Brought to you by
Forensic Analytics

Outdoor DAS TIMING ADVANCE

PRECISION & RELIABILITY TESTING

July 2025

v1.3

Contents

1. Executive Summary

Timing Advance (TA) is a type of geolocation data provided by cellular networks. It is generally regarded as providing evidence of a target phone's approximate distance from the serving cell tower when connections were made.

A previous report published by this team outlined some considerations related to timing advance in general; this paper deals with TA as it relates to one specific deployment case, known as 'outdoor DAS'.

DAS – or Distributed Antenna Systems – provide a deployment method where the base station equipment that serves a location or an area is deployed centrally, with cables carrying radio signals out to a distributed set of remote antennas in the served area.

DAS is typically deployed indoors or underground, in places like shopping malls, sports stadiums and subway networks, but there is a particular class of DAS site known as 'outdoor DAS' that this research focuses on.

For TA data to reliably indicate a suspect phone's approximate distance from the tower, practitioners need to have confidence in a number of factors:

- the measurements on which TA are based must be stable – a phone at a given distance from a particular tower will always generate similar TA values
- the location of the antenna the phone is connected to must be clearly understood
- the location of the base station at which the TA measurements are captured must be clearly understood

Unfortunately, our research so far has indicated that none of the above factors can necessarily be relied upon in relation to some outdoor DAS sites.

TA measurements captured for some outdoor DAS cells appear to fluctuate across a wide range of delay/distance values even for test phones that are stationary. Some outdoor DAS cells are transmitted via up to four antennas deployed at some distance from each other (up to 1 mile or 1.5km) and the tower list entries for some outdoor DAS sites make it difficult to understand where the base station is located.

All these factors add up to TA from some outdoor DAS sites being demonstrably unstable, with distance estimates varying and a degree of uncertainty about where to map any derived TA distance arcs from.

Further testing is required, but our interim recommendation is that TA data derived from outdoor DAS sites specifically should be treated with extreme caution and shouldn't be considered to be precise enough to use as the basis of an evidential location estimate.

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)
- Joe Noyes – Palm Beach SO (US)
- Sam Chan – Orange County SO (US)

2. Timing Advance

Please see our previous report - *Timing Advance Precision & Reliability Testing* - for a detailed explanation of how TA works and how TA values can be interpreted. A brief summary of this information is provided below.

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

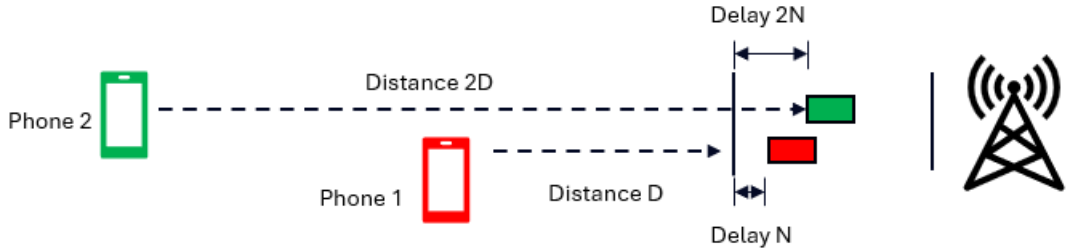

At the real distances that cellular signals are required to travel (usually up to a few miles), the differences in travel times at different distances would be in the microseconds, but even time differences of that level can have unwanted effects. Most cellular multiple access techniques, used in different 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.

TA 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.

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

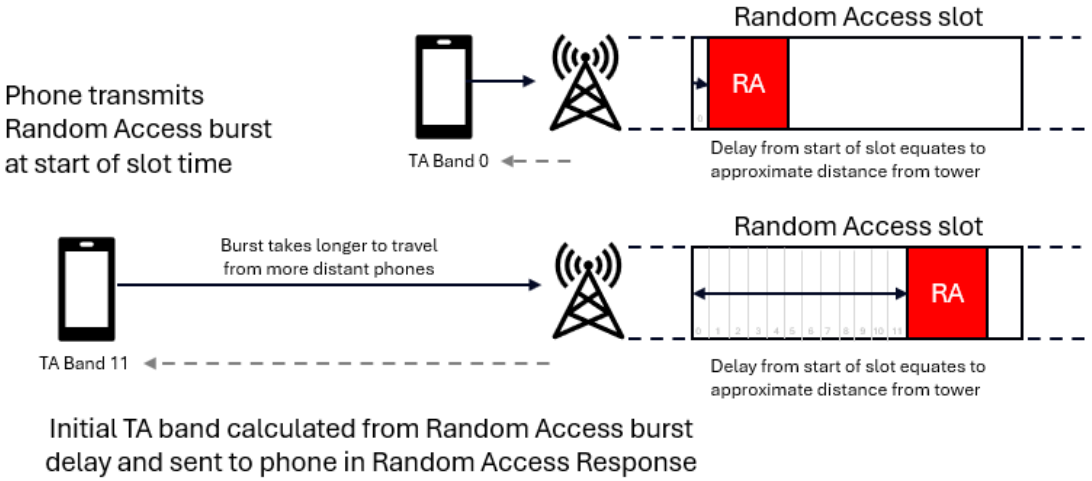

In the diagram above, the device at the top is close to the tower, its transmission arrives early in the allocated period and the base station determines that no TA adjustment is required.

The second device is much further away from the tower – its transmission therefore takes much longer to travel to the base station and arrives later in the allocated time 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.

The initial TA 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 TA 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, for example, 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.

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'.

LTE employs a fixed band size of 78.07m (256ft or 0.05 miles), the values vary in 5G, with band sizes of 78.07m, 39.035m and 19.51m depending upon an individual cell's configuration.

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

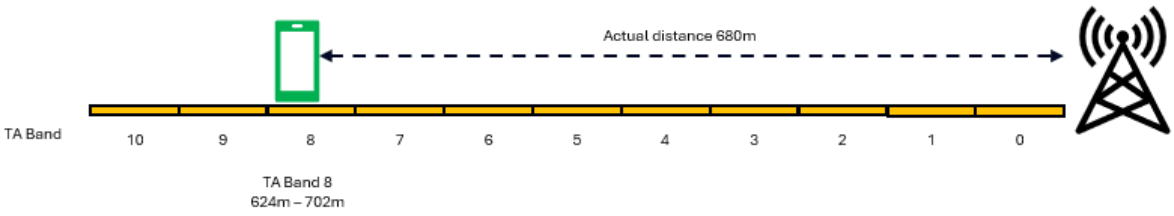

Imagine, for example, that an LTE device is 680m away from the serving tower – the actual two-way delay associated with that distance is equivalent to between 8 and 9 TA steps.

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.

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?

Signals connecting via very tall towers, signals that travel via longer routes between a phone and tower where there is no direct line of sight and signals carried by more complex deployment methods such as DAS, could all suffer more delay than just that which can be accounted for by the straight line distance between the phone and the tower.

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. But it means that sometimes, the distance inferred from TA data can be overestimated and any TA arc that is plotted on a map could be up to a few hundred metres further away from the tower than the suspect phone's actual location at that time.

3. DAS – Distributed Antenna Systems

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.

An example of a typical 'indoor DAS' deployment could be at a shopping mall.

Most shopping malls are designed to have an interior walkway, surrounded by the shops – this type of dense construction is often impervious to radio signals generated by towers situated outside the structure – the radio signals have difficulty penetrating through the outer ring of shops to reach customers on the interior walkways and the inner ring of shops.

In these environments, the cellphone networks often provide indoor coverage. It would be expensive to deploy a full base station at each point where indoor coverage was required, however, so in most cases, the shopping mall's architects would have set aside a specific room within the structure to house the base station equipment (often known a little whimsically as a 'base station hotel') and the providers would then run cables from the base stations to a network of remote (or 'distributed') antennas. The base stations contain the cellular radios and the cables to distribute the radio signals to and from the distant 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 indoor DAS deployments are typically carried by coaxial copper cables. Each antenna serves an area with a radius of a few tens of metres, meaning that TA measurements are largely irrelevant as evidence of geolocation, as anyone using one of those cells would have had to have been located within a few metres of the antenna.

3.1 Outdoor DAS

Some networks – but specifically T-Mobile – have adapted the DAS concept for outdoor use, hence this report's focus on 'outdoor DAS'.

Outdoor DAS can be seen as a cost-saving measure for deploying cellular coverage in dense urban environments, as, just as with indoor DAS, it means that operators don't have to deploy a base station to every 'tower' or cell site.

