Terragraph Mesh

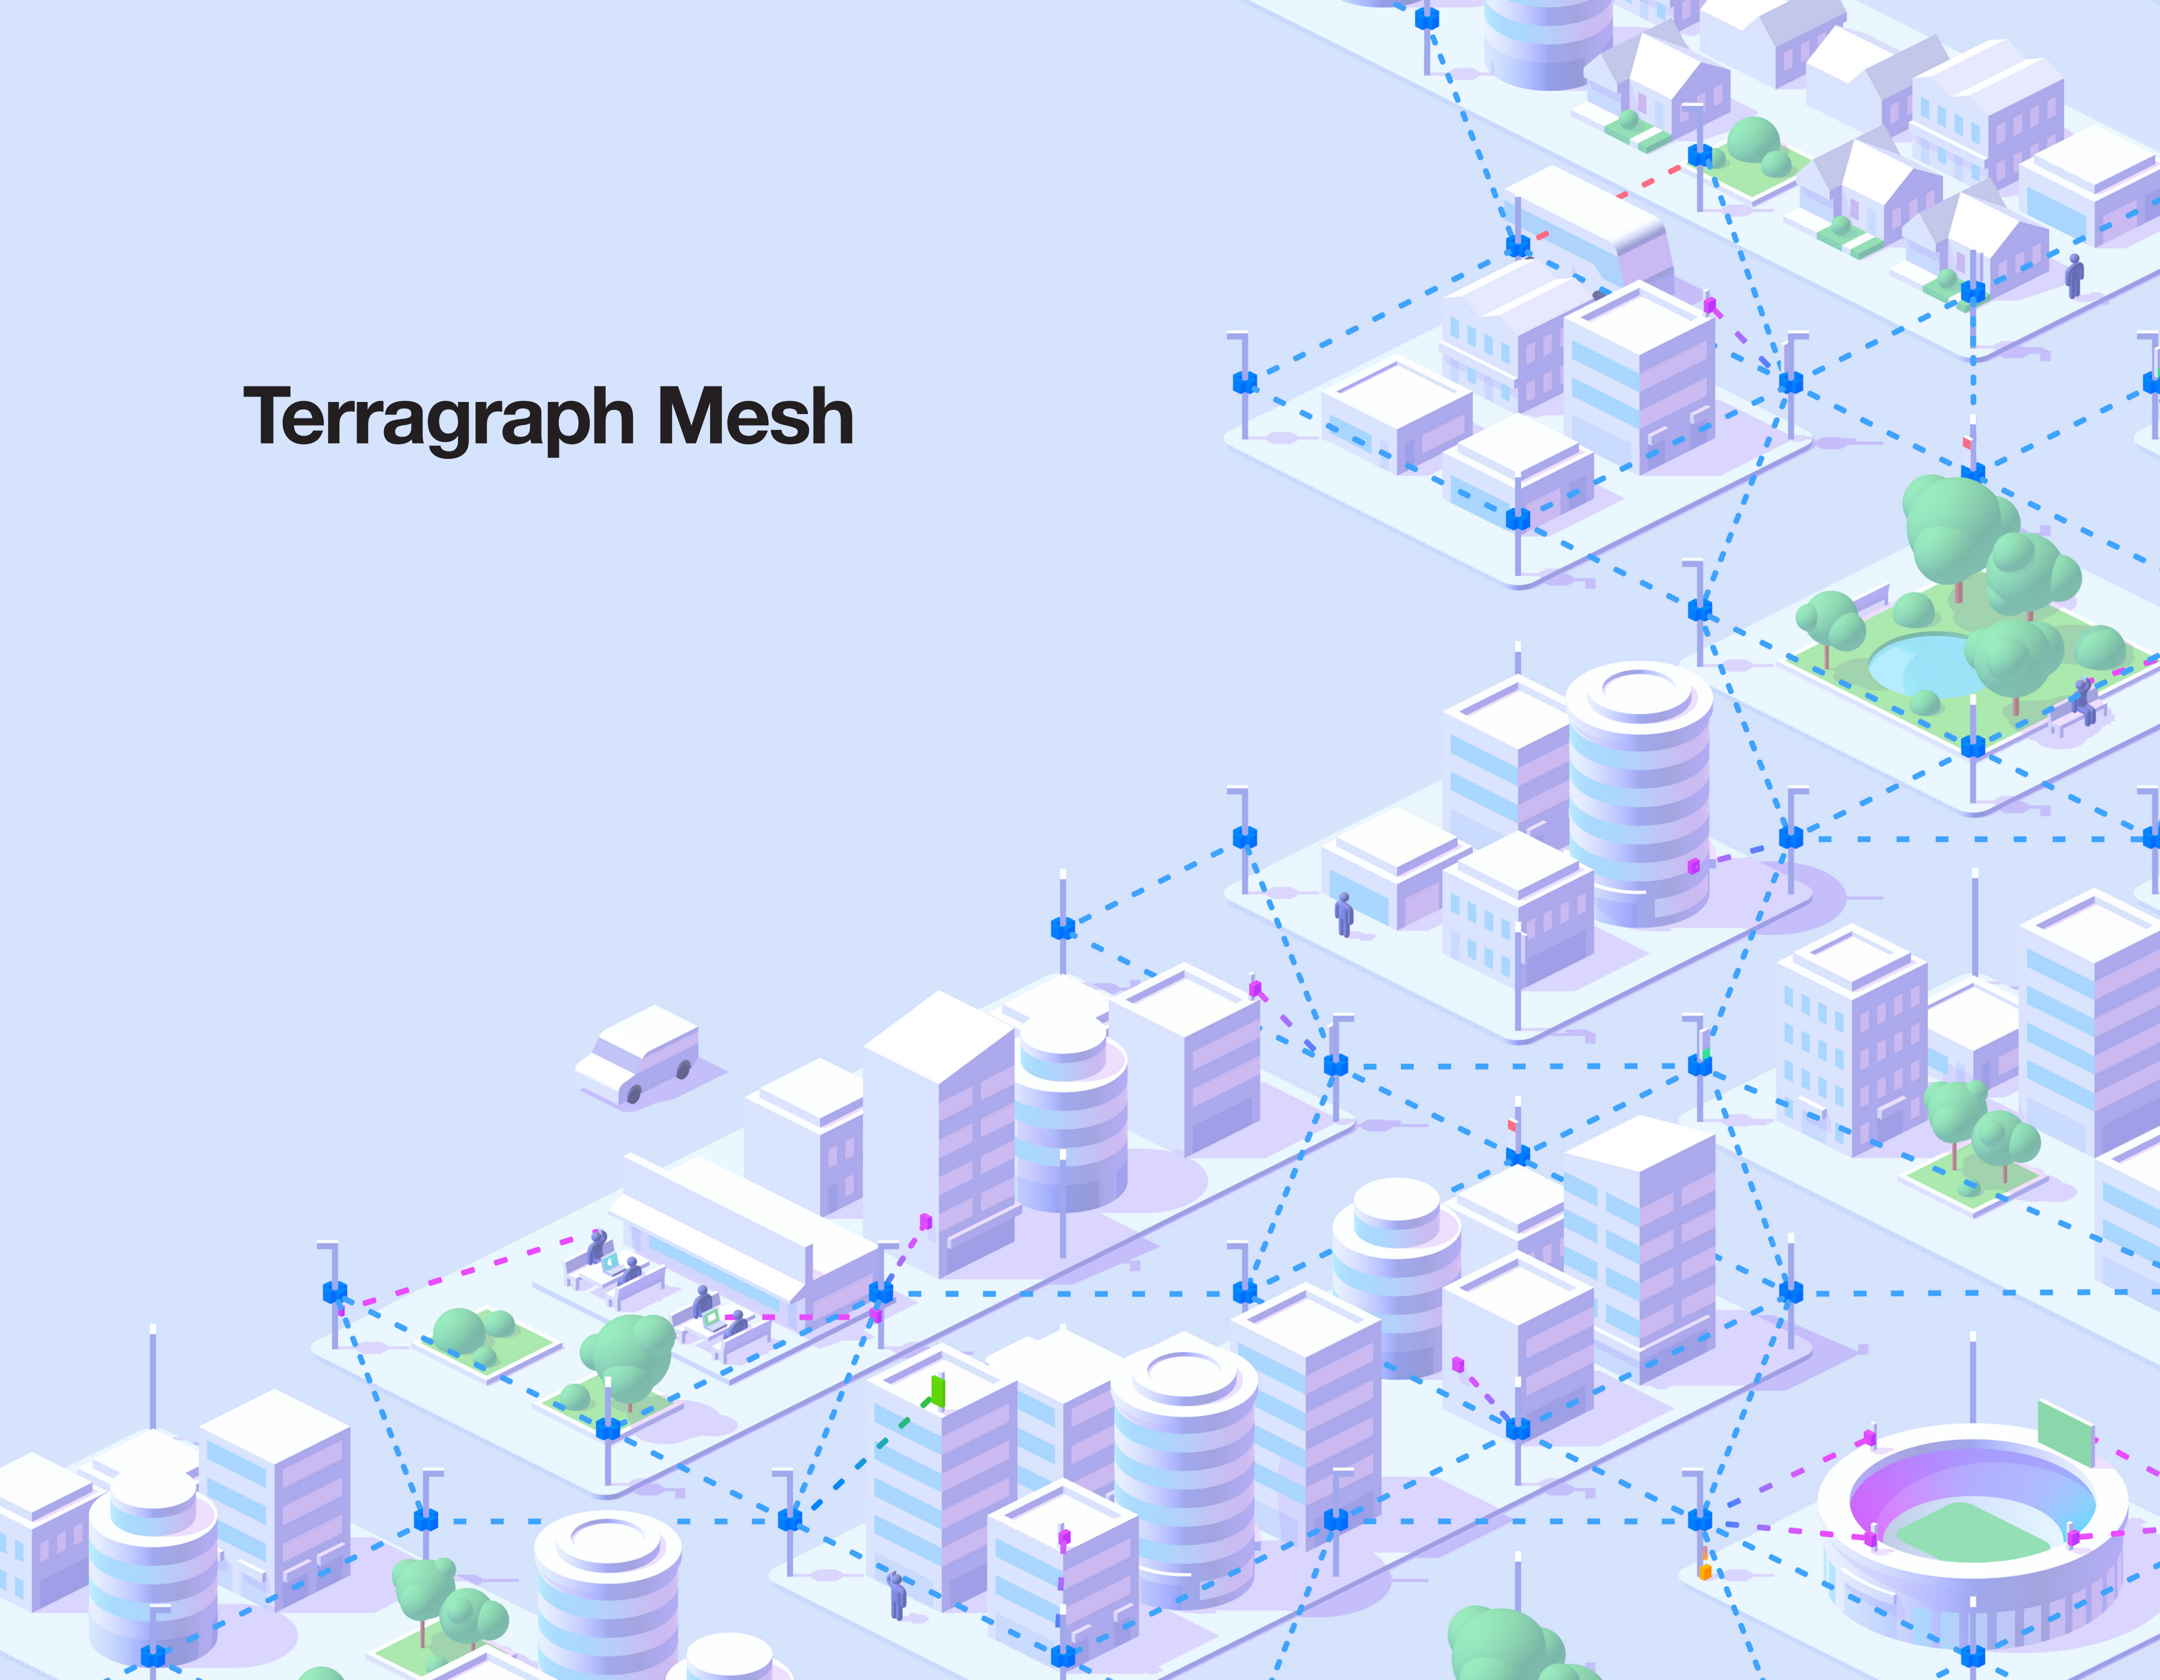

Scope

In this white paper we present the case for the performance and economic benefits offered by 60GHz Terragraph mesh design over a 60GHz point to multipoint (P2MP) design, in meeting the high network availability requirements of a fixed wireless access service provider.

What is Terragraph?

Terragraph is a multi-node, wireless, mesh technology designed for delivering a reliable and high-speed internet. It uses the wide unlicensed (or lightly licensed) spectrum available in 60 GHz frequency band to deliver gigabit speeds. Terragraph uses a unique mesh solution built upon OpenR open source standard to achieve high network reliability. Terragraph solution is much much lower in costs compared to competing solutions and is very fast to deploy, it can be brought to market in a matter of weeks.

