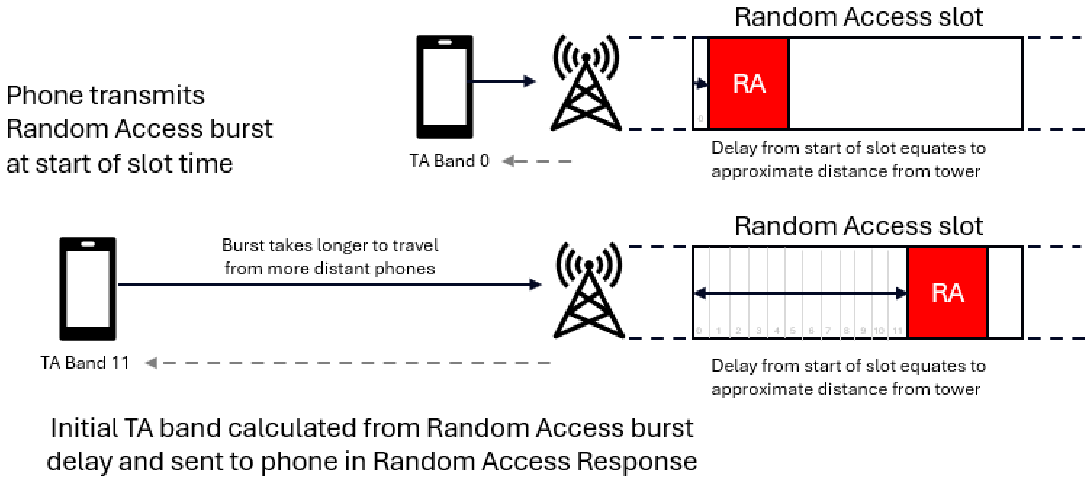
Phone transmits Random Access burst at start of slot time
Random Access slot
RA
TA Band 0
Delay from start of slot equates to approximate distance from tower
Random Access slot
Burst takes longer to travel from more distant phones
RA
TA Band 11
Delay from start of slot equates to approximate distance from tower
Initial TA band calculated from Random Access burst delay and sent to phone in Random Access Response

The RACH burst will travel at the speed of light (*c*) and will arrive at the base station at some point during the random access period, the device's distance from the tower will determine how late in the period the burst arrives.

In the diagram above, the device at the top is close to the tower, its RACH burst arrives early in the random access period and the base station determines that no Timing Advance adjustment is required.

The second device is much further away from the tower – its RACH burst therefore takes much longer to travel to the base station and arrives later in the random access window. The amount of delay suffered by the burst allows the base station to calculate the amount by which the device needs to compensate for uplink delay.

It is not necessary for the base station to capture the *exact* delay experienced by each phone's transmissions; the process doesn't need to be that granular. Instead, delay is measured against a set of fixed steps or 'bands'. For example, in 4G LTE, each step period is 0.52µs (0.52 microseconds) long. If the delay measured for a phone is between 0 and 0.52µs, it is regarded as being in TA Band 0, if the delay is between 0.52µs and 1.04µs it is regarded as being within TA Band 1 and so on.

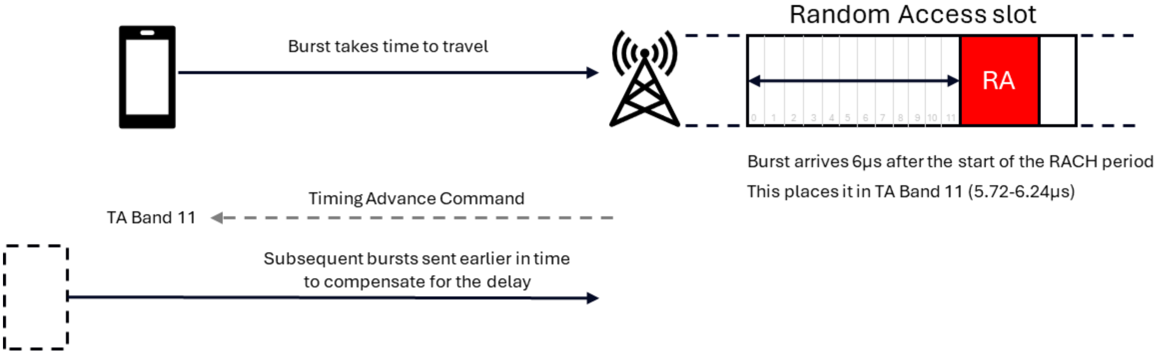
Random Access slot
Burst takes time to travel
RA
0
1
2
3
4
5
6
7
8
9
10
11
Timing Advance Command
TA Band 11
Subsequent bursts sent earlier in time
to compensate for the delay

Burst arrives 6µs after the start of the RACH period
This places it in TA Band 11 (5.72-6.24µs)

In the example above, the RACH burst from a phone arrives at the base station 6µs after the start of the Random Access period. That level of delay falls within TA Band 11 (which ranges from 5.72µs to 6.24µs).

The base station sends the device an initial Timing Advance command in the Random Access Response (RAR) message – this instructs the device to start transmitting any further bursts slightly earlier in time (earlier than the start of an allocated subframe), to compensate for the distance-related delay and to ensure that those bursts arrive within the allocated time period.

The initial Timing Advance calculation for a device is made when it first performs random access, but mobile devices are, well, mobile and can move around during the lifetime of a connection. If the device moves further away or closer to the tower, there's a chance that it will eventually slip into a different Timing Advance band and that its TA setting may need to be adjusted.

Further timing 'advance' commands are sent to the device to instruct it to transmit earlier if the delay on the uplink increases (if bursts begin to arrive too late due to moving further away from the tower) or timing 'retard' messages can be sent to instruct it to transmit later if the delay on the uplink decreases (if bursts begin to arrive too early due to moving closer to the tower).

In LTE, TA instructions are sent as frequently as several times per second or as infrequently as every few seconds, as required, dictated by changes in the measured delay. TA is only available if a phone is active – it needs to be sending data to a tower for TA to be measured.

3. Inferring Distance from TA

The 'steps' or bands signalled by TA commands indicate a unit of two-way delay, with each band in LTE equal to a delay of around 0.52µs. The full 4G TA calculation is:

$T_{TA} = 16 \times TA \times T_s$

TA is the current Timing Advance band number, which is multiplied by 16 to reduce the granularity of the process. T_s is (1/ (FFT x SCS)) and is the basic timing unit of LTE, where the FFT is the maximum theoretical number of subcarriers in a channel (2048 in LTE) and SCS is the subcarrier spacing across that channel (15000Hz in LTE).

As radio signals travel at the speed of light, it is possible to calculate the distance travelled by a radio signal in a given period of time – the distance travelled in one TA band period is 156.14m (0.52µs x 299,792,458m/s).

This is the two-way delay, however, as it includes the downlink delay (the time for the leading edge of the RACH subframe to reach the device) and the uplink delay (time for the RACH burst to travel back to the tower), so we divide it by 2 to get an approximate one-way distance, giving a TA band size of 78.07m (256ft or 0.05 miles) for LTE . The TA calculation takes into account the subcarrier spacing applied to the channel, so the values vary in 5G, as the SCS varies, with 5G band sizes of 78.07m with SCS of 15kHz, 39.035m for SCS 30kHz and 19.51m for SCS 60kHz.

In theory, we can use the delay associated with Timing Advance to infer an approximate distance for the device from the serving tower.