The diagram below illustrates a typical urban Outdoor DAS deployment.

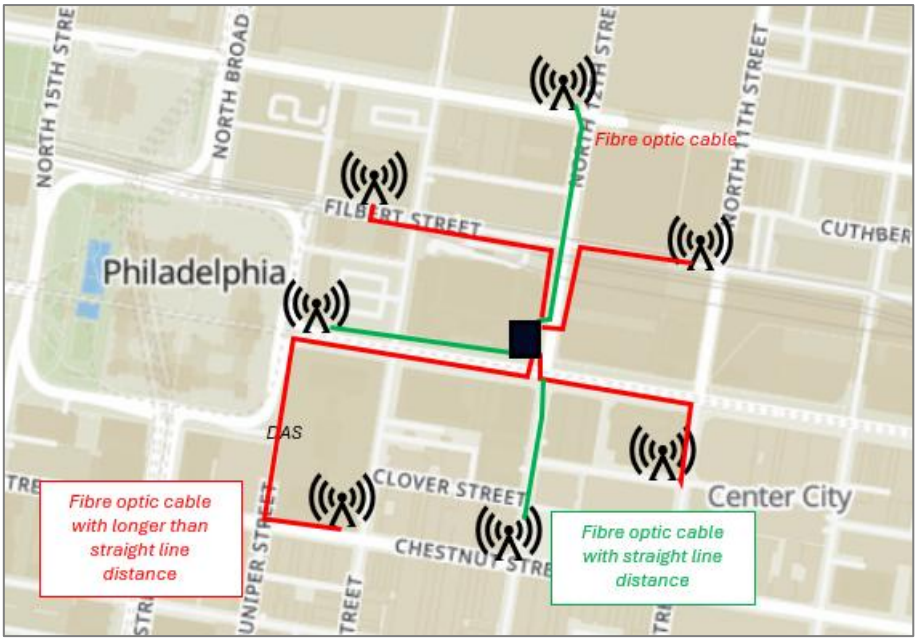

The base station is deployed at a central location. Antennas are deployed in neighbouring blocks, all served by radios installed at the central base station. Fibre optic cables link the base station to each antenna.

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. So the effect of using a DAS cell on the TA calculation for a phone is that the measured delay is the sum of the journey over the radio connection and through the fibre.

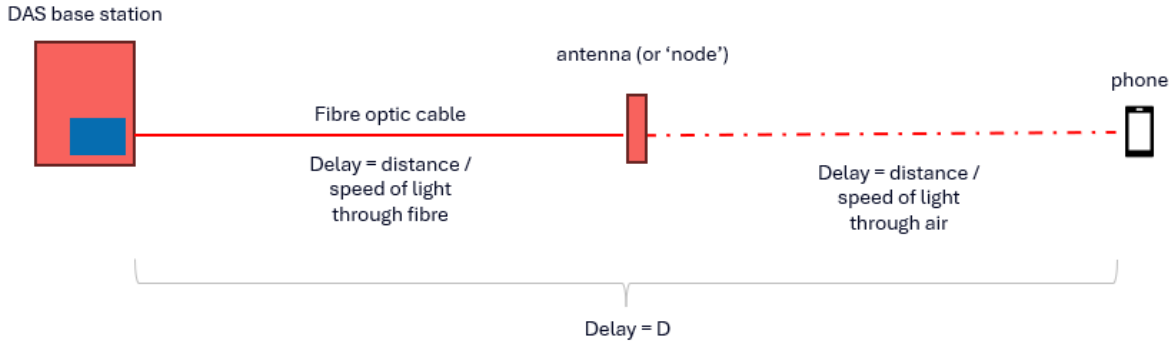

Light moves through fibre at a different speed than it moves through a vacuum or through air.

In fact, light moves through a fibre at about 2/3rds of the speed of light through a vacuum. This in turn means that cellular signals suffer more delay – travel more slowly – through the fibre connection than they do through the radio connection.

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

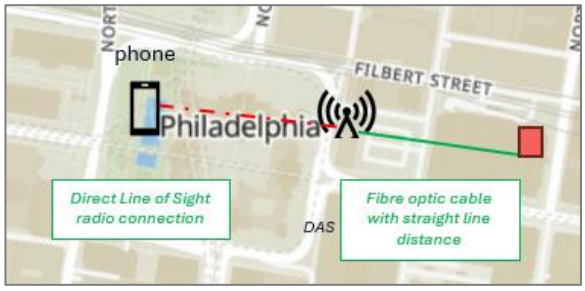

Firstly, if the fibre cable between a specific DAS antenna and the base station follows a direct path, then the delay suffered by cellular signals will be roughly comparable to the straight-line distance between those elements (albeit with slightly more delay as light travels more slowly through glass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

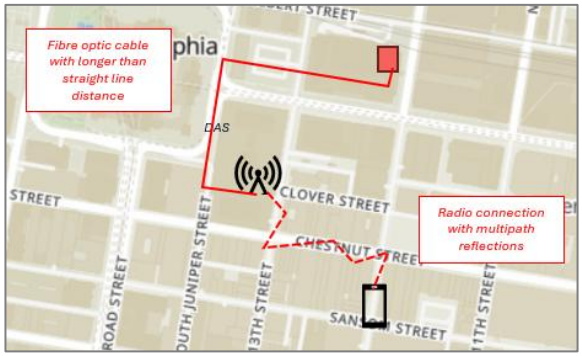

If the fibre cable follows a less direct route, the delay will increase, but if the length of that fibre path isn't known, then any distance inferences made from TA data could over-estimate the suspect device's distance from the tower.

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 radio transmitter/receiver, which is mounted 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.

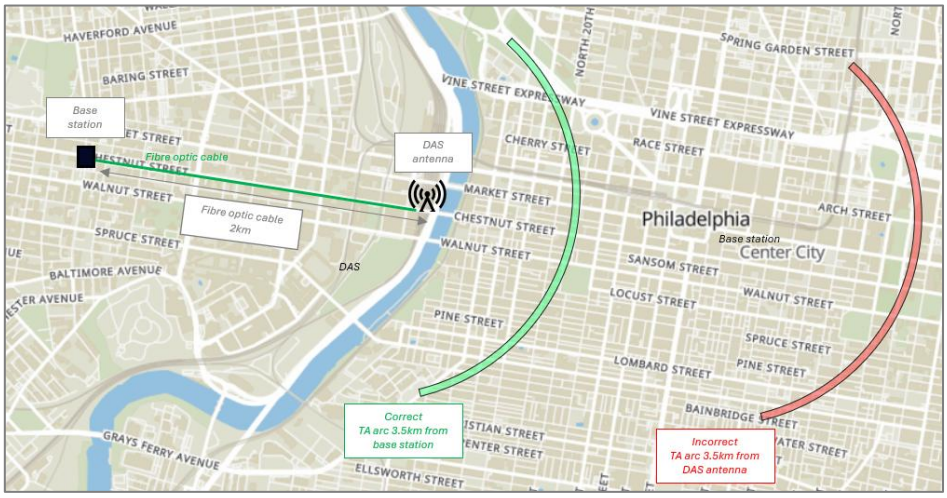

For example, imagine that the DAS antenna used for a call is 2km from the base station. An investigator is informed that a suspect's phone was measured as being in TA band 46 (3.59-3.66km) but they are only provided with the antenna's location. If they map the arc from the antenna, it will cause them 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, which is 1.5km from the antenna) and not 3.5km from the antenna (the red arc).

To recap a point made earlier: to be able to reliably plot TA distance inferences obtained from DAS, investigators need to have confidence in a number of factors:

- the measurements on which TA are based must be stable – a phone at a given distance from a particular tower will always generate similar TA values
- the location of the cellular antenna the phone is connected to must be clearly understood
- the location of the base station at which the TA measurements are captured must be clearly understood

In theory, if these three factors are known, it should be possible to compensate for the complexity associated with urban outdoor DAS deployments by estimating or determining the length of the fibre connection between the central base station and the DAS antenna, subtracting this distance from the inferred TA distance and plotting the remainder from the antenna as a TA arc.