Unlicensed Spectrum
60 GHz

Low Costs
Fast TTM

High Speeds

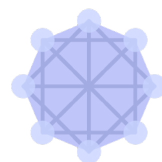

Resilient Mesh

Facebook's mission around Terragraph

Fast and reliable access to the internet is absolutely vital to our digital economy as new applications in education, medicine, transportation and entertainment depend on it. One of the biggest challenges for delivering broadband connectivity is last mile connection. There are more than a billion people still reliant on outdated non-scalable infrastructure to access the internet. FTTH has been steadily expanding in footprint but in most parts of the world it is prohibitively expensive and slow to deploy as they need permitting, trenching, laying cables etc. This in turn makes it too expensive for the end consumer.

Facebook connectivity is precisely trying to address this gap between demand and supply and Terragraph is a technology developed by Facebook for this purpose. With Terragraph, Facebook's mission is to not become a service provider or an OEM. Our main goal is to work with current industry partners such as service providers, OEMs, system integrators etc. to enable an sustainable and competitive ecosystem around Terragraph innovation. Facebook's work around Terragraph span across a) Technology innovation and licensing b) spectrum advocacy c) Ecosystem engagement.

Key operator benefits delivered by Terragraph

Terragraph leverages the existing IEEE 802.11ad standard and improves upon it to help build a more dense, reliable, scalable and easily deploy 60 GHz network. Terragraph's unique contributions and enhancements have been standardized and now a part of IEEE 802.11ay standard. The diagram on the next page maps the technical features Terragraph relies on and the operator benefits it delivers.