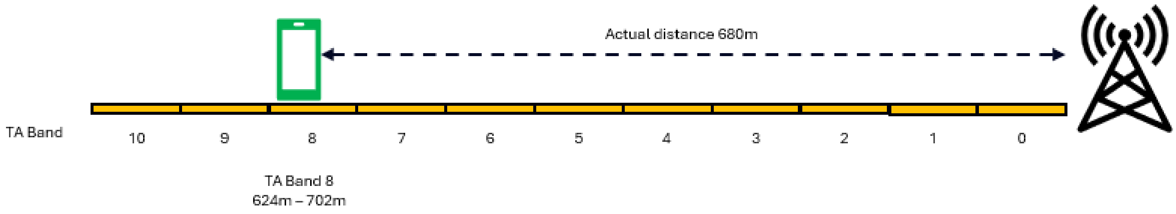

Imagine, for example, that an LTE device is 680m away from the serving tower – the actual two-way delay associated with that distance would be 4.54µs, which is equivalent to between 8 and 9 TA steps.

So, without knowing the device's actual distance from the tower, we could use TA to infer an approximate distance of between 624-702m (TA Band 8) and we could plot this on a cell site map as follows:

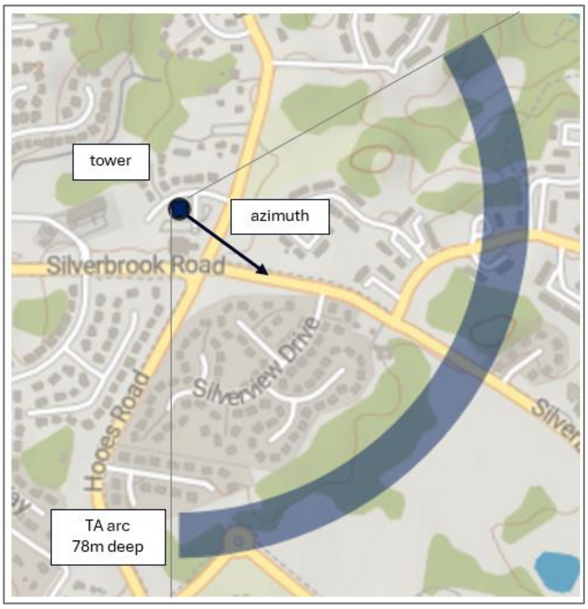

4. What if distance isn't the only delay component?

This form of analysis, inferring a suspect device's distance from a tower based on TA data, is widespread, extensively used by law enforcement and commonly accepted as evidence by US courts to aid the geolocation of devices.

The basis of the distance inference is that the only component in the measured delay is the device's straight-line distance from the tower, but what if straight-line distance isn't the only delay component?

The diagram on the following page shows three scenarios in which additional factors could contribute to the measured delay:

- The cell is transmitted from a very tall tower or from the top of a tall building – if the base station (where the TA measurement is taken) is at ground level, then the additional distance up to and down from the antennas will add to the delay without increasing the straight-line distance.
- The cell is transmitted by a DAS (Distributed Antenna System) with a significantly long fibre optic cable running between the base station (where the TA is measured) and the antenna. If the map plotting the TA value shows the arc location in relation to the DAS antenna site instead of the base station, then the arc could be wrongly positioned by a value equal to the length of the fibre connection.
- If the transmission suffers from multipaths. Multipaths are caused by reflections when a radio signal has no clear line of sight between transmitter and receiver. In dense urban areas (and also in mountainous, rural areas) the path taken by a radio connection might suffer multiple reflections and end up being longer than the straight-line, LoS distance between the tower and the device.

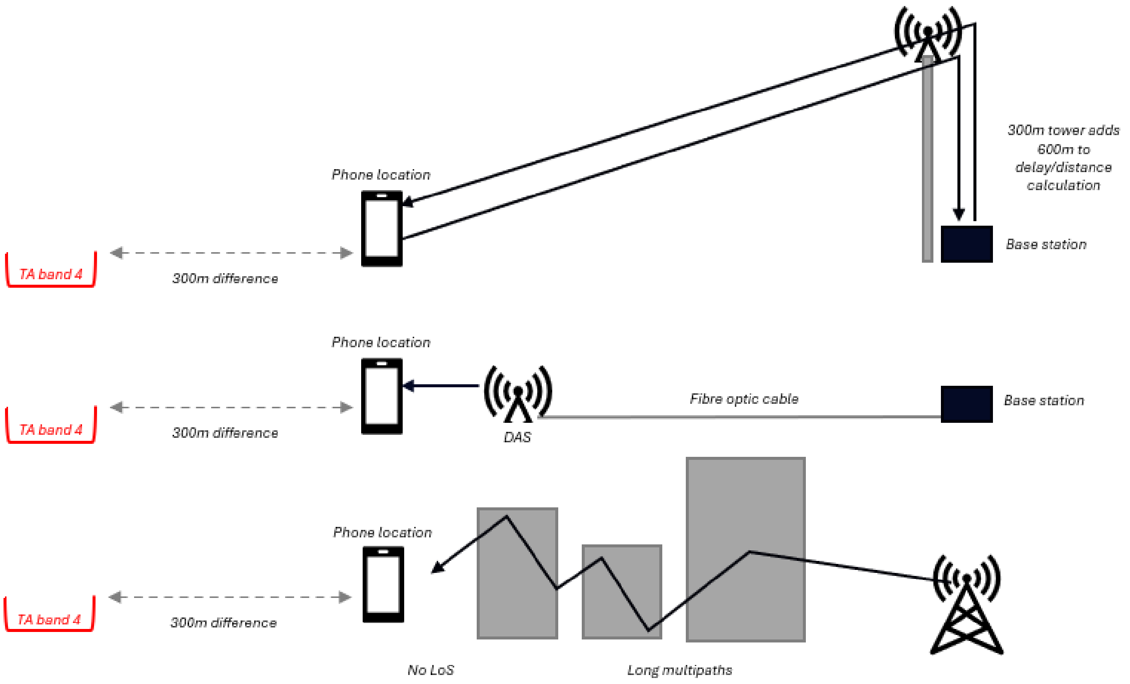

To add some context to the scenarios outlined above, it should be pointed out that the majority of towers, in urban and suburban areas at least, a rarely very tall, so the level of additional TA delay is likely to be small for most sites; there are ways to include the length of the fibre connection employed for a DAS site to compensate for the additional delay in the TA-distance calculation (see below); in the small cells used for the majority of urban and suburban cells, there is not enough distance involved for long multipaths to be employed, again minimising the effect that such a phenomenon would have on the TA-distance calculation.

In the majority of cases, we can expect the additional causes of delay suffered by the majority of cellular connections to amount to a few extra microseconds, which in turn amounts to a few tens or low hundreds of metres of difference to the distance inferred from TA data.

4.1 DAS & TA

DAS (Distributed Antenna Systems) are extensively employed by some operators. In a DAS, a base station is deployed at a central location and feeds service to a number of remote antennas.

The antennas don't process any traffic themselves, they relay downlink traffic from the base station to the phones and pass received uplink traffic from the phones back to the base station. The connections between the centralised base station and the distributed antennas in traditional DAS deployments are carried by fibre optic cables.