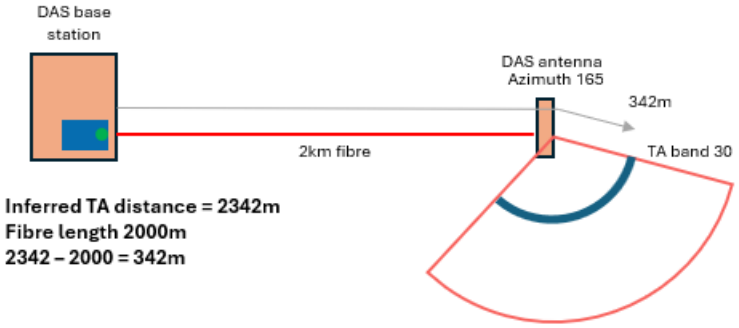

This method, whilst it might produce a more precise geolocation estimate, relies on several assumptions and may be considered to be opinion, and so may not be a form of evidence that every practitioner would be permitted to submit or would be comfortable relying upon.

The following sections provide more detail on the specific deployment model adopted by T-Mobile outdoor DAS sites and looks at some of the potential causes of delay in that model, before moving on to look at the potential effect on Timing Advance for phones using such cells.

4. T-Mobile Outdoor DAS

4.1 T-Mobile DAS Model

T-Mobile use a specific deployment model for some of their outdoor DAS sites.

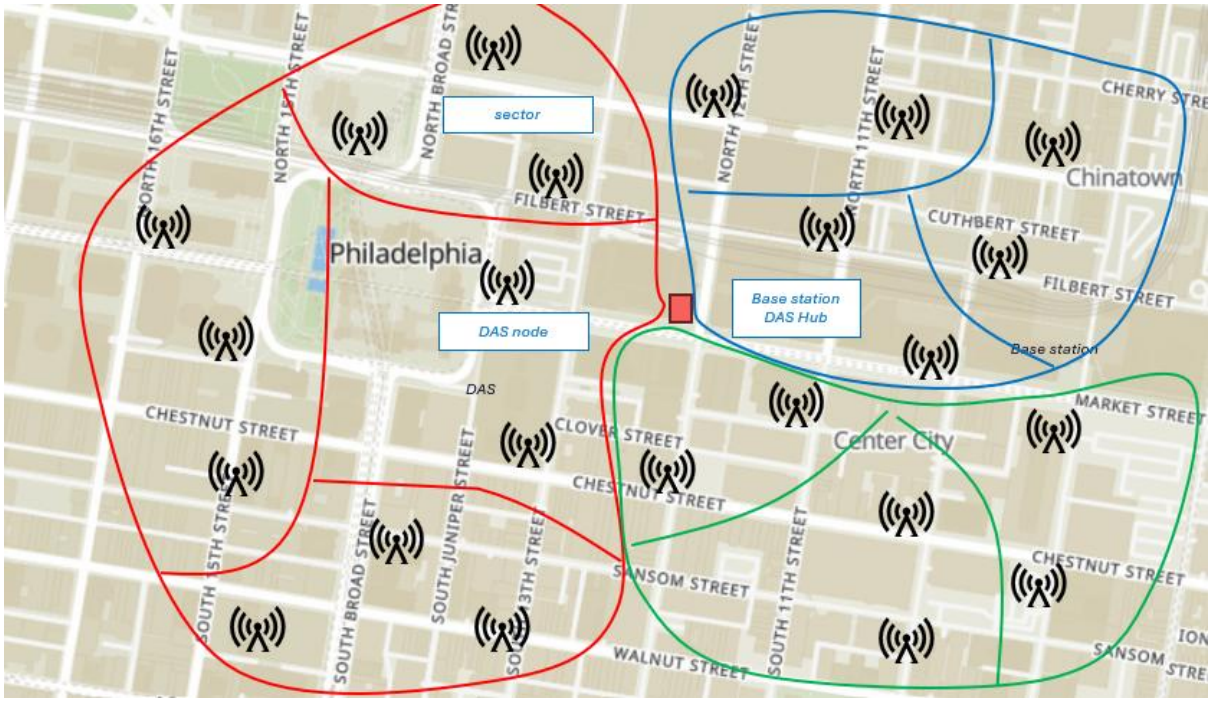

In an area served by this model of outdoor DAS, they will deploy:

- A central site to house the base stations known as a hub
- One hub might house several base stations, so each base station has its own service area – in the diagram, the hub hosts three base stations, each with a separate service area (coloured red, blue and green)
- The distributed antennas are known as 'nodes' – each hub serves a specific set of nodes
- The nodes are organised into 'sectors' – one sector will host between one and four nodes – a sector is analogous to a single cell
- All nodes in a sector propagate the same cell – if there are four nodes in the sector, all four nodes will transmit the same signal, which offers a form of 'transmit diversity'
- To be clear, in this model, up to 4 nodes will transmit exactly the same cell – if a phone uses that cell, it will not be easy to determine which of the four antennas it used from moment to moment
- The nodes are connected to the hub via fibre optical cable. Cables may be shared by different nodes. The transmission method employed on the cables is digital, meaning that RF signals need to be converted from analogue to digital and vice versa as they enter and exit the cables

4.2 T-Mobile Tower Lists

T-Mobile tower lists for outdoor DAS sites contain, amongst other things, the following details:

- Cell ID (eCGI)
- Cell azimuth
- Site address
- Site lat/long
- Sector lat/long

The naming convention used for these items leads us to expect that the 'site lat/long' coordinates link to the 'site address' location, but examination of the data shows that this isn't always the case.

4.3 Example of T-Mobile Outdoor DAS Site

Take the site in Denver, Colorado depicted in the diagram as an example:

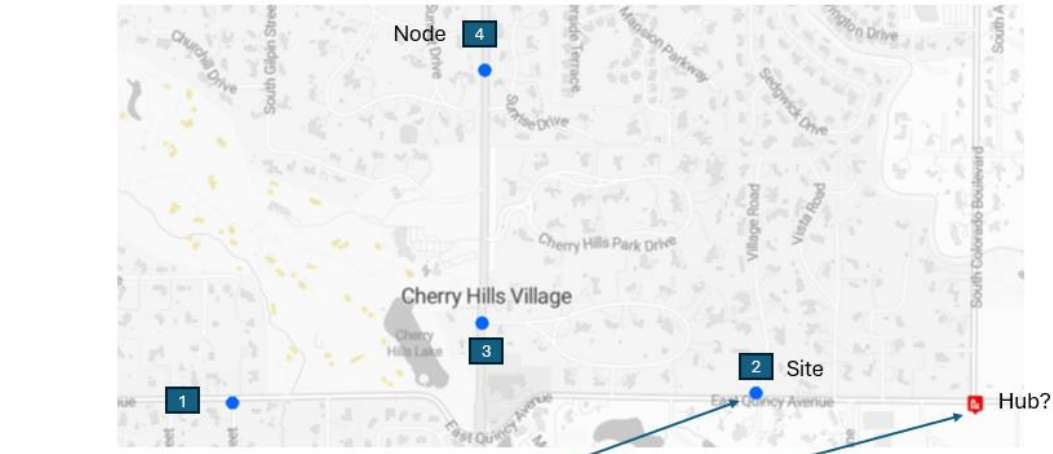

Cell_	NodeB_	CGI	Sector_C	Azimu	Latitude	Longitude	Site_Address	Sector_Latitu	Sector_Longitu
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.638524	-104.9691
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.638806	-104.94915
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.640863	-104.95958
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.648273	-104.959484

1 Node
2 Node & Site
3 Node
4 Node

LTE Cell ID 310-260-8618763 is listed four times in the tower list, with the same site address, the same site lat/long but different sector lat/long details. The obvious assumption to make is that the site address and site lat/long relate to the hub and the various sector locations are the nodes.

However, closer examination of the address data shows that the site lat/long doesn't match the site address. Instead, the site lat/long points to a location some 720m east of the site address.

A further obvious conclusion to draw is that the various 'sector' locations show different sets of historical data for the cell, that the antenna has been moved several times in the past and that only one of those sets of details shows the antenna's current location.

RF survey data (captured using Lima Cell Monitor) for the area shows different – this indicates that three of the four nodes are currently active and that all of them are

simultaneously propagating cell ID 310-260-8618763.