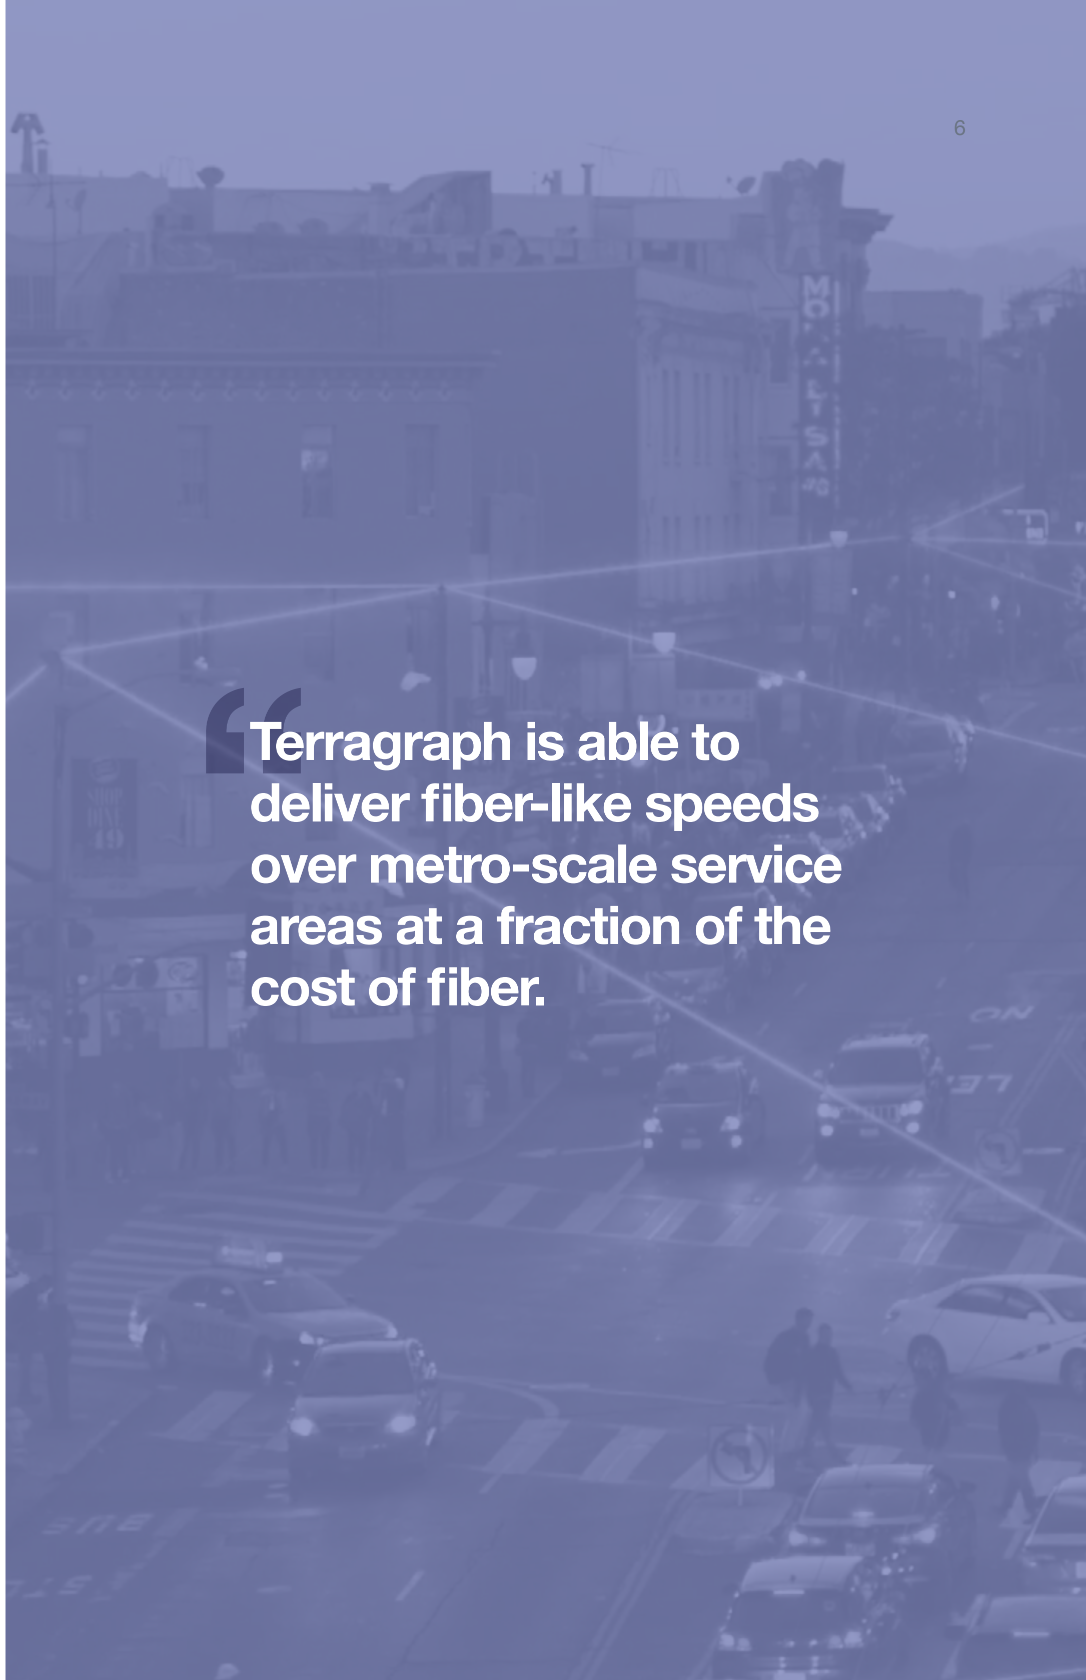

> **"Terragraph is able to deliver fiber-like speeds over metro-scale service areas at a fraction of the cost of fiber.**

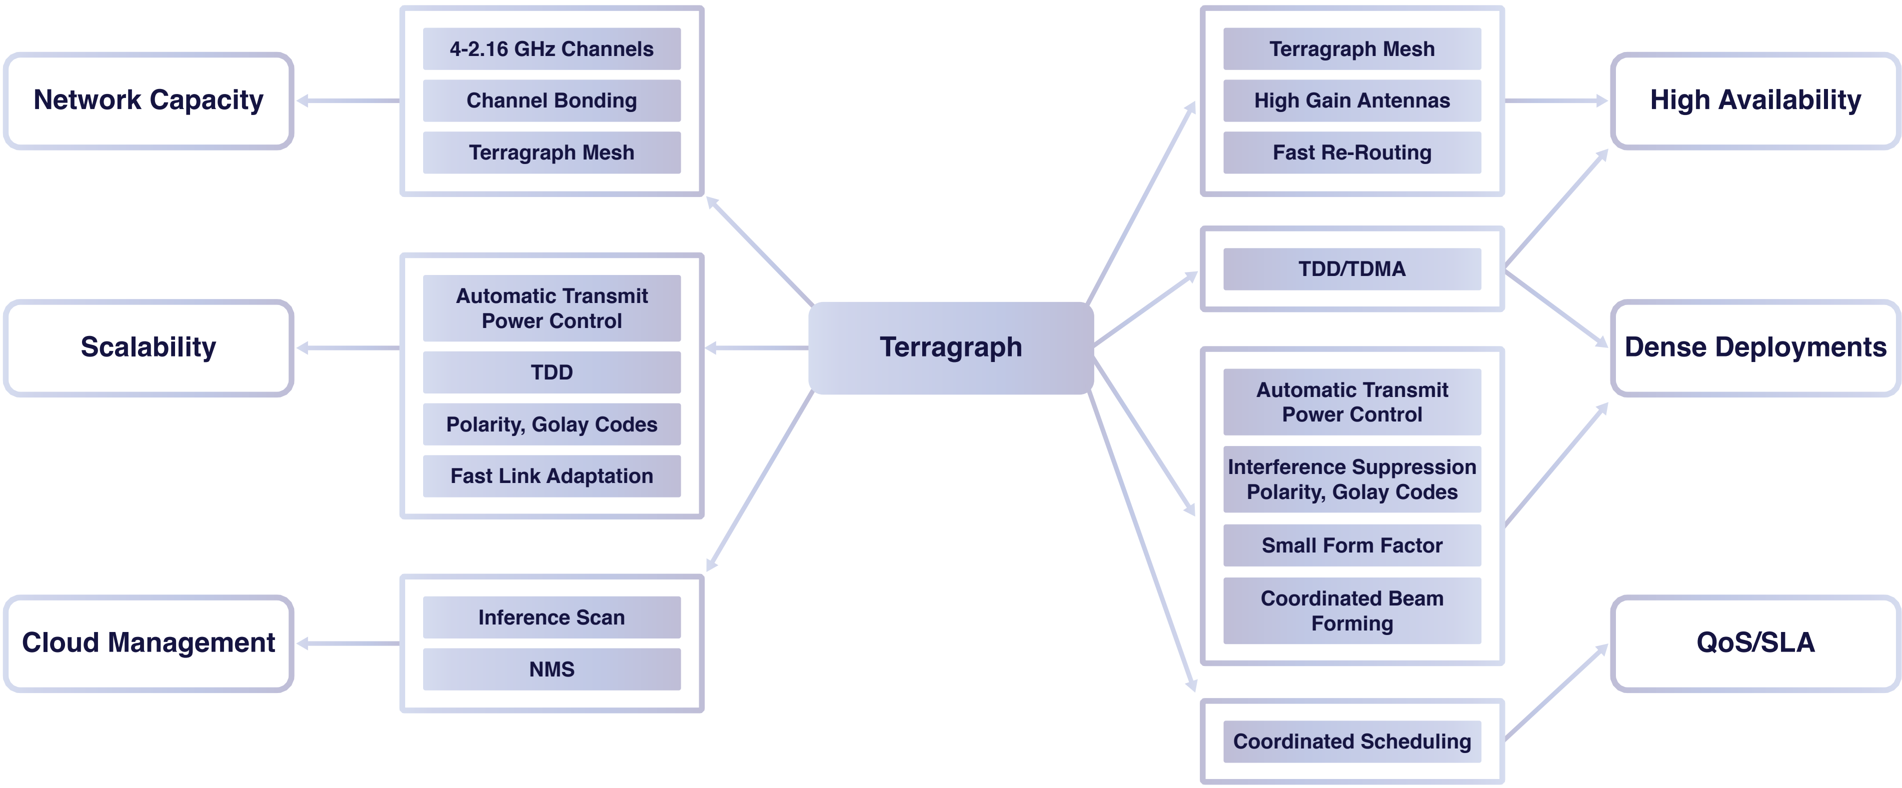

60GHz spectrum communication needs line of sight (LOS) and as a consequence is more prone to link outages. Terragraph mesh was developed to workaround this challenge and to deliver high network level availability (e.g. > 99.9%). In the subsequent sections, we will describe how the Terragraph mesh network delivers high availability in a more robust and economical way than existing P2MP solutions.