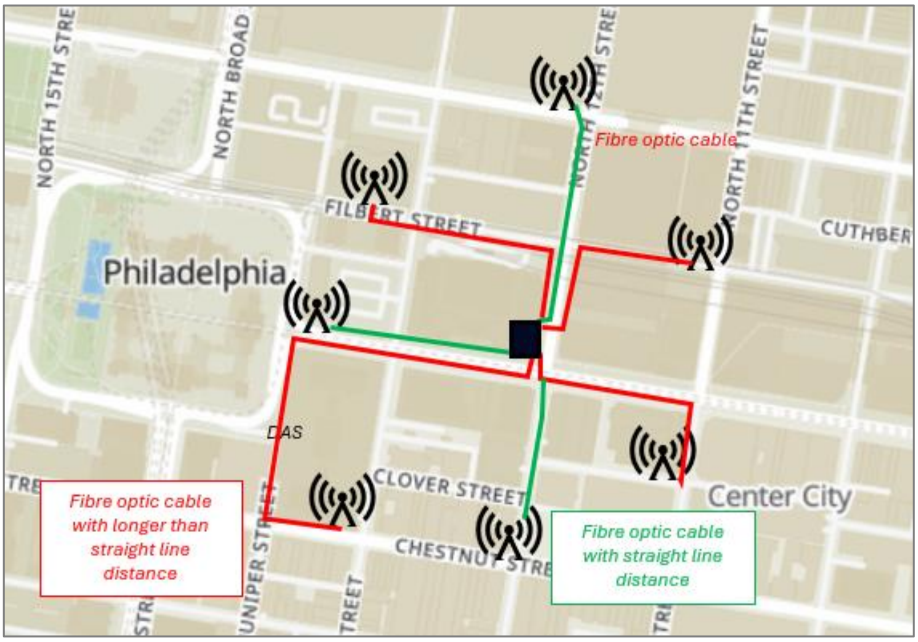
Fibre optic cable
Philadelphia
DAS
Fibre optic cable with longer than straight line distance
Fibre optic cable with straight line distance
Center City
FILBERT STREET
CLOVER STREET
CHESTNUT STREET

In effect, the additional run of fibre between the antenna and the base station adds to the delay suffered by signals transmitted by cell phones – the further the antenna is from the controlling base station, the more delay is suffered.

There are two potential issues for TA calculations when DAS is employed.

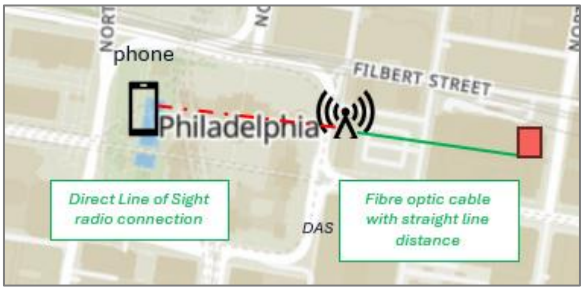
phone
Philadelphia
FILBERT STREET
Direct Line of Sight radio connection
DAS
Fibre optic cable with straight line distance

Firstly, if the fibre cable between a specific DAS antenna and the base station follows a direct path, then the delay suffered by cellular signals will be roughly comparable to the straight-line distance between those elements (albeit with slightly more delay as light travels more slowly through glass than through air).

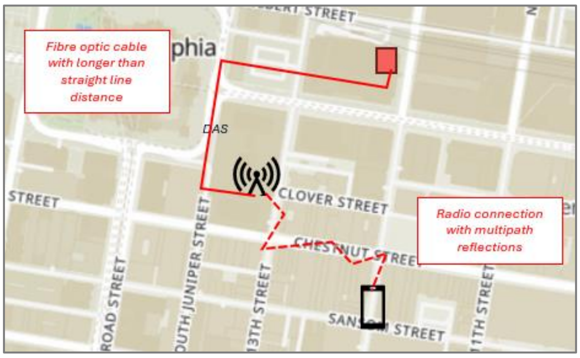
Fibre optic cable with longer than straight line distance
DAS
CLOVER STREET
CHESTNUT STREET
Radio connection with multipath reflections

If the fibre cable follows a less direct route, the delay will increase, but if the length of that fibre path isn't known, then any distance inferences made from TA data could over-estimate the suspect device's distance from the tower.

The difference in the speed of light through optical fibre compared its speed through a vacuum is significant – lights travels at around 66% of the usually quoted speed when it travels through glass, meaning that if a signal was calculated to have travelled 100m through vacuum during a specific time period, it would only have travelled 66m through optical fibre in that same time. This has the effect of pushing the calculated TA arc further away from the measurement point and potentially over-estimating the target phone's distance from the tower.

Secondly, when inferring TA distances, it is important to understand where the TA calculation is being made – for traditional DAS deployments it is made in the base station. But if the investigator is provided with 'tower' location details for the DAS antenna, which could be hundreds of metres away from the base station, then any inferred TA arc values mapped from that antenna location could be wrong.

For example, imagine that the DAS antenna used for a call is 2km from the base station. An investigator is informed that a suspect's phone was measured as being in TA band 46 (3.59-3.66km) but they are only provided with the antenna's location. If they map the arc from the antenna, it will cause them to infer a distance for the suspect phone that is around 3.5km from the DAS antenna site. As the TA measurement is actually captured at the base station, the TA arc should be mapped with reference to that location instead, meaning that it would be shown 3.5km from the base station (the green arc in the diagram) and not 3.5km from the antenna (the red arc).

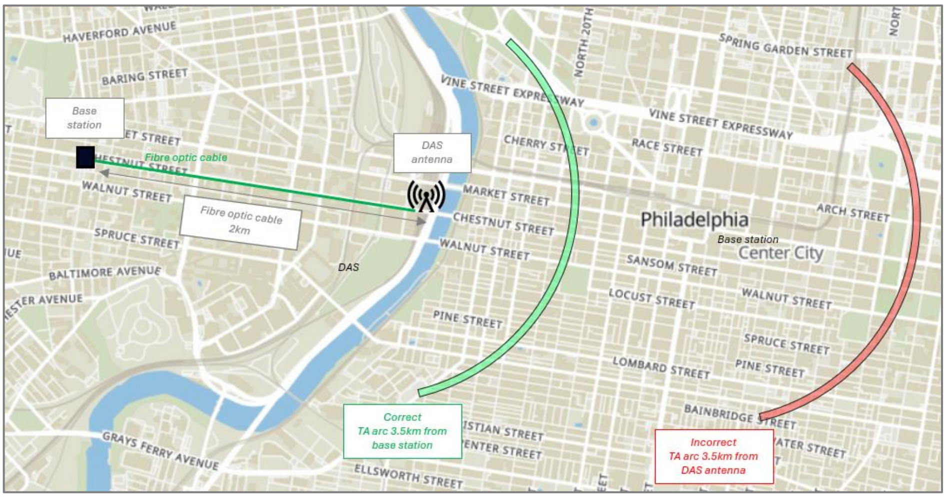

The recommended method for plotting TA arcs associated with urban outdoor DAS deployments – at least, as far as T-Mobile outdoor DAS sites are concerned – is to estimate or determine the length of the fibre connection between the central base station and the DAS antenna, to subtract this distance from the inferred TA distance and to plot the remainder from the antenna as a TA arc.

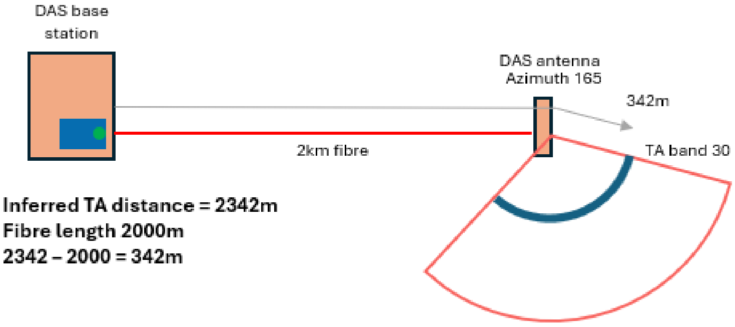

This may not be easy to achieve, so an alternative might be to say that you could consider not mapping TA arcs from DAS sites if the exact fibre length is not known.

4.2 Mapping TA data

All of these scenarios (and more) can affect the delay associated with a signal without changing the straight-line distance between the device and the tower, but without knowing if any of those additional delay components exist, we are forced to map TA data as if the straight-line delay was the only component.

