

Why geographic diversity of OGS sites is essential: proof from Southeast Asia



Figure 1: Miratlas LaserCom Pathfinder, Cloud Cover and Turbulence Analysis

The structure of the global Internet and its infrastructure has evolved considerably in recent years, with the addition of an extra layer at both the transport and access levels in low Earth orbit. Long marginal because it was confined to geostationary orbit — with everything that implies in terms of throughput and, above all, latency — the space-based network layer has become a major player in the global telecommunications infrastructure. It brings with it both ubiquitous access with near-global coverage and a direct access point to the fixed user (and soon to mobile users), but also a transport layer that is highly performant and extremely secure.

Interconnecting the constellations with the ground

Nevertheless, this network running parallel to the terrestrial network cannot, for the time being, compete head-on with the latter, particularly in terms of throughput. A first step is being taken with the interconnection of satellites through optical links, which offer performance comparable to that of optical fibre on the ground. Moreover, the low Earth orbit network provides an extremely dense mesh thanks to its thousands of satellites. The technologies used for the links between satellites belong to the same field and rely on the same foundations as optical fibre, using wavelength-division multiplexing and throughputs that can reach tens or hundreds of gigabits.

The stumbling block of this new network is its interconnection with the ground infrastructure. Pending hypothetical data centres in orbit, all of the constellations' data must be sent up to

and brought back down from them. At present, the access segment is necessarily radio-frequency: 5G meets the need and the framework appears relatively clear on this front. However, for the trunking segment and the access of large corporate clients, this radio-frequency solution is intrinsically limited. First, it is limited in bandwidth, even with the opening of new frequency bands. We are dealing with multiples of 1 Gbps, or even 10 Gbps, with very little hope of going much higher. Indeed, the highest-performing spectral bands are in turn caught up by physics and fading effects. Next, these spectral bands have the drawback of being regulated and of requiring operating licences that are often costly. In addition, ground and space infrastructures are not very efficient in terms of energy consumption and compactness. Finally — and this is perhaps the most problematic point at present — these radio-frequency links are relatively easy to jam and to interrupt. It is also fairly simple to identify the emission points in a military context. The alternative to these radio-frequency trunking links between the ground and the satellites is of course the optical link, using the laser as a carrier, in the same way as the links between satellites operate. Nevertheless, it is more difficult to cross the atmosphere than space or an optical fibre, and clouds and atmospheric turbulence strongly affect the performance of these links, or even directly their availability in the case of clouds.

If we approach the problem in a very theoretical and macroscopic way, focusing on the need to interconnect these two network layers — space and ground — optically, one could well imagine placing ground stations at highly favorable locations, for example in deserts, and thus covering the entire planet with a dozen ground stations. Even a very varied (and privileged) region such as Europe can be covered fairly simply thanks to its Mediterranean zone. Obviously, this textbook case remains highly theoretical, since one cannot imagine a network that depends, at any given moment, on a single link, and the resulting worldwide bandwidth of these few interconnections would be quite insufficient. Finally, placing a ground station in the Atacama Desert to interconnect a server in Asia with a client in the United States makes no sense. The trend that is emerging very clearly is that access and interconnection infrastructures with the space layer must be close to the services being used and to the servers running them. Yet the deployment sites of these data centres have not, to date, taken into account this need for a clear, low-turbulence, sky. The most active zones today for data-centre deployment — South Asia in particular — are, on the contrary, fairly unfavorable to optical links with space.

From a telecom standpoint, the infrastructure that makes sense consists in reasoning by regional plate, with the best possible resilience over a zone of a few hundred or a few thousand kilometres. The definition of this zone is underpinned by the availability of the optical-fibre links that will complete the mesh. For example, such a zone would be Western Europe, North America or Southeast Asia. The aim is to serve the constellation as close as possible to its users, in order to reduce latency and transit across the global Internet. The logic is of course the same in the case of sovereign networks and defence networks.