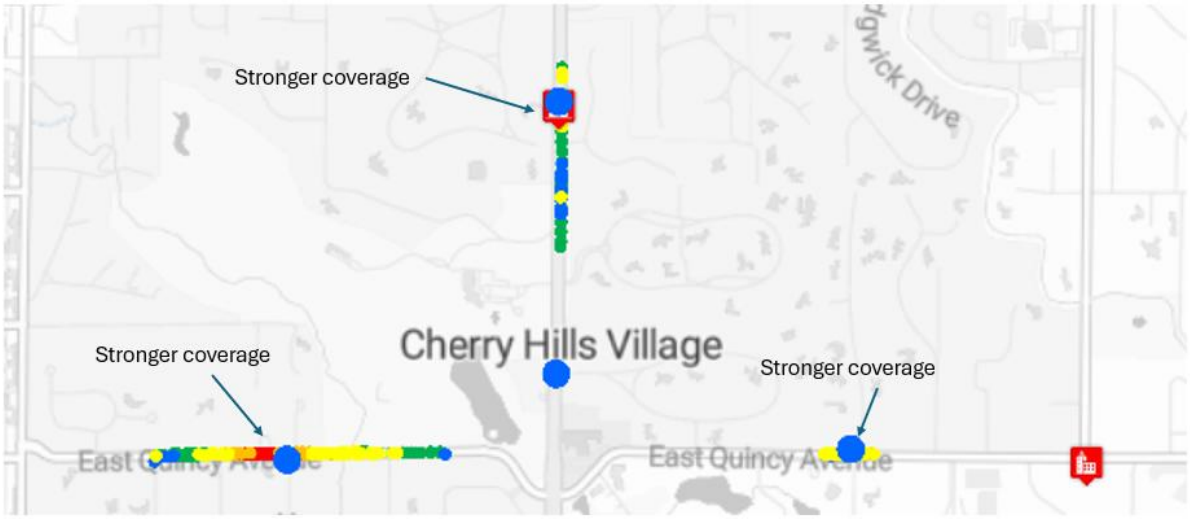

The localised nature of the RF coverage in the immediate vicinity of each node indicates that this isn't one coverage area from one antenna that covers the whole area, it is three separate coverage areas for the same cell ID.

This is, we believe, an example of a DAS hub supplying service to a sector that has four nodes (even if we only detected signals from three of the nodes).

4.4 Second Example

A second example of the T-Mobile Outdoor DAS site, this one in Palm Beach, Florida, shows an alternative configuration model.

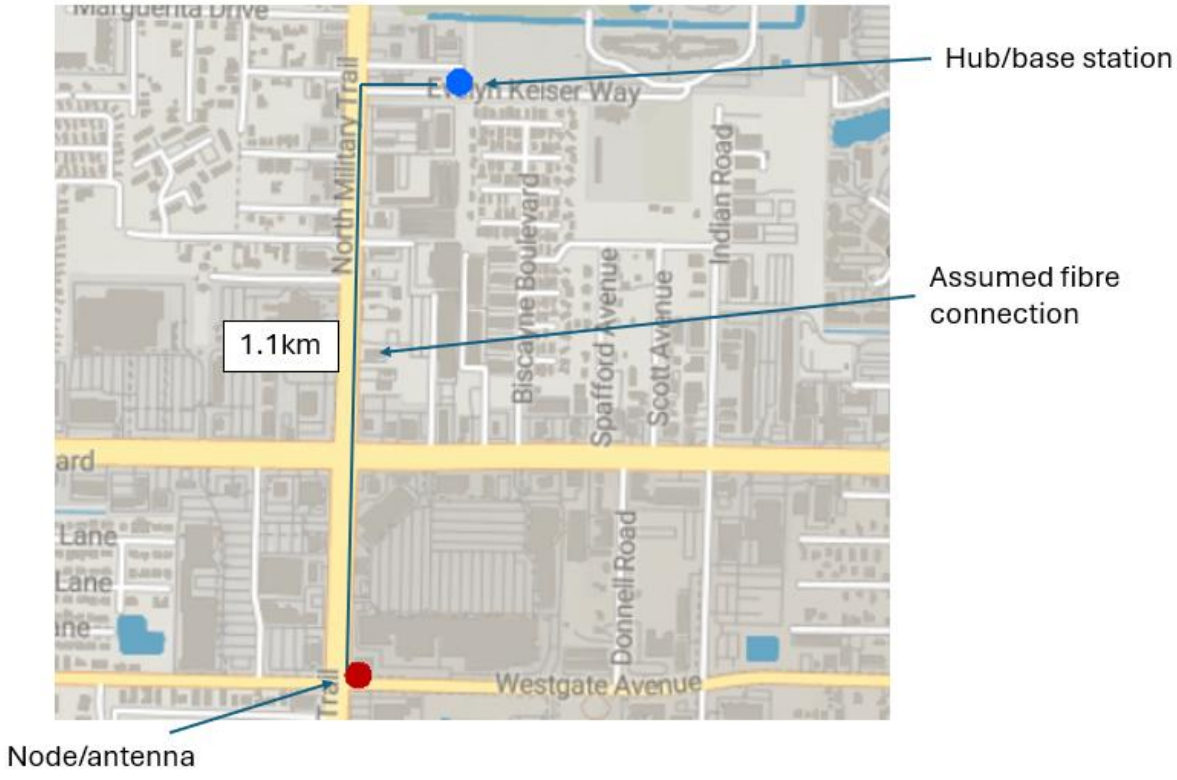

In this example, the hub serves just one node for a specific cell ID. Only one 'sector

lat/long' is listed for the cell ID in the tower list and only one coverage area for the cell was detected during an RF survey.

This would seem to be an example of a sector containing just one node.

4.5 Implications for Timing Advance

The 'hub-node' or 'hub-multiple node' deployment models employed by T-Mobile for outdoor DAS sites are both perfectly valid from a technical point of view.

In sectors with multiple nodes, phones will detect different versions of the same cell, each transmitted from a different antenna, but will presumably treat them as if they were multipaths (copies of the same signal that travelled over different length paths to reach the phone at slightly different points in time) and will deal with it as part of normal signal processing. On the uplink, each phone will transmit just one signal, which will be separately received by the different nodes and fed back to the base station, which will, in turn, treat them as multipaths or 'echoes' of the same signal.

TA will be calculated from the delay associated with the earliest arriving of the multiple copies of the signal received at the base station. This means that TA continues to perform its primary purpose, of synchronising the transmission time of each active phone to ensure that traffic arrives at the expected point in time.

The implications for any distance inferred from those DAS TA measurements are less clear.

Additionally, there is a similar lack of clarity in terms of determining the point from which any TA inferred distance measurement should be mapped.

The TA measurement is taken in the base station, that's where the cellular radios are. We've been told that the base station is at the hub and we can assume that the hub is at the site listed in the tower list – but do we use the site lat/longs or the site address to identify the hub location? We have numerous examples where those two sets of coordinates point to different locations.

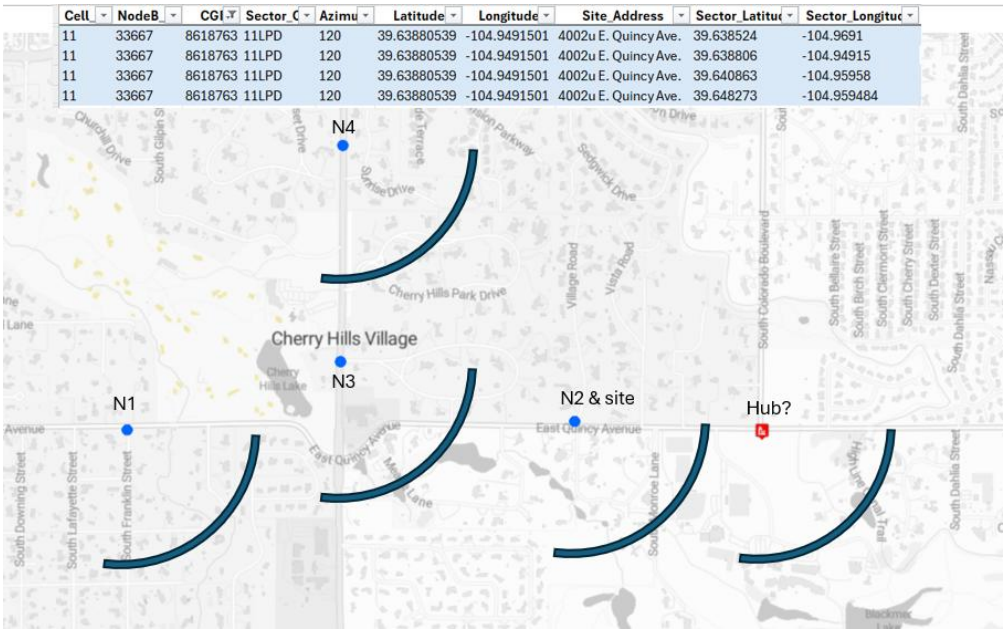

4.6 Recognising T-Mobile Outdoor DAS Systems

Understanding that the records you are reviewing are derived from an outdoor DAS system is important. Through conversations with T-Mobile engineers, the following current naming convention is used by T-Mobile in cell site lists to identify sector ID's, displayed as 'Sector_CD' in current T-Mobile cell site lists.