This will occasionally lead to maps showing TA arcs that are further away from the tower than the device's actual location, or to put it another way, the TA arc will generally show the suspect device's maximum distance from the tower, but it might have been closer.

In urban areas, where each cell generally covers only a small area, the effects of tower height, DAS connections and multipaths are generally also small – they maybe add a few tens of metres to a signal's journey and therefore add only a few fractions of a microsecond to the measured delay. In urban areas, the effects of additional delay may only add a few tens of metres to the distance calculation, this may be multiplied in suburban or rural areas.

4.3 Interpreting TA data

TA evidence is routinely provided by the major US network operators, usually in csv, Excel or txt file formats.

T-Mobile currently provide TA data in the form of a distance value, like this:

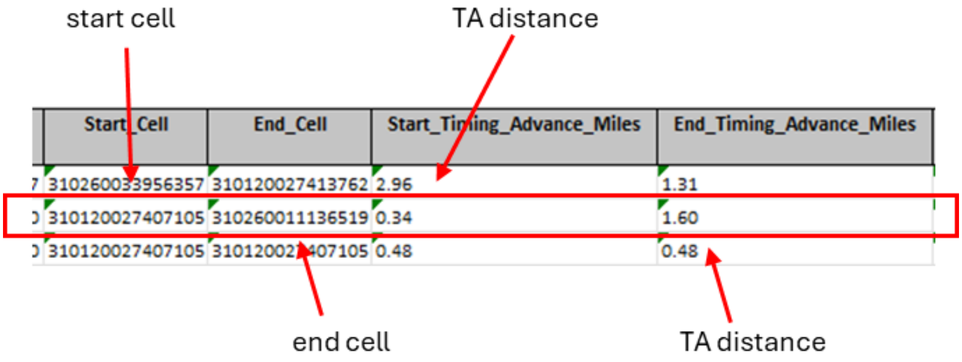

Start_Cell	End_Cell	Start_Timing_Advance_Miles	End_Timing_Advance_Miles
310260033956357	310120027413762	2.96	1.31
310120027407105	310260011136519	0.34	1.60
310120027407105	310120027407105	0.48	0.48

In this example, T-Mobile states that the target phone connected to cell ID 27407105 at the start of an event with an inferred TA distance of 0.34 miles.

0.34 miles isn't the exact distance measured from the phone to the tower, instead it can be interpreted as the distance to the inner edge of the TA band that begins at that distance from the tower.

Allowing for rounding differences, 0.34 miles is equal to 546.5m, which is the start of LTE TA band 7 (546.5-624.6m or 0.34-0.39mi). The T-Mobile and Verizon 'distance' values seem to all line up with the metric TA band distance values.

AT&T, in their most recent versions of TA disclosure, include the TA band number and not a distance. In their provided Records Key, AT&T describes this as a "unitless integer" and instructs the end user to multiply by 78.07 for 4G LTE and by 39.04 for 5G/NR cell sites to obtain the distance in metres.

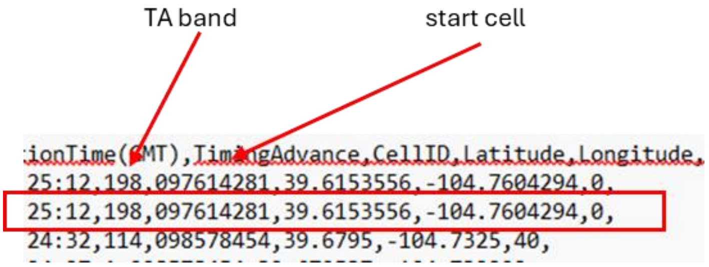

```
ionTime(GMT),TimingAdvance,CellID,Latitude,Longitude,
25:12,198,097614281,39.6153556,-104.7604294,0,
25:12,198,097614281,39.6153556,-104.7604294,0,
24:32,114,098578454,39.6795,-104.7325,40,
```

In this example, the phone was using cell ID 97614281, with a level of delay that placed it in TA Band 198 (15,458.0-15,536.1m).

Finally, Verizon, in one of their disclosure types, provides TA values as distances from tower, like T-Mobile.

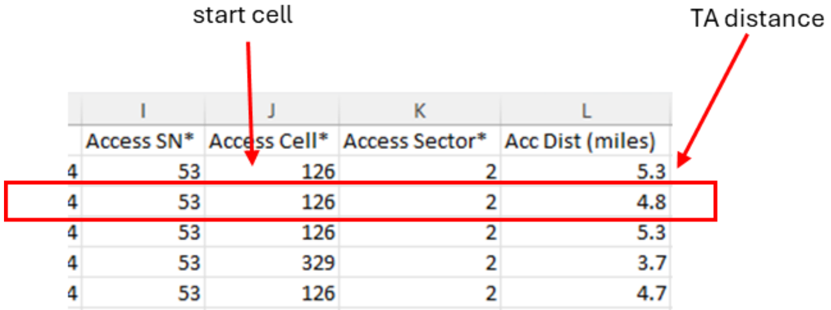

	I	J	K	L
	Access SN*	Access Cell*	Access Sector*	Acc Dist (miles)
4	53	126	2	5.3
4	53	126	2	4.8
4	53	126	2	5.3
4	53	329	2	3.7
4	53	126	2	4.7

In this example, the target phone had a level of delay that equated to TA band 99 (7,729.0-7,807.0m or 4.8-4.85mi).

It can be seen that none of these disclosures appears to be providing an exact distance measurement to the target phone. Indeed, Timing Advance doesn't work with exact distances, so the best we can get is the approximate area covered by a TA band.

5. How precise can inferred TA distances be?

The minimum granularity of the TA inferred distance calculation is about 5m (for LTE) – based on the value of T_s or 1/(FFTxSCS).

The LTE TA calculation is $T_{TA} = 16 \times TA \times T_s$, but the sizing of a single TA band is 16 x T_s.

That means that the TA band depth is calculated from 16 times the T_s value. In theory, each TA band could be subdivided into 16 sub-bands.

The depth of the sub-band will be $(T_s \times c)/2$ – it is divided by two due to the calculation initially working out the two-way delay.

This therefore works out as:

$T_s = 1/(\text{max FFT} \times \text{SCS}) = 1/(2048 \times 15000) = 3.26\text{x}10^{-8}$

The distance covered by one sub-band $(T_s \times c)/2 = (3.26\text{x}10^{-8} \times 299{,}792{,}458)/2 = 4.88\text{m}$

This means that the minimum granularity, the smallest variation in inferred distance that can be calculated, is around 5m. So if a target device is within ~5m of a TA band boundary, it is possible that it could be reported as being on the other side of that boundary, in the 'wrong' TA band.