In this white paper, we are primarily focused on Terragraph mesh's ability to deliver a fixed wireless access (FWA) network with high network availability at a cost significantly lower than that of legacy P2P/P2MP FWA networks.

Terragraph Network Modeling

Site survey

The network chosen for this analysis is in an Urban city with a dense distribution of multi-story multi-dwelling units and small to midsize businesses. The network is designed to provide coverage to all 1200 buildings in the area. This densely built city makes it difficult to deploy fiber due to cost and disruption considerations.

This network consists of 115 poles with 4 POP locations spread over 1sq km area. Existing 6m utility poles were used for this deployment.
The network shown below uses existing poles to form a Mesh network serving the city.

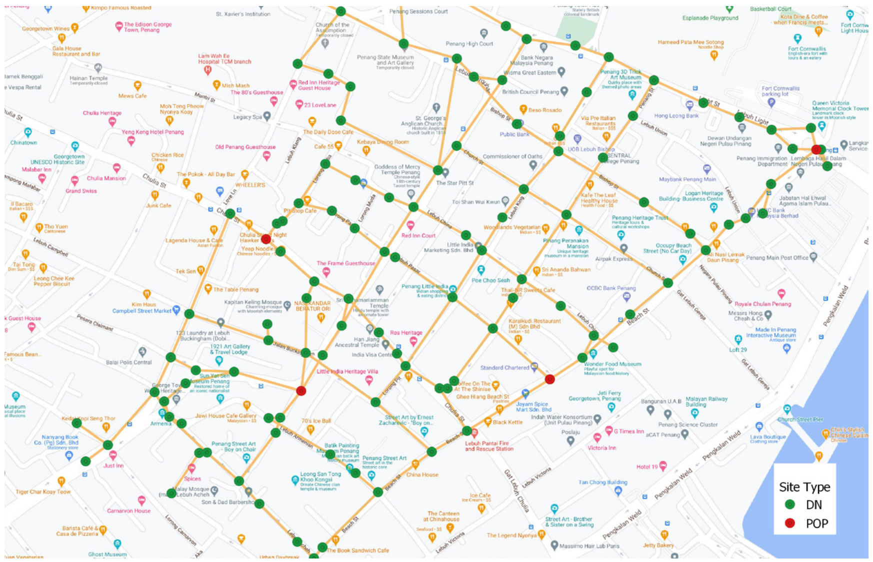

Network Planning Details

The network topology modelled consists of the following:

1. Fiber Points of Presence (PoP) - The network consists of multiple Fiber POPs and the physical location of these sites is carefully selected for physical diversity and at the same time to break up coverage in similar sized groups.
2. Distribution nodes (DN) - DN sites can be mounted on building rooftop and street poles. The DN site locations in this case study are all mounted on Street poles.
3. Consumer nodes mounted on building rooftop or on side walls to provide Broadband connectivity.

Link Level Availability Modeling

For this simulation we have assumed the DN specification to be the same as Puma (Facebook's proof of concept DN). The simulation parameters are captured in the table below.

Transmit Power	11 dBm
Antenna Gain	30dB
EIRP	41dBm
Peak Data Rate (MCS12)	3.6 Gbps (Bi-Directional agg.) 1.8 Gbps DL/1.8 Gbps UL
Oxygen Loss	15 dB per Km
Rain Fade attenuation	7 dB

The chart below captures the relationship between downlink speeds and link distance for different link availability and rain fade conditions.

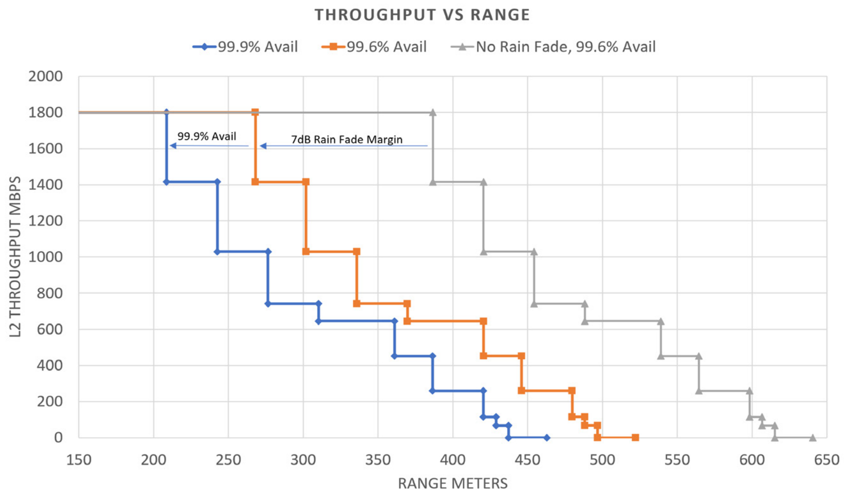

In field deployment link level availability varies on the above parameters and is measured by monitoring MTTF and MTTR.

For this analysis we have assumed Mean time to Failure (MTTF) of 720mins & Mean Time to Repair of (MTTR) 3 mins. This translates to link availability of 99.6% . This is a conservative number for regions with heavy rain fades. Most commercial equipment will exceed these targets.

The curve for 99.6% Link availability shows that we can expect to get >250m range on DN-DN links at MCS12. For links closer to 200m we can expect availability of 99.9%.