Using the sample Sector_CD value of **11LPD**, each character in the name maps to one specific type of data as follows:

Digit	Information Type	Sample	Result	Other Possible
1st digit	Sector Number	1	Sector 1	
2nd digit	Carrier	1	Carrier 1	
3rd digit	Technology	L	L for LTE	N= 5G NR; G= GSM; U= UMTS; T= NBIoT
4th digit	Band	P	P for PCS	A= LTE AWS; U= UMTS AWS; F= 5G AWS
5th digit	Site Type	D	D for DAS	A=outdoor macro; Y=indoor DAS; 1=mixed/special DAS; B=outdoor micro; Z=indoor micro; W=indoor pico; E=outdoor temporary

Where NBIOT = Narrow Band Internet of Things and AWS = Advanced Wireless Services band.

Other possibilities may exist, but the naming convention should stay the same. Check the relevant T-Mobile cell site list to determine what the characters mean if something other than the above is seen.

The most important part to know is that the last character of the cell site lists Sector_CD ending in D represents an outdoor DAS system. If uncertain as to what site type was being used, it is recommended to contact T-Mobile for confirmation before making evidential conclusions.

It's not guaranteed that all outdoor DAS sites will actually be marked as such in the tower lists – anecdotally, we've had reports of sites that are clearly deployed using outdoor DAS techniques, and that are associated with unusually large TA values, that are listed as 'outdoor micro' sites.

5. T-Mobile Outdoor DAS Issues

5.1 Key Issue – how do we plot this data?

As can be seen in the previous section, there are multiple potential issues associated with the interpretation of TA data when associated with outdoor DAS sites. All of these issues contribute to uncertainty about how best to plot TA arcs associated with that type of site.

Some of these uncertainties relate to determining the location of the elements involved in a TA transaction – where is the radio, which antenna was used, how far are those elements from each other?

5.2 Fibre Delay

The first uncertainty relates to the additional delay suffered by signals as they travel over fibre optic cables. In a DAS, fibre is employed between the radio and the antenna and so forms part of the delay measured by TA.

The difference in the speed of light through optical fibre compared to 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.

When a distance is inferred from TA, it is common to calculate that distance by multiplying the delay by the speed of light (c) – the value of c that is commonly used is the speed of light through a vacuum (299,792,458m/s), although the speed of light through air at sea level (299,702,547m/s) is also sometimes used. The speed of light through fibre is much lower (approx. 200,000,000m/s) and this speed disparity 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.

An example of the effect of the slower transit of light through fibre is shown in the diagram. The symbol µs represents microseconds.

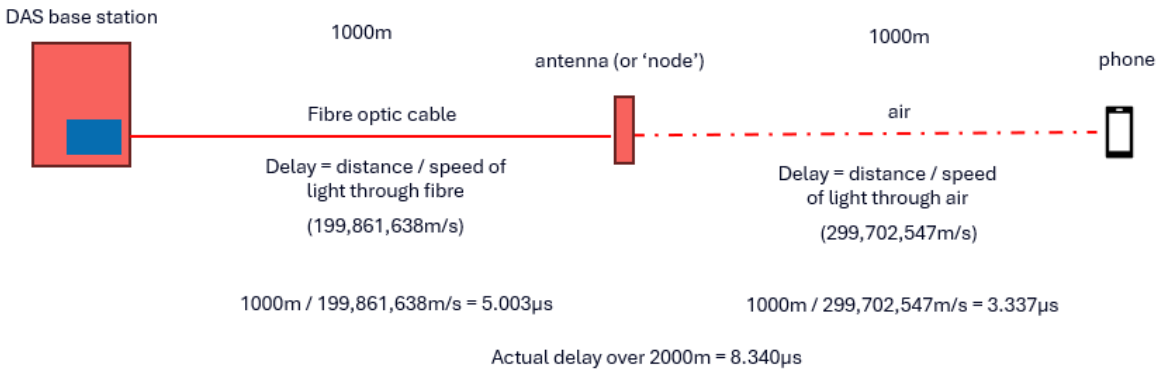

Imagine that Phone A is 1000m from a node antenna with a direct line of sight connection.

The node has a straight 1000m cable connection to the hub.

Phone A is therefore 2000m from the base station which would ideally be measured as

being within TA band 25.

The actual measured delay for the connection is 8.34µs. What can't be easily measured and isn't known is that the overall delay consists of two separate parts: delay through air and delay through fibre.

Light travels more slowly through 1000m of fibre (5.003µs) and much faster through 1000m of air (3.37µs), but both add up to the overall measured delay (8.34µs).

When inferring a distance from the TA value, it is common to multiply the delay by the speed of light through a vacuum - 8.34µs x *c* = 2500m – which suggests that the phone was further away from the base station than its actual distance of 2000m.

This means that the phone is reported as being in TA band 32 instead of TA band 25 – a 500m difference.

One way of compensating for the effect of the DAS fibre connection, is to compensate for the slower transit of the signal through fibre by multiplying the length of the fibre by x1.5 (which compensates for the 66% slower speed of light), then subtract the compensated length of the fibre from the inferred TA distance and then plot the remainder as the TA arc distance from the node.

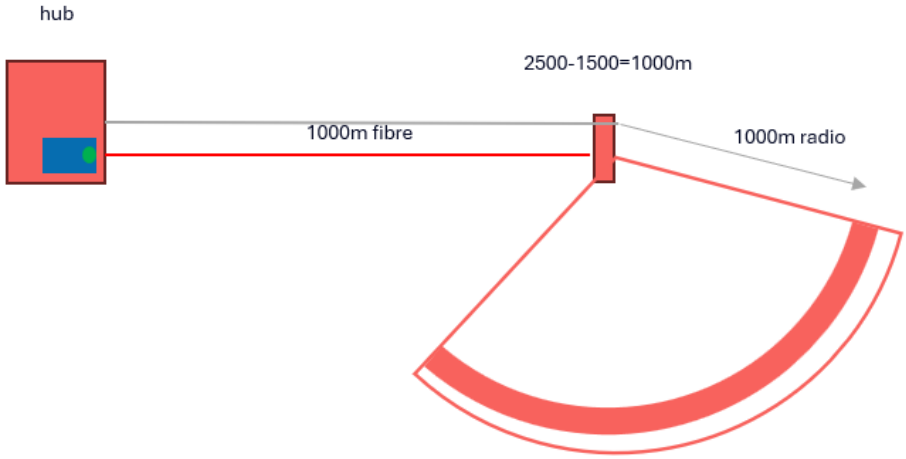

If the exact length of each fibre connection was known, it would be possible to compensate for the additional delay – by multiplying the exact fibre length by 1.5 before subtracting that value from the inferred TA distance.

So:

- if the inferred TA distance was 2500m
- the exact fibre length was 1000m
- multiply the fibre distance by 1.5
- 1000m x 1.5 = 1500m
- Subtract the compensated fibre distance from the inferred TA distance
- 2500 – 1500 = 1000m
- plot the phone's position based on the remainder – 1000m from the antenna node

Whilst this method makes mathematical sense, it introduces assumptions and opinions into the process of plotting a TA arc that many practitioners may feel uncomfortable with.

5.3 Digitisation Delay

In simple, traditional DAS systems, like those employed in shopping malls, the connection between the radio and the antenna is carried by a basic coaxial copper cable. The RF signal remains an analogue signal in this method, but coaxial cables have a limited transmission distance of just a couple of hundred metres, unless signal amplifiers are employed along the cables.

The longer distance connections required by some outdoor DAS systems requires a more complex transmission system, based on digital methods.

In digital DAS transmission, the radio in the hub produces an analogue RF signal, as usual, and that radio signal is converted to an optical signal so it can travel over a fibre optic connection. Fibre has a greater maximum transmission distance than coaxial cables, with even basic fibre cables theoretically being able to carry signals for up to 20km, without using signal amplifiers.