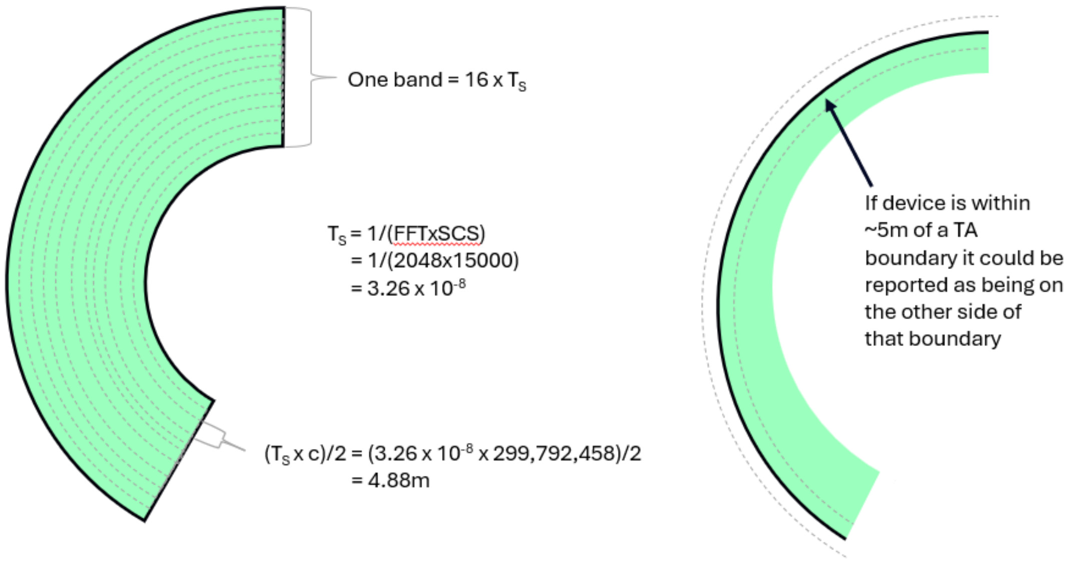

This means that TA data should really be interpreted as offering an *approximate* indication of a device's distance from the tower and that instead of indicating that the device could only have been in the reported TA band, it might be more realistic to acknowledge that the phone's actual location may have been within one or two TA bands of the stated band.

The question then becomes 'if TA data isn't always exactly precise, can we put any bounds on how often it is imprecise and by how much?'.

To fully understand the question, we must provide some interpretation of the word 'precise'. The context in which this word is used in this document is in the sense that a set of TA values have indicated a TA band that corresponds to the exact distance a suspect's phone was from the tower when those details were captured by the network. 'Precise' would mean that the TA data reported the band that corresponded with the suspect's phone's actual distance from the tower. 'Imprecise' means that the suspect's phone was at a different distance from the tower than was indicated by the stated TA band.

To answer this question, we set up a test project. There are two propositions that can be tested – the level of precision of TA data presented by a test device and the level of precision of TA data captured by the network. Both propositions require knowledge of the test device's actual location at the time that connections were made.

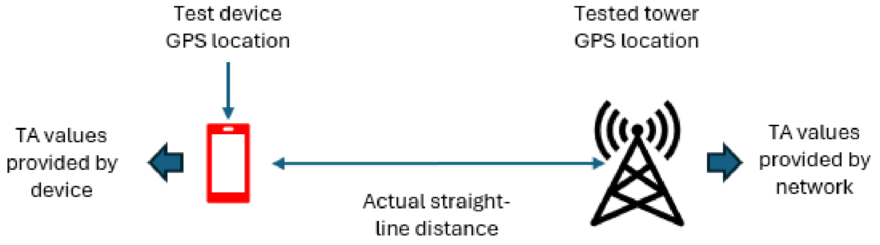

6. Device-side Testing

For device-side testing, we used the following:

- Lima Cell Monitor and a separate test rig consisting of an LTE modem and a Python script which sent AT commands to the test device
- A GPS receiver
- A tower with a known location

AT commands provide a method for controlling the actions of a modem from a connected PC. In our test rig, we connected a Windows laptop to a cellular radio modem (essentially the radio part of a cell phone, without the rest of the components that make up a complete phone) via a USB cable. Our Python script sent commands to the modem (e.g. set up a data connection, measure the signal strength of the serving cell, request the current TA value from the modem, etc) and then captured the responses.

For example, the AT command to request TA data from the modem we used was:

```
AT+QNWCFG="lte_time_advance"
```

And the response that was returned looked like this:

```
+QNWCFG="lte_time_advance",1,10
```

Which indicates that the delay experienced by signals to and from the device during its last Random Access event fell somewhere within TA band 10.

AT commands are a standard method for controlling modems, but each specific type of modem has its own specific AT command set, so the commands listed above might not work on other modem types (we used a Quectel RM502Q-AE).

Based on the local control of the test modem, the device-side test methodology we followed was:

- Select a specific tower location and capture the cell ID(s) of the cell(s) broadcast by that tower
- Start the survey equipment and start driving away from the tested tower
- The survey equipment automatically generates a connection (SMS or data event) every 10 seconds or so – for each event we capture or calculate:
    - Device's current GPS location
    - Current straight-line distance to tested tower
    - Current TA value provided by the device
    - The inner and outer band edges of that TA band
    - An indication of whether the device's current actual distance from the selected tower is within the reported band
    - If not, the number of bands difference between the reported band and the device's actual location

The testing process was extended as it took some time to work out the exact set of steps required to force the device to provide a regularly updated TA value – initially we found that 'old' TA values continued to be reported for an unpredictable period of time, which invalidated the results.

Once the optimum set of actions was determined, we found that the reported TA value was getting updated for most measurement events.

In testing undertaken by Joe Hoy and Martin Griffiths, in Letchworth, UK, on the O2 4G network in December 2024, we obtained the following results:

Reported bands difference	Description	%
-1	Reported one band further in than actual	11%
0	Actual distance within reported band	25%
+1	Reported one band further out than actual	39%
+2	Reported two bands further out than actual	16%
Other values		9%

These results show that 25% of events had device-side GPS locations that fell within the reported TA band. 55% of events showed that the reported band was one or two bands further out from the tower than the device's actual distance – this means that there was more delay associated with the connection for those events, more than could be accounted for by the straight-line distance to the tower, which caused the inferred distance to be slightly further than the actual distance.

Interestingly, 11% of events reported the test device to be up to one band closer to the tower than the reported TA band. This doesn't make logical sense on first viewing – it means that there was less than the expected amount of delay in relation to the test device's actual distance from the tower, or to put it another way, the signal from the device got to the tower faster than expected, which could only happen if the signal travelled faster than the speed of light!

In reality, the majority of these 'negative' results are examples of where the test device was close to the inner edge of a TA band, within the 5m minimum granularity, and got reported as being just into the next closer TA band. A band difference of -1 band can be explained in this way, differences of more than -1 (which place the device even closer to the tower than it actually was), if they were ever reported, would require a different explanation.

Finally, 9% of events returned unusable TA values – a couple of measurements provided TA band values of 1312 and 1328, for example, whereas the maximum reportable TA step value in LTE is 1282, so these results were disregarded. Results like this are a consequence of pulling the data straight from the modem, they are not something we would expect to see when using TA measurements supplied by the network as values of that kind wouldn't be measured on the network side.

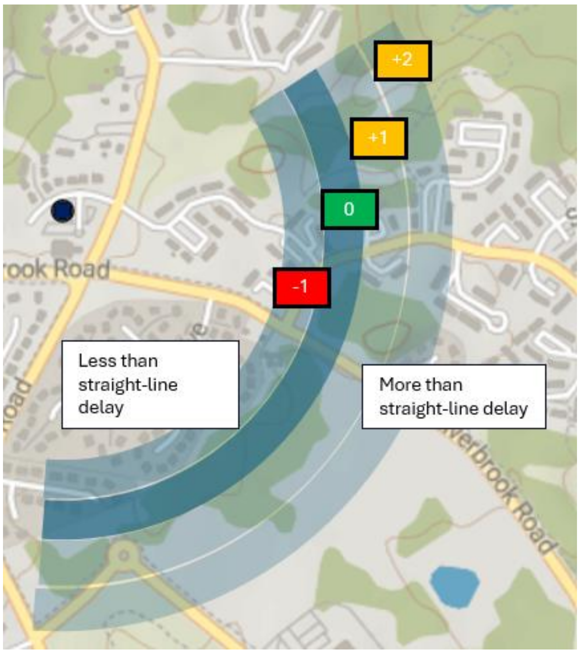

If the results of this limited testing can be extrapolated to provide a general conclusion, it seems that 91% of device-side TA reports place the suspect device within +2 to -1 band steps of the device's actual distance from the tower, which offers a precision of around 200-250m.