To model a worse case scenario with longer links, we have assumed that each link on this network has an independent availability of 99.6%.

Mesh and Tree design and availability modeling

Our baseline network is assumed to be a well connected Mesh with 4 POP locations shown in the image on the previous page.

Network availability simulation modeling

We run a Monte carlo simulation for a simulated duration of 10 Years to capture a sufficient number of individual link failures and their impact on overall network availability. The probability of link failure and link recovery in the simulation are modeled via an exponential distribution with mean values defined by MTTF and MTTR.

Each minute in the 10 Year duration 365 * 24 * 60 * 10 = 5.25 million mins is simulated with independent link failure based on MMTF and link restoration based on MTTR. Whenever a link fails the network attempts to reconnect using an alternate route before the next minute. In case no alternate route is possible, the DNs depended on that link for connectivity are identified and recorded.

At the end of the simulation we calculate the % of time each DN had connectivity to POP. This represents node level availability. Network availability is an average of this node level availability.

The simulation duration of 10 years was chosen to allow the results to converge to 4 places of decimal. The 10 year simulation duration should approach the lifetime of the product.

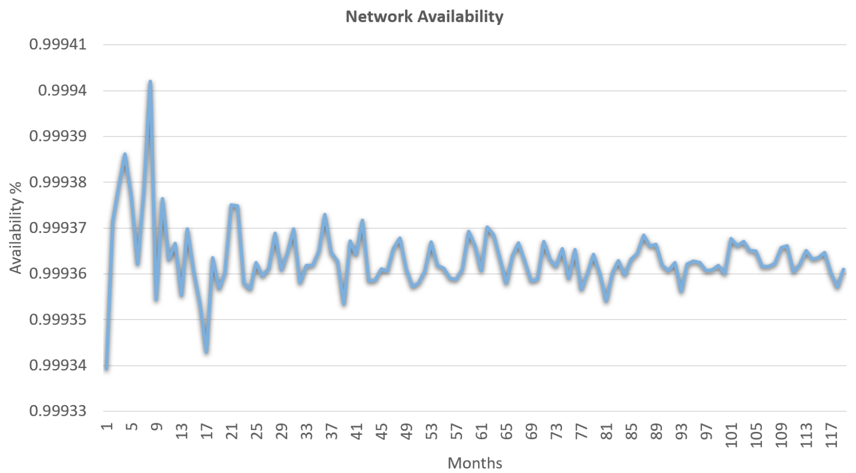

The chart above shows network availability results with a simulation run from 1 - 120 months. At approximately 10 Years of simulation duration the results start converging to 4 places of decimal.

Since this analysis focused on Network Availability contrasting a Mesh vs Tree deployment scenario, We only included POP and DN nodes in our analysis. Each DN node was assumed to be required.

Mesh Design

A Mesh network provides most DN with multiple routes to POP's. Therefore in the case of random link failure there are alternative routes that are available resulting in no loss of connectivity for the end user.

Mesh networks are designed with more redundant network paths than what is needed and hence might require more sectors per site to provide alternate links. The alternate links under normal circumstances do not carry any traffic and are available only as a backup.

In a Mesh network the availability of a Mesh Node will always be greater than link availability allowing operators to meet their network availability targets with well designed network layouts.

Operators can also augment the network with additional POP locations increasing overall capacity available and reducing the number of hops. This would also result in higher network availability.

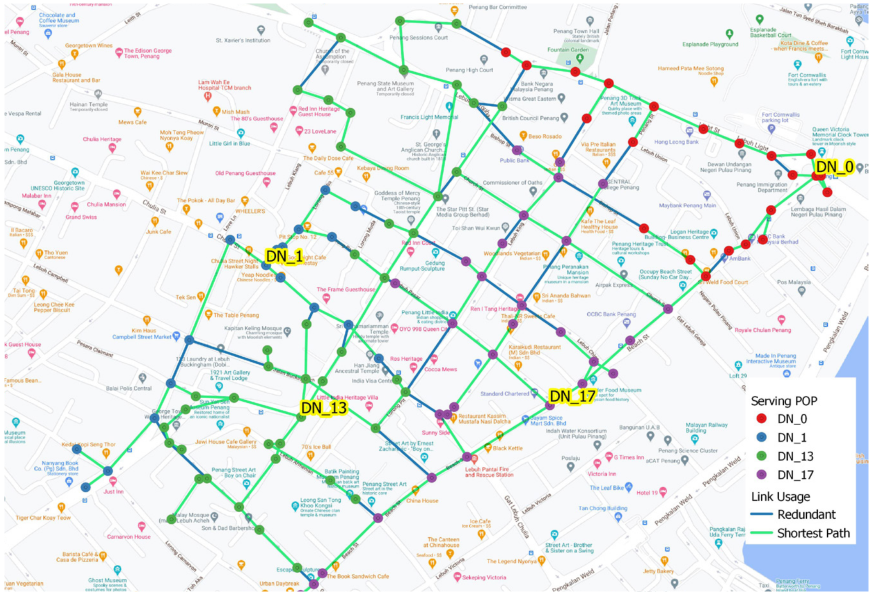

The plot above shows DN's with four different colors based on the serving POP. Each of the clusters operate independently while all the shortest path links are available. The moment a Shortest Path link fails, The network will start routing traffic to the nearest POP using the Redundant links. In case of a failure of a POP, the traffic in a cluster will get routed to other available POP's.

Tree Design

A Tree design is focused on minimizing cost and uses the minimum number of sectors possible and provides each DN with only 1 path to POP. In case of a link or a POP failure all the nodes downstream lose connectivity till the link is restored.

In a Tree network the availability of a Mesh Node will always be lower than link availability. To improve node availability the only option operators have is to deploy equipment which offers higher link availability. This is usually achieved by buying higher cost equipment, using lower inter site distance or operating at a lower guaranteed data rate.

Operators can add additional POP locations to increase network availability by reducing the number of links each node needs to traverse. This would also result in higher overall capacity available in the network.