The analogue to digital (or RF to optical) conversion process creates what's known as an amplitude modulated (AM) light signal, where the amplitude of the light sent down the fibre varies in time with the variations in the radio signal.

Again, some very simple DAS systems simply carry the analogue AM light signal along the fibre to the antenna, but this has distance and efficiency limitations, so 'carrier grade' digital DAS systems employ an additional step that converts the AM optical signal to digital before transmitting it.

Analogue to digital conversion of this kind is usually achieved by 'sampling' the analogue signal, taking amplitude measurements very frequently (millions of times per second) and sending those measurements as strings of digital information (1s and 0s) along the fibre. At the receiving end, digital to analogue conversion is achieved by creating rapid pulses of RF energy that match the amplitude of each of the received samples – if the pulses are rapid enough the resulting signal is perceived to be a continuous analogue wave.

There is a common standard for this process called CPRI (Common Public Radio Interface) which was designed to support the digital DAS use case. CPRI was enhanced to create eCPRI in the late 2010s. CPRI samples the analogue light signal and sends those samples across a fibre connection; eCPRI takes the samples and packs them into Ethernet frames for the journey across the fibre. Ethernet allows a transmission link to be shared by more than one node and also allows for more complex management messaging to be exchanged between the hub and the nodes.

Several DAS manufacturers have developed their own proprietary solutions as well and it one of these that is apparently used for T-Mobile outdoor DAS.

The process of RF to optical conversion causes little or no delay to the transit of the signal, but the analogue to digital 'digitisation' process potentially does impose additional delay.

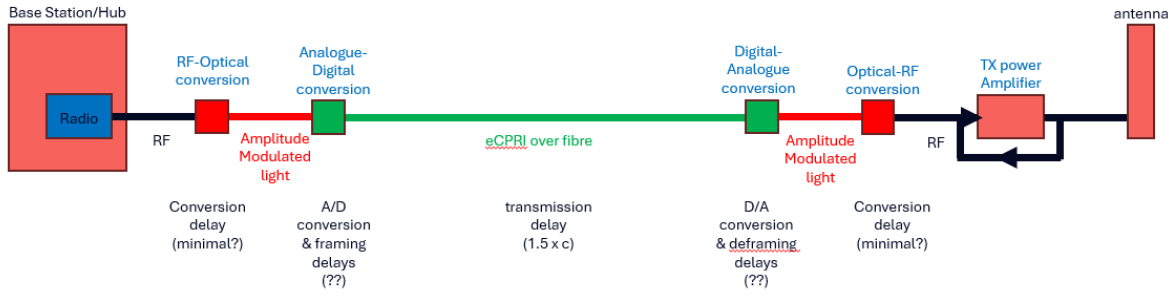

eCPRI, as an example, has managed delay classes of 50/100/200/500µs. Delays of this level would be imperceptible to the users of cellular connections that travelled over such a transmission link – humans don't tend to notice delay until it gets to the hundreds of milliseconds level (300ms or 1/3 of a second is the usual lower limit of delay we'd notice), so digitisation delays would have no effect on the quality of a DAS connection. TA is also able to deal with delays of this kind with no ill effects.

But even a 10µs delay would add to the delay measured by TA and would lead to an extra 1.5km being added to the inferred TA distance for a phone.

To fully understand the effect of TA via a DAS cell, it would be necessary to understand the additional delay added by any analogue-digital-analogue conversion process.

5.4 Shared Transmission Delays

In systems like eCPRI, once the RF signal has been digitised, the resulting digital information needs to be transmitted across a fibre connection. In very simple systems (like CPRI), the data is just sent in a simple stream over the cable, but more complex systems (like eCPRI) use an established Layer 2 transmission protocol like Ethernet to manage the flow of data over the fibre optic connection.

Ethernet is the principal technology employed for local area networks (or LANs) and in addition to managing the flow of data, it also manages the sharing of network connections by multiple devices.

If Ethernet was employed on the DAS fibre connections, it could allow each connection to act like a LAN and be shared by more than one node. This in turn could make the fibre transmission system more efficient, by allowing multiple nodes to share the same connection and not requiring a separate fibre to be run from the hub to each node.

If Ethernet is used to share the capacity of the fibre connections, this could introduce additional – and variable – delays, as each node may have to wait a few microseconds before it could transmit its next Ethernet frame full of digitised RF data if the connection was already in use by a different node.

The other forms of delay explored so far – the speed of light through fibre, digitisation and conversion – can all be expected to contribute fixed amounts of delay; if additional delay could be caused by congestion within a network connection, then that would introduce variable delay, meaning that the TA measured for a phone at a fixed location might vary (as the network delay varies), which in turn could cause the distance inferred from those TA measurements to also vary.

5.5 Multiple Nodes per Sector

As discussed in section 5.1, T-Mobile's version of outdoor DAS seems to follow a hub/sector/node configuration, where a central hub serves several sectors and each sector could support between 1 and 4 nodes.

In this model, if there are multiple nodes in the same sector, they each seem to broadcast the same cell (that cell is created in a single radio in the base station, but the signal is split, to be transmitted from multiple antennas), so the nodes in a sector can be thought of as separate antennas serving the same cell.

Examination of the Outdoor DAS sites mentioned in T-Mobile tower lists shows that the network consists of a mix of single node sectors and multi-node sectors spread around the network. Multi-node sectors seem to be more heavily deployed in some areas (Philadelphia, Denver) rather than others (South Florida).

A 4-node sector was described in section 5.3 and the layout of the elements in that group is shown again in the diagram below.

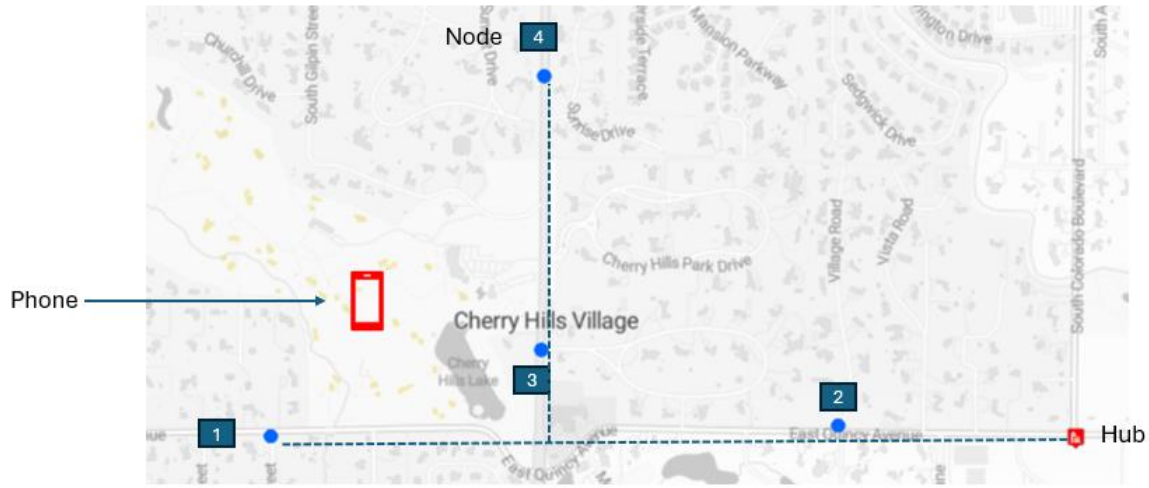

Let's imagine that a phone (in red) is situated in an area where it could connect to any of three nearby nodes (1, 3 or 4). It will be receiving copies of the same signal from all three nodes on the downlink, and those nodes will be receiving copies of the same uplink signal and passing it back to the base station (which is presumably in the hub).

The fibre connection back to the hub from each node has a different length and therefore a different delay profile.