Although this is a long way from being able to state that TA values appear to be 100% precise, it does provide some reassurance that the imprecisions are explainable and are within reasonable levels of accuracy.

However, it must be borne in mind that the results of this phase of testing were obtained by interrogating the test device locally, whereas the TA results that are provided for evidential purposes are obtained from the network-side.

Example TA band difference scenarios:

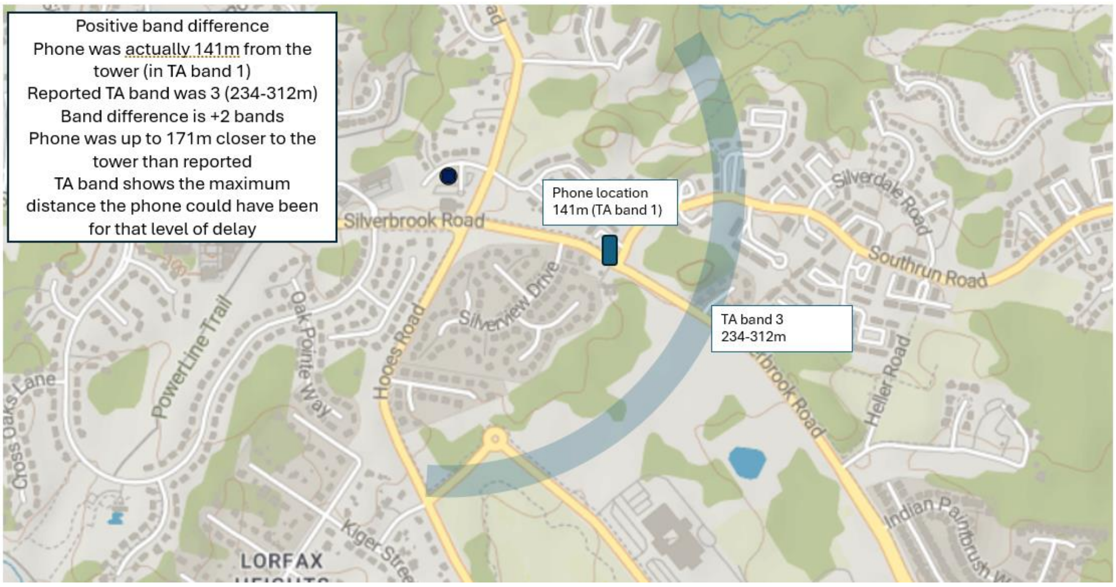
Positive band difference
Phone was actually 141m from the tower (in TA band 1)
Reported TA band was 3 (234-312m)
Band difference is +2 bands
Phone was up to 171m closer to the tower than reported
TA band shows the maximum distance the phone could have been for that level of delay
Phone location 141m (TA band 1)
TA band 3 234-312m
Silverbrook Road
Silverdale Road
Southrun Road
Silverview Drive
Hooes Road
Oak Pointe Way
PowerLine Trail
Cross Oaks Lane
Heller Road
Indian Paintbrush W
Kiger Street
LORFAX HEIGHTS

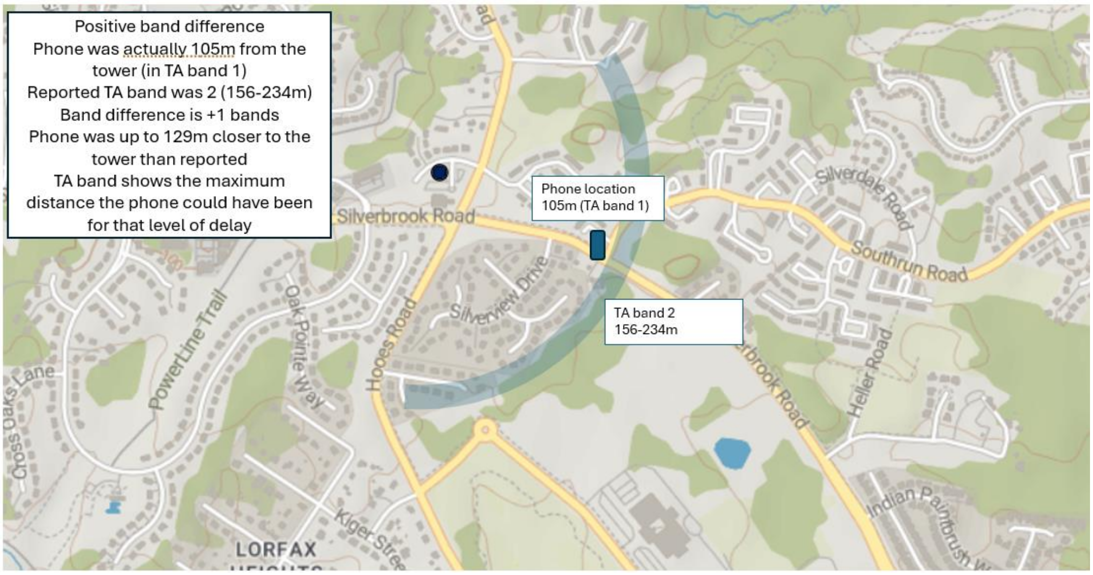

Positive band difference
Phone was actually 105m from the tower (in TA band 1)
Reported TA band was 2 (156-234m)
Band difference is +1 bands
Phone was up to 129m closer to the tower than reported
TA band shows the maximum distance the phone could have been for that level of delay

Phone location
105m (TA band 1)

TA band 2
156-234m

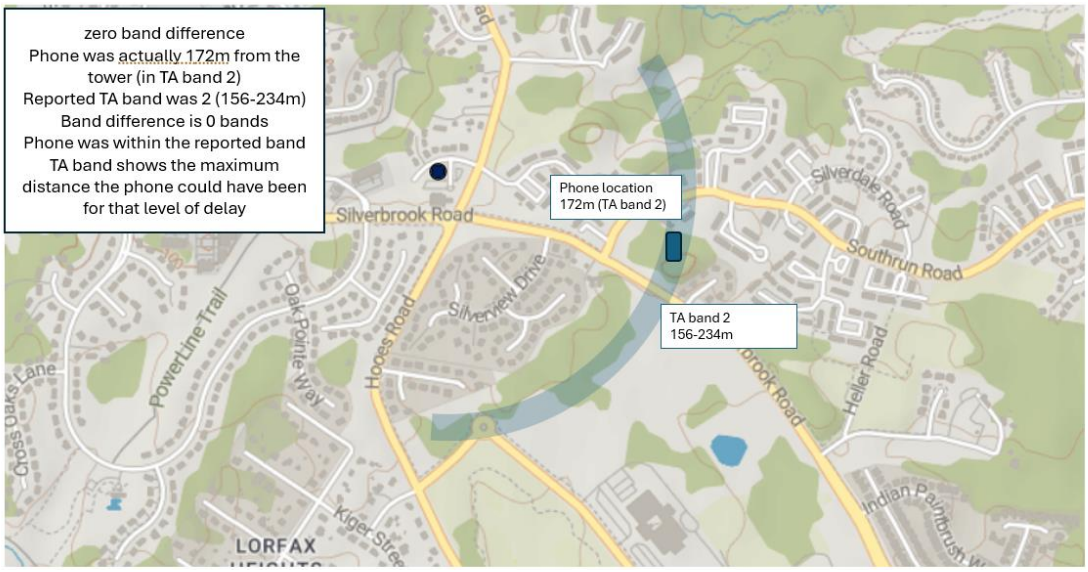

zero band difference
Phone was actually 172m from the tower (in TA band 2)
Reported TA band was 2 (156-234m)
Band difference is 0 bands
Phone was within the reported band
TA band shows the maximum distance the phone could have been for that level of delay

