# Deploying Command and Control Services to Crisis Areas more Rapidly with Cloud Environments

Juha Järvinen, Deepak Choudhary, Jukka Manner
*Department of lnformation and Communications Engineering*
*Aalto University*
Postal Address: PO Box 13100, FI-00076 AALTO
Email: {Juha.Tapio.Jarvinen, Deepak.Lalith, Jukka.Manner}@aalto.fi

*Abstract*—**Cloud environments have become increasingly prevalent. They are advantageous and convenient, particularly when a digital service necessitates rapid deployment, scaling, or relocation to a closer geographical site, for instance, in response to delays. Concurrently, last-mile mobile communications connections, along with connections provided by Low Earth Orbit (LEO) satellites, have recently emerged as a feasible technology, attributed to enhanced bandwidth, reduced latency, and decreased costs, particularly in remote regions. In crisis situations, such as natural disasters, prompt coordination of assistance, specifically Command and Control (C2) capability, is imperative. Traditionally, C2 equipment has been dispatched to the site in shipping containers, a process that is notably slow.**

**This paper presents a novel approach to the secure establishment and operation of C2 services, particularly in geographically dispersed and extensive environments. An environment has been established to measure latency and jitter, which are critical factors for the success of C2 operations. Finally, the performance of two services — voice and chat — is analyzed across various scenarios studied these findings.**

## I. Introduction

Cloud environments have become increasingly familiar to the populace with phrases such as ”use cloud services” and ”your banking service is in the cloud.” Frequently, service providers and companies appear to exaggerate the novelty and superiority of their offerings by emphasizing the term ”cloud” in their marketing efforts. For the average citizen, the concept of the cloud is intended to signify that services are perpetually available and accessible from any location; this notion should be an inherent characteristic of the Internet, independent of the terminology employed.

Critical services vary by individual; some see them as access to websites, email, and TV, while others view them as access to grocery stores and social media. In this paper, critical services refer to Command and Control (C2) services over an IP data network that support society or organizations during emergencies, like life-