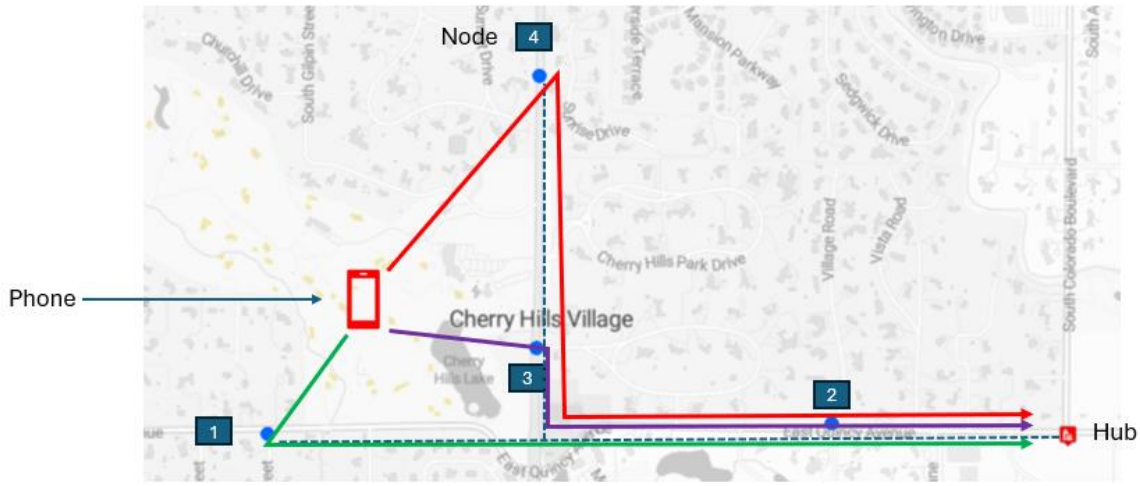

The base station will calculate TA from whichever is the dominant version of the signal from moment to moment. If the signal routed via node 1 is used to calculate TA, there will be a lower amount of delay than if the signal routed via node 4 was used.

The amount of delay suffered by the phone, and therefore the distance that would be inferred by TA for that phone, might vary moment by moment, as the base station processed the combined uplink signals it was receiving.

5.6 Issues & Mitigations

In an Outdoor DAS system, as we've seen, there may be a combination of elements contributing to the delay suffered by signals from phones served by those cells. Some of these causes of delay can be understood and quantified and may also be capable of being mitigated to ensure that TA data can be used reliably. Other causes, particularly those of a variable nature, may not be so easily mitigated.

Cause 1 – distance

The variable delay suffered by a phone will increase as the distance between it and the serving base station increases. This is mitigated by the standard working of the TA system, which increases or decreases the level of compensation that is applied as the delay varies.

Cause 2 - fibre connection

The fixed delay suffered by a phone will increase (compared to the delay suffered by a non-DAS connection) if the connection between the antenna and the radio travels over optical fibre, as light travels more slowly through fibre than through air. This can possibly be mitigated by determining the exact length of the fibre and multiplying that distance by 1.5 before subtracting it from the inferred TA distance – this may provide a more precise inferred distance for the phone from the serving antenna.

Cause 3 – Digitisation delay

The fixed delay suffered as the RF signal is converted to digital for its journey across a longer distance fibre connection. This can possibly be mitigated if the vendor of the DAS equipment is prepared to share figures for the delay imposed by their technique.

Cause 4 – shared transmission delays

The variable delay suffered when DAS transmission links are shared by multiple cells/nodes. This cannot be mitigated as it would be impossible to know how much delay each link suffered at specific moments in time.

Cause 5 – multi-node sectors

The variable delay suffered by a connection as the radio signal is transmitted by a set of geographically separate nodes that serve the same sector. This cannot be mitigated as it would not be possible to know which node the dominant signal travelled via from moment to moment.

5.7 Solutions & Mitigations

It seems apparent that some of the additional delays associated with outdoor DAS sites could be mitigated if additional information was made available – length of fibre cables, delays associated with the transmission systems employed on those cables, etc.

Some of these values are essentially unknowable, however, and the process required to calculate mitigations for the others would require assumptions to be made and might rely on expert opinion, which makes these methods unsuitable for some practitioners. The use of assumptions in the calculations would make the results of those calculations open to challenge.

6. Testing

6.1 Single Node DAS

Testing of a single node sector site in Palm Beach, Florida was undertaken in Spring 2025.

A node/antenna was selected that was situated in a low-rise suburban area and was connected to a hub/base station situated approximately 1km (0.6 miles) away.

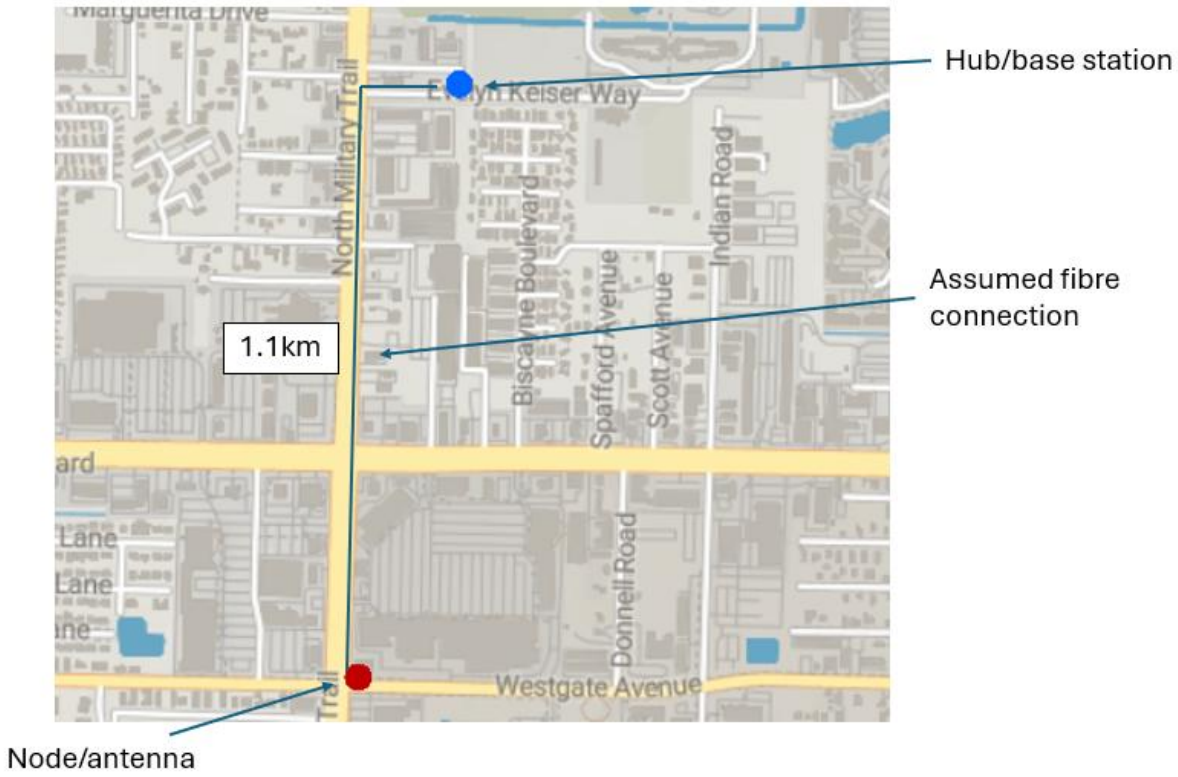

The exact length of the fibre connection between hub and node was not known.

Network TA data was not available for this test, so all TA values were derived from the survey device (a Lima Cell Monitor), so the results obtained should be regarded as provisional and subject to further testing and validation.

The GPS linked to the survey device provided details of its location that could be compared to the inferred distance provided by TA.

An illustration of the data captured during this test is presented below.

At a certain point in the test, the surveyor's vehicle was parked in a car park some 86m from the node/antenna and was stationary for several minutes.

The TA values captured from the survey device during this period fluctuated from TA Band 30 (2342m) to TA Band 35 (2732m).

timestamp	phone lat	phone long	TA	RF distance
02/05/2025 15:38:09	26.70288833	-80.11007167	33	86
02/05/2025 15:38:11	26.70288833	-80.11007167	34	86
02/05/2025 15:38:12	26.70288833	-80.11007167	34	86
02/05/2025 15:38:13	26.70288833	-80.11007167	35	86
02/05/2025 15:38:20	26.70288833	-80.11007167	34	86
02/05/2025 15:38:22	26.70288833	-80.11007167	34	86
02/05/2025 15:38:22	26.70288833	-80.11007167	35	86
02/05/2025 15:38:23	26.70288833	-80.11007167	33	86
02/05/2025 15:38:24	26.70288833	-80.11007167	35	86
02/05/2025 15:38:47	26.70288833	-80.11007167	33	86
02/05/2025 15:38:50	26.70288833	-80.11007167	34	86
02/05/2025 15:38:56	26.70288833	-80.11007167	30	86
02/05/2025 15:38:58	26.70288833	-80.11007167	35	86
02/05/2025 15:39:15	26.70288833	-80.11007167	34	86
02/05/2025 15:39:18	26.70288833	-80.11007167	35	86
02/05/2025 15:39:37	26.70288833	-80.11007167	34	86
02/05/2025 15:39:40	26.70288833	-80.11007167	35	86
02/05/2025 15:39:58	26.70288833	-80.11007167	34	86
02/05/2025 15:39:59	26.70288833	-80.11007167	34	86
02/05/2025 15:40:01	26.70288833	-80.11007167	35	86
02/05/2025 15:40:04	26.70288833	-80.11007167	35	86

The reasons for this level of fluctuation, which covers some 500m (0.3 miles) are difficult to explain.