Phone location
172m (TA band 2)

TA band 2
156-234m

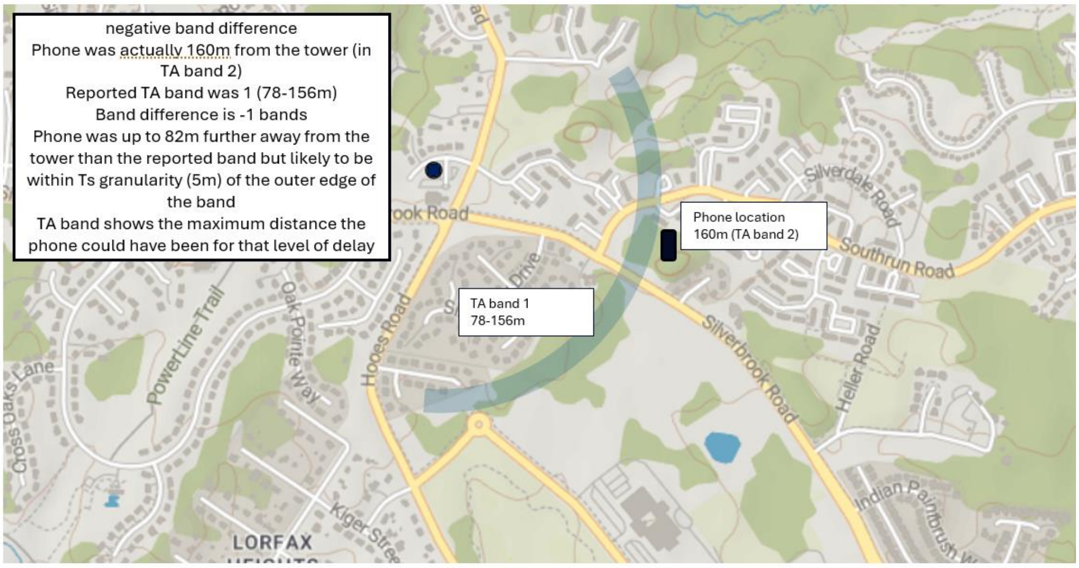

negative band difference
Phone was actually 160m from the tower (in TA band 2)
Reported TA band was 1 (78-156m)
Band difference is -1 bands
Phone was up to 82m further away from the tower than the reported band but likely to be within Ts granularity (5m) of the outer edge of the band
TA band shows the maximum distance the phone could have been for that level of delay

Phone location
160m (TA band 2)

TA band 1
78-156m

7. Network-side Testing

The next phase of testing followed the same methodology as before but obtained the network TA results for the test device as well.

In this phase, testing was undertaken in the US, as we don't have access to this type of disclosure from UK networks yet. The testing was undertaken by Investigator Cory Brodzinski of the Denver District Attorney's Office and other Denver-area law enforcement officers between March and April 2025 on the AT&T 4G network.

Two locations were chosen, a suburban location in north Denver and a rural location to the east of Denver. The area in north Denver was chosen as a typical suburban community. The rural location on the eastern plains of Colorado was chosen as it is as close to flat and obstruction-free as could be practically found.

In both cases, the testers had one or more AT&T test phones and an RF survey device (a TSMA for the urban test and a Lima for the rural one). The RF survey device was employed to capture GPS data showing the device locations at the time that TA events were captured.

During each survey, the test phone was used to establish connections and other actions that would hopefully generate some TA data captured by the network.

After the surveys were completed, the TA data was requested from AT&T. A comparison could then be made of the survey device and test phone's location (they were all in the same car) at the time of each TA event and the TA band registered by the network.

Our analysis calculated the straight line distance between the survey device's GPS fix and the lat/longs for the tower that provided the TA measurement. We then worked out the TA band that the straight line distance falls within and compared that to the TA band provided in the TA data at the same time.

The analysis was looking to see how many 'bands difference' there was between the stated TA band and the TA band inferred from the measured straight line distance.

7.1 Suburban testing

The suburban test was conducted in an area north of Denver.

The test equipment consisted of a Rohde & Schwarz TSMA survey device and a handset with an AT&T test SIM.

The survey lasted for around 45 minutes, during which there were 29 TA events captured in the AT&T data. Those TA events were time aligned with the closest matching TSMA event to establish the location of the devices at that time.

Reported bands difference	Description	%
-1 (less delay than expected)	Reported band was 1 further in than actual	17%
0 (expected delay)	Actual distance within reported band	61%
+1 (more delay than expected)	Reported band was 1 further out than actual	22%
+2 (more delay than expected)	Reported band was 2 further out than actual	0%
Other values		0%

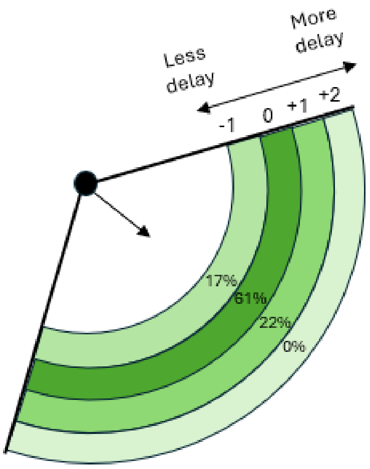

The results showed that 100% of TA events agreed with the test device's GPS location to within two TA bands of the stated band, with 61% of TA events accurately bracketing the phone's measured location.

7.2 Rural testing

The rural tests were conducted in an area around the town of Strasburg, to the east of Denver, which was chosen due to the flat geography.

The intention was to determine how well TA data agreed with the device's GPS location over longer distances.

The image alongside provides the reader with an idea of exactly how wide-open and flat the area is. The tower was relatively tall, there were no large buildings or hills and line of sight was maintained with the tower for most of the survey duration.

The test equipment consisted of a Lima cell monitor survey device and three handsets with AT&T test SIMs.

The survey lasted for around one hour and 6 minutes, during which time there were between 23 and 37 TA events captured in TA data for the three phones. Those TA events were time aligned with the closest matching Lima event to establish the location of the devices at that time.

Phone 1 – stated TA band versus GPS location

Reported bands difference	Description	%
-1 (less delay than expected)	Reported band was 1 further in than actual	16%
0 (expected amount of delay)	Actual distance within reported band	54%
+1 (more delay than expected)	Reported band was 1 further out than actual	19%
+2 (more delay than expected)	Reported band was 2 further out than actual	11%
Other values		0%

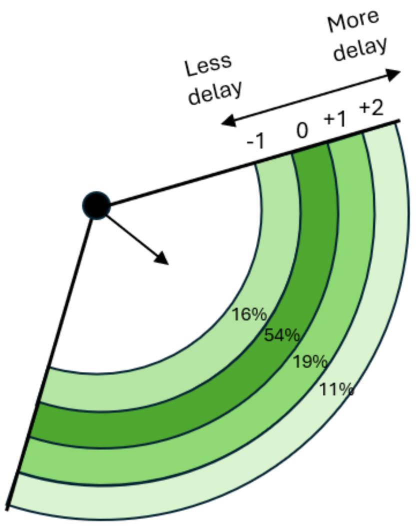

The results showed that 100% of TA events agreed with the test device's GPS location to within two TA bands of the stated band, with 54% of TA events accurately bracketing the phone's measured location.