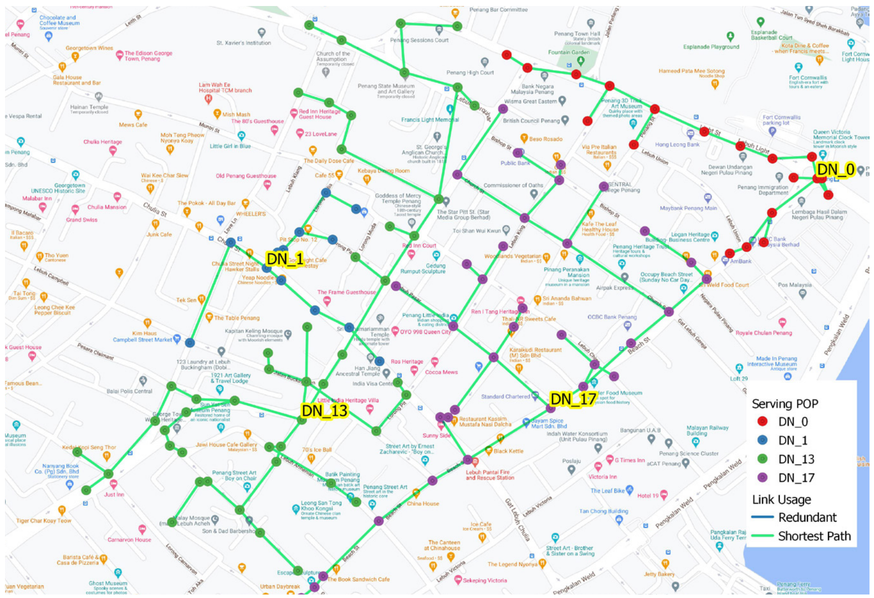

In the plot on previous page there is only one route from DN's to POP with no redundant links. The four different colors are based on the serving POP. The moment a Shortest Path link fails, The DN's on that chain will lose access to POP. In case of a failure of a POP, all DN in the cluster will lose connectivity.

Transitioning from Mesh to Tree Design

A fully connected Mesh and a Tree design represent opposite ends of deployment architectures. An operator does not need to choose one vs. another as it's possible to seamlessly transition from a Mesh to Tree design.

An operator should initially plan a Mesh design as that would allow them to benchmark network availability and cost. Once the Mesh network plan is ready, they can gradually eliminate redundant links to meet their Budget & Availability targets.

Any operators network choices are constrained by their Budget & Availability targets. Terragraph architecture allows operators to vary the number of Nodes, Links, POP location to design a network within these constraints.

Impact of Varying POP

Along with varying the number of links to transition from a Mesh to Tree, service providers also have the flexibility to vary POP count and locations.

In this study for each Mesh & Tree architecture we test the network availability with a range of POP locations. As expected, greater number of POP locations increases network reliability but at a cost.

The chart on the next page shows the impact on mesh & tree availability as the number of POP's goes from 1 to 8.

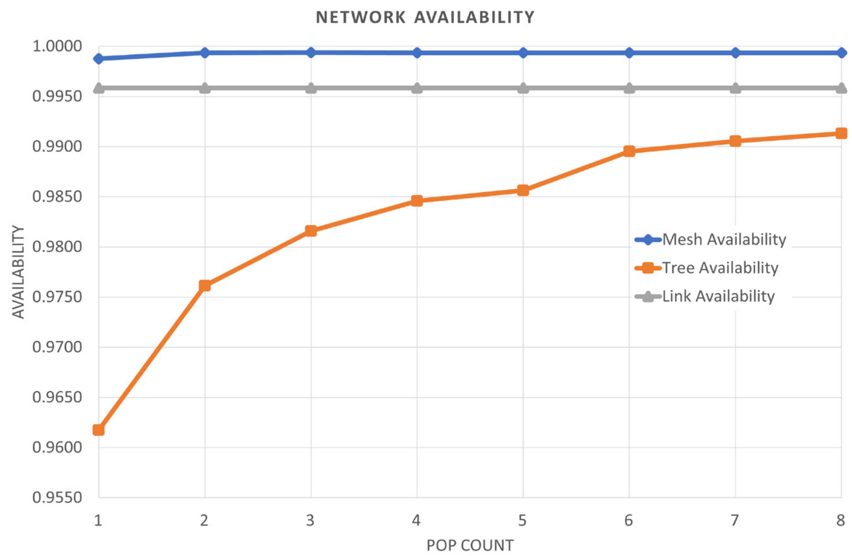

As seen above, in a Terragraph mesh network the mean network availability stabilizes at 0.9994 with just 2 POPs. With 2 POP's a Tree network offers 0.9761 network availability.

With a Mesh architecture operators can design networks to exceed the link availability offered by their equipment by building in redundant paths.

However, a tree network's mean availability is not able to go beyond 0.991 even when the number of PoPs is increased to 8. Therefore for an operator targeting a mean network availability of 0.999 or higher a Terragraph mesh becomes essential.

A significantly greater benefit of Terragraph mesh is observed on an individual node basis. The node availability plots between Terragraph mesh and tree design for number of PoPs increased from 1 to 8 is shown on the next page.