Our expectation was that, if we knew the straight-line distance between the survey device and the node/antenna, we could subtract that distance from the overall inferred TA distance and get a reasonably stable estimate of the delay introduced by the fibre connection.

This did not turn out to be the case – with the distance between the device and the node remaining static, the result of subtracting that distance from the inferred TA distance provided to be highly variable.

A potential reason for this is variable delay over a shared fibre connection – traffic being sent by other nodes over a shared fibre was causing traffic from the tested node to have to wait a variable amount of time to access the connection. Additional delay would have been caused by the digitisation process and the slower transit speed of light through fibre.

The result of this testing was that a TA arc could not be reliably plotted for the test device.

6.2 Multi Node DAS

Separate testing of a DAS deployment in Denver, Colorado was undertaken, also in Spring 2025.

This deployment appeared to use the 'multi-node' model, as the same cell ID appeared to be linked to four different node/antenna locations.

Cell_	NodeB_	CGI	Sector_C	Azimu	Latitude	Longitude	Site_Address	Sector_Latitu	Sector_Longitu
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.638524	-104.9691
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.638806	-104.94915
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.640863	-104.95958
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.648273	-104.959484

The relationship between these locations can be mapped as follows:

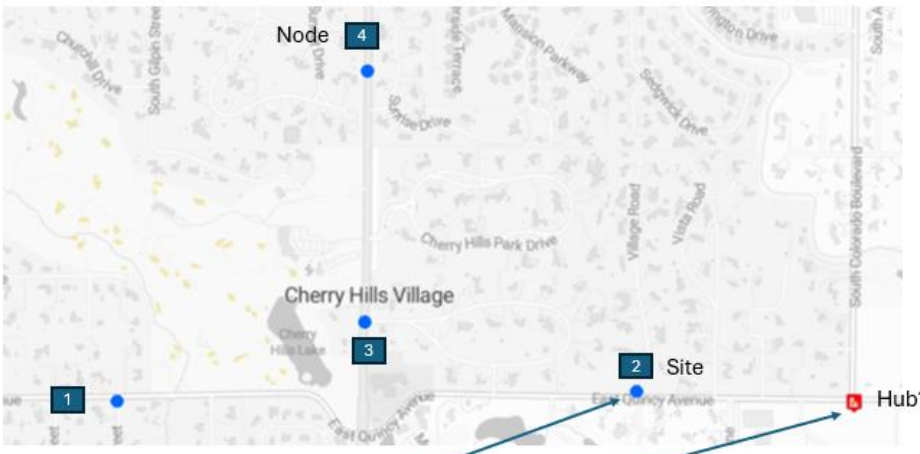

Cell_	NodeB_	CGI	Sector_C	Azimu	Latitude	Longitude	Site_Address	Sector_Latitu	Sector_Longitu	
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.638524	-104.9691	1 Node
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.638806	-104.94915	2 Node & Site
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.640863	-104.95958	3 Node
11	33667	8618763	11LPD	120	39.63880539	-104.9491501	4002u E. Quincy Ave.	39.648273	-104.959484	4 Node

RF surveys captured in the area showed that three of the four nodes appeared to be broadcasting signals for the same cell ID – 310-260-8618763.

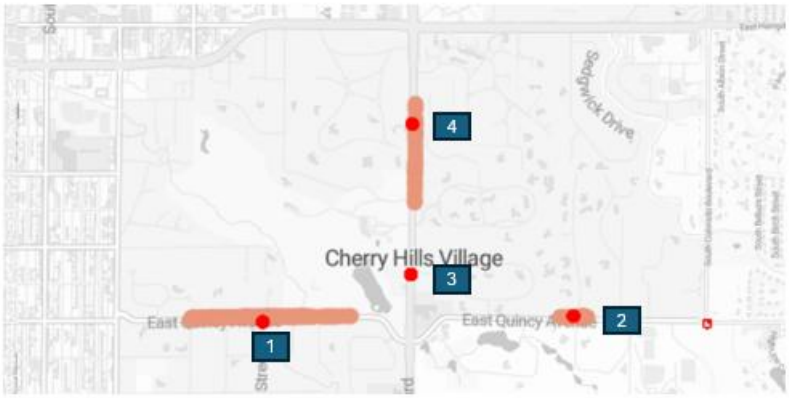

The orange areas on the map above show where that cell ID was detected as providing serving coverage during the RF survey.

This indicates that the antennas are more likely to be providing omnidirectional coverage, rather than the 120° sectorised coverage mentioned in the tower list and it also clearly shows that each node/antenna (apart from node 3) is generating its own separate area of coverage for the same cell ID.

The TA values captured during this survey – both those captured by the Lima survey device and those generated by the network for the test phone – showed a large degree of variability. The lowest TA band captured in the network data was 9 (702-780m or 0.44-0.49 miles) and the highest TA band was 304 (23.73-23.81km or 14.74-14.79 miles).

As there was no way of determining which of the three active nodes/antennas (only 1, 2 and 4 appeared in the RF survey data) carried the signal from which the TA was

measured, we calculated the straight line distances between the Lima's GPS location and each node to see if there were any correlating measurements. Confusingly, node 3 came out as the node with the closest distance correlations, but even for that node, the minimum difference to the network TA value was 2 bands and the maximum was 283 bands.

Instead, we analysed the data in relation to which node the Lima was closest to at the time of each network TA event. We estimated the likely length of the fibre connection between the closest node and the hub and added that distance to the measured Lima-node distance derived from the Lima GPS. Even when the estimated fibre length was multiplied by 1.5, to compensate for the slower speed through fibre, the majority of TA values still did not correlate with the estimated distance to the hub.

It is difficult to draw any definite conclusions from such variable data; suffice it to say that TA data obtained from multi-node DAS sites does not appear to be precise enough to use as evidence of the geolocation of the associated cellular device.

6.3 Subway/Metro DAS

Further testing is underway on both the New York City Subway and the London Underground network, both of which employ DAS to offer coverage in stations and tunnels.

The results of these tests will be added to a future version of this report.

6.4 Testing on Other Networks

In the US, AT&T and Verizon both seem to deploy sites using the Outdoor DAS model. Testing of DAS deployments for these networks is also planned and results will be shared once they are available.

Contacts in Verizon have suggested that there is no way to determine of one of their sites is DAS, from examining the tower lists.

AT&T tower lists do appear to distinguish between site types, with a 'build_type' of WIFI/DAS potentially being used to indicate DAS – whether this relates to indoor DAS, outdoor DAS or both is not yet fully understood.

7. Interim Recommendations

Only a limited amount of testing has been undertaken on this topic so far – a single-node DAS site in Florida and a multi-node site in Denver – but combined with the anecdotal evidence of unusual TA results obtained from other DAS sites, we feel confident in expressing an interim recommendation.

TA data obtained via DAS sites of all kinds – indoor or outdoor – appears to be subject to additional sources of delay. Due to the shorter cable connections typically employed by indoor DAS systems, it is likely that these additional delays will be fairly negligible. The longer cable connections employed by some outdoor DAS systems, plus the other causes of fixed and variable delay, mean that these differences can be significant.

Anecdotally, we hear reports of TA values fluctuation by several miles within a few seconds as a phone is passed between a traditional macro site and a DAS site. We have seen evidence of similar behaviour in our limited testing. We have also seen evidence of the TA values and their inferred distances fluctuating markedly in very short time frames, even when the phone stays connected to just one cell and when the phone is known not to be moving.

Our interim recommendation is that TA obtained via DAS sites should be treated with a great deal of caution, unless it can be corroborated with TA from traditional sites or with other types of geolocation data.

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