Phone 2 – stated TA band versus GPS location

Reported bands difference	Description	%
-1 (less delay than expected)	Reported band was 1 further in than actual	0%
0 (expected amount of delay)	Actual distance within reported band	33%
+1 (more delay than expected)	Reported band was 1 further out than actual	51%
+2 (more delay than expected)	Reported band was 2 further out than actual	16%
Other values		0%

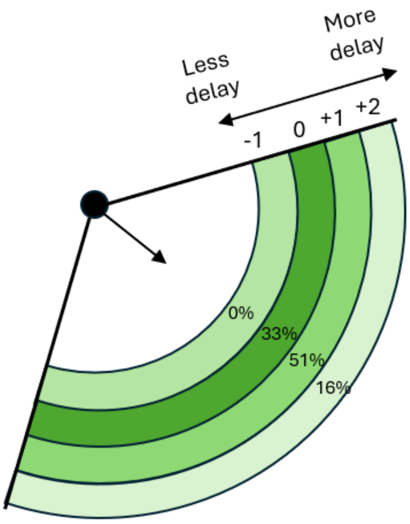

The results showed that 100% of TA events agreed with the test device's GPS location to within two TA bands of the stated band, with 33% of TA events accurately bracketing the phone's measured location.

Phone 3 - stated TA band versus GPS location

Reported bands difference	Description	%
-1 (less delay than expected)	Reported band was 1 further in than actual	0%
0 (expected amount of delay)	Actual distance within reported band	17%
+1 (more delay than expected)	Reported band was 1 further out than actual	63%
+2 (more delay than expected)	Reported band was 2 further out than actual	20%
Other values		0%

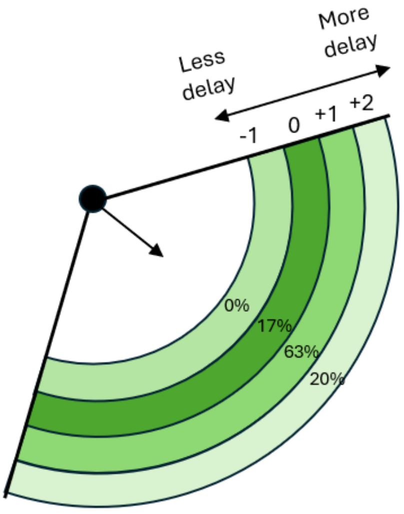

The results showed that 100% of TA events agreed with the test device's GPS location to within one TA band of the stated band, with 63% of TA events accurately bracketing the phone's measured location.

It can be seen that, in both suburban and rural environments, based only on the tests we have conducted so far, the target phone seems to be within the stated TA band or the next TA band closer to the tower the majority of the time.

The TA process would report the phone's location as being one or two bands further from the tower if there was more than just straight-line delay involved in the connection – given the flat geography of the tested area, that additional delay is unlikely have been caused by multipaths and can probably be attributed to the height of the tower and processing delays in the base station.

Whilst this is not as precise as being to state that a phone is always within the stated TA band, it is significantly more precise than stating that the phone could have been anywhere within that cell at that time.

The raw data from the tests we conducted can be made available on request.

Anecdotally, we've also been told about a dataset where the investigators requested TA data for an individual who was also wearing a GPS ankle monitor. In this dataset, the GPS locations were within 500ft (roughly two TA bands) of the stated TA band 99% of the time and were within 250ft (roughly 1 TA band) 70% of the time. This seems to correlate well with the results described from our testing above and we'll try to include this data in a later version of this paper, once the related case has been through court.

8. Interim Conclusions

Timing Advance data **does not** provide evidence of a suspect device's exact distance from the serving tower with a guaranteed level of precision – it does not necessarily 'pinpoint' the phone's location.

Instead, Timing Advance data provides evidence of the delay suffered by the two-way connection with the suspect device, from which a straight-line distance from the tower can be **inferred**.

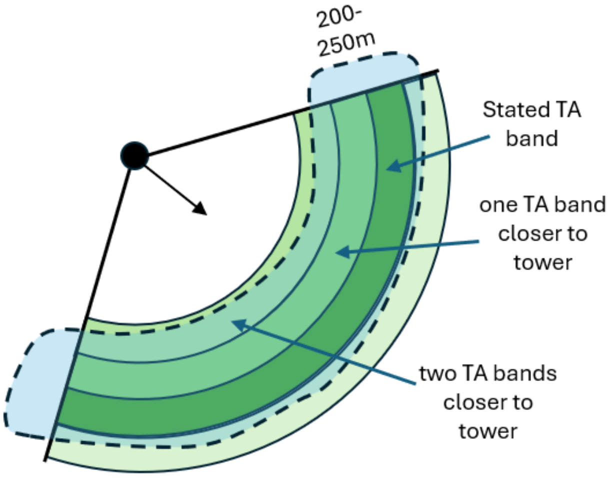

There is no guarantee that the inferred distance is completely precise, but testing so far indicates that TA is generally precise to within around 200-250m (650-820ft), which is approximately the distance covered by three TA bands.

Additional factors, beyond just the straight-line, horizontal distance from the tower, can cause additional delay, which in turn can cause the distance inferred from TA data to be slightly greater than the suspect device's actual distance was (which places the target phone in TA arcs that a little further away from the tower than where the device actually was).

In general, we can assume, based on our testing, that the reported TA band shows the maximum distance the suspect device would have been from the tower, with a high likelihood that the device was actually in the stated band or was in either of the two bands closer to the tower from the reported band at the time.

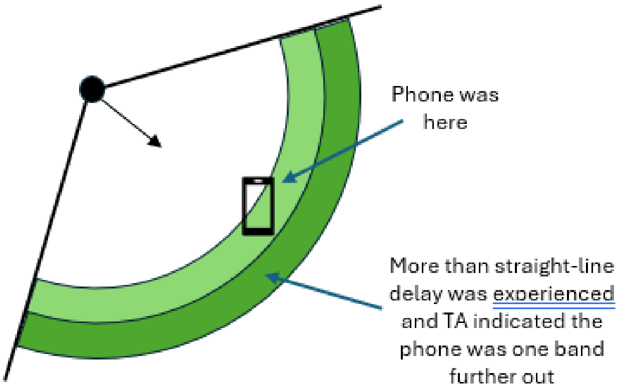
Phone was here
More than straight-line delay was experienced and TA indicated the phone was one band further out

In testing, the phone's actual location corresponded most often with the TA band next closest to the tower from the stated TA band, meaning that a small amount of additional delay had pushed the inferred TA distance out by one band.

With a limited sample size, the applicability of these results in all environments can't be guaranteed and further testing is required in additional types of area (dense urban, mountainous rural, etc.) to see if TA events in these areas are any more or less likely to be precise. Testing the repeatability of the results in areas already tested is also required.

Timing Advance has proven to be an immensely useful form of geolocation evidence. It is often more granular, in terms of the number events it shows, than CDRs, often including multiple TA events that occurred in between the calls, SMS and data session records shown in CDR data.

When coupled with knowledge of the cell site/tower location and the cell azimuth, it is possible for TA to determine a more precise potential location for a device than knowledge of just the site location and azimuth allows.

The testing undertaken as part of this report has shown that TA can be regarded as offering reliable geolocations for suspect devices, as long as the appropriate level of granularity is taken into consideration – TA may not always be able to reliably geolocate a device to within 78m (e.g. within one specific TA band), but it can reliably geolocate a device to with 200-250m (two TA bands) of the stated TA band.

Brought to you by
Forensic Analytics

CSAS

CellView

CDAN Nexus

CSAS Find

Lima

Forensic Analytics Ltd
Pixmore Centre
Pixmore Avenue
Letchworth Garden City
SG6 1JG
United Kingdom

+44 (0) 800 158 3830
www.forensicanalytics.io