Node Availability Distribution and Mean Network availability

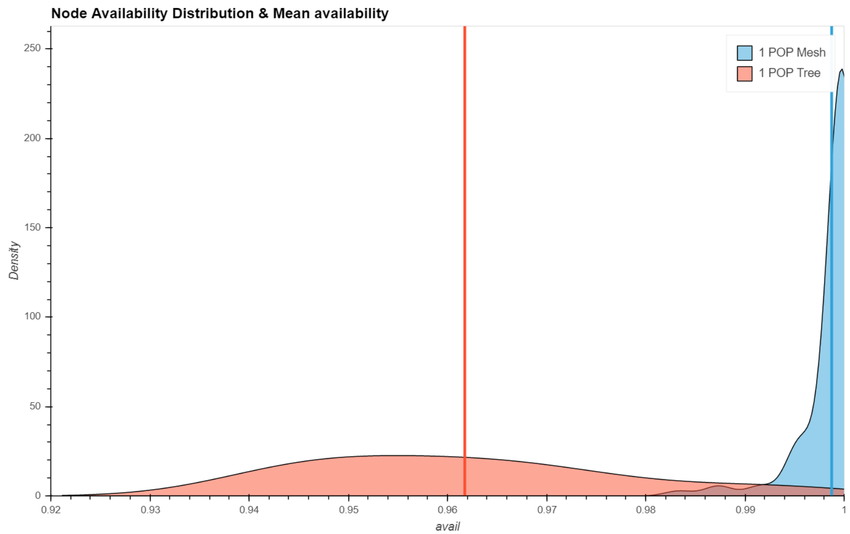

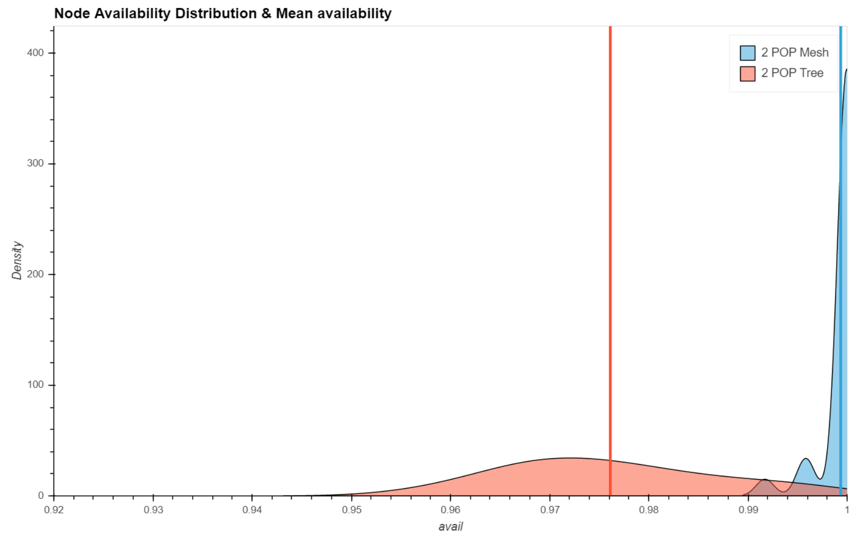

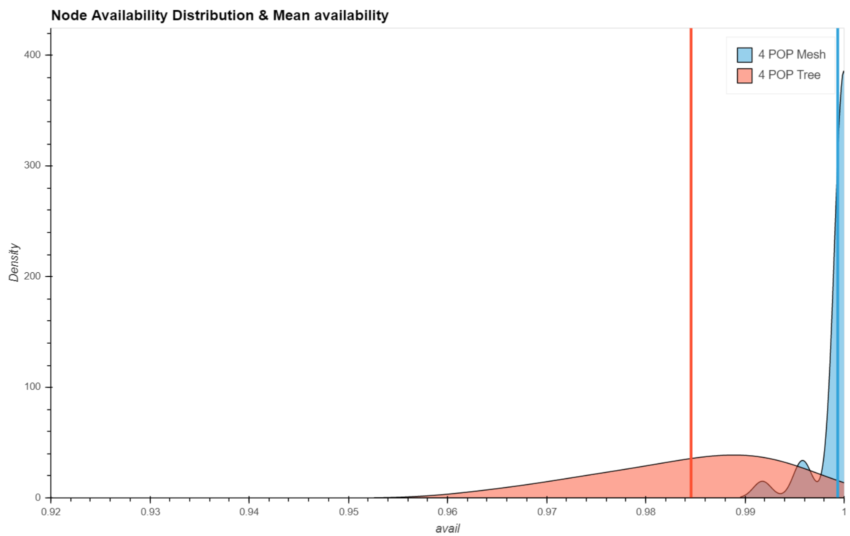

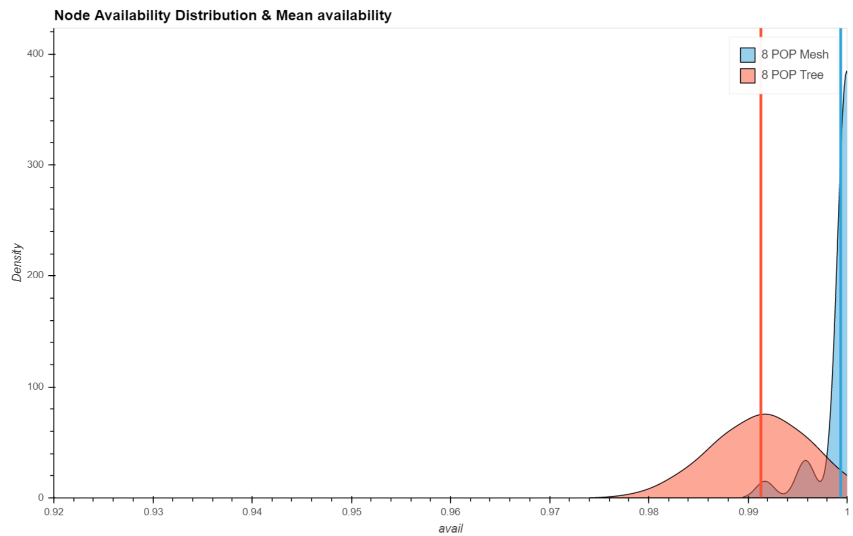

Observations:

1. Even a single PoP configuration network using Terragraph mesh is able to deliver a mean node level availability close to 99.9% with a very tight variance. However a tree network has a much wider availability spread across its nodes and individual node availability can be as low as 92%.
2. Increasing the no. of PoP sites to 8 improves tree network node level availability but the variance is still too large and the vast majority of the nodes have lower than 99% availability.

Conclusion

Existing 60 GHz P2P and P2MP solutions provide adequate equipment availability but cannot control for external factors. This limits availability of a daisy chain network to the product of each link's availability and hence makes it infeasible to design a 60 GHz network that exceeds the individual link level availability.

Terragraph mesh has been designed to precisely combat this limitation. It provides superior network availability by designing a mesh with an optimum number of redundant links needed in case of individual link failures with minimal impact on individual node availability.

This design trade-off becomes vital for fixed wireless service providers using this technology to meet wireline reliability of 99.9% or higher.

terragraph

learn more at terragraph.com