Case study

In this document we will study a fairly difficult case: Southeast Asia. We will try to serve this region as well as possible by deploying ground stations across several countries and by attempting to optimize the cloud-cover and atmospheric-turbulence criteria in order to achieve the highest possible availability. In this context, it is clear that we must give up the five “nines” that constitute the availability criterion of conventional operators. We must accept that this orbital layer cannot, with optical links, offer such a level of availability, and in all cases keep radio-frequency trunking to be able to fill the gaps — sometimes seasonal — in the optical coverage.

To carry out this simulation we used Miratlas's tool, LaserCom Pathfinder, which makes it possible to compute the availability of an aggregate of sites and to optimize the choice of sites and their location. This model uses a database with a history of around 5 years, in order to take into account the evolution of the climate, which is occurring over ever-shorter timescales. The model accounts for both cloud cover — whether for a geostationary or a low Earth orbit application — and atmospheric turbulence, which affects both the investment level of the solution and its throughput performance.

This vast South Asian region, between India, China and Australia, is of course extremely active economically and is also a new, very active area for data-centre deployment. The region is fairly widely covered by optical-fibre networks.




Figure 2: Submarine Cable Map ASEAN, Source Telegeography

On the other hand, this region is the scene of numerous geopolitical tensions, and the submarine optical-fibre links have fairly obvious points of concentration that are as many points of vulnerability, notably in the Strait of Malacca.



Figure 3: Submarine Cable Map Singapore, source Telegeography

Between the tropics and the equator, this region is subject to the monsoon regime, with a very wet period between June and September. This monsoon regime will be a recurring difficulty in the deployment of optical links and penalizes the annual availability rate with a difficult, yet relatively brief, period. On its own, the monsoon requires keeping radio-frequency access in parallel with the optical link.




Figure 4: Seasonal maps with a very different cloud cover depending on monsoon.

We therefore ran a model over the following zone:




Figure 5: Site selected for analysis

The blue points represent the sites considered in the optimization, and those circled in yellow are the ones selected among the 25 stations.

One can note the presence of two sites in Bangladesh and one site in northern Australia. The other sites are relatively well distributed across the oceanic region, with the notable exception of Borneo, owing to its rather wild environment and the presence of a single interconnection point.

By default, the sites considered are IXPs. Obviously, in the case of a real study, these would be the user's sites — whether data centres, interconnection points or teleports. This is therefore a very synthetic exercise here, but one that gives a very good idea of the difficulty and of the access results in this region, which is among the most challenging.




Figure 6: Time availability chart

The chart above shows the link availability as a function of the day and the hour. The minimum number of available stations is studied between 0 and 3. Obviously, on many days a greater number of stations are available.

Unsurprisingly, the non-availability points concern the monsoon episode, with a peak at daybreak and another in the late afternoon, at the time of the heavy rainfall. Over a large part of the year, from September/October until the beginning of June, the availability rate is very high. In practice, this means that over the summer period a terrestrial or radio-frequency backup would be necessary.

Obviously, broadening the region slightly to include drier environments such as eastern India or northern Australia would have a very significant upward impact.

In the end, the average annual availability with these 25 sites — which are not really optimized — is 98.32%. This constitutes a very respectable score, given that most of the non-availability is concentrated in the monsoon season.




Figure 7: Cumulative network availability versus the number of ground stations (Southeast Asia). Source: Miratlas LaserCom Pathfinder.

Regarding turbulence, an availability threshold was set as a balance between the isoplanatic angle and the Fried parameter.

The assumption does not include adaptive optics.

For cloud cover, a site is ignored at a given moment if the total cloud cover (TCC) exceeds 30%.

Conclusion

In conclusion, this brief study — carried out on theoretical points but over a region that is very difficult owing to its tropical and equatorial climate — shows that one can achieve an availability rate that is, overall, very good over a large part of the year. Nevertheless, the specific conditions during the monsoon period mean that a purely optical network is not sufficient to guarantee an availability level adequate for most applications, and a radio-frequency backup to the satellite constellation remains necessary.

Planning and modelling an interconnection between terrestrial networks and networks in low Earth orbit or GEO therefore requires prior statistical-analysis work based on past meteorological data, but also on-site analysis in order to avoid falling into inconsistencies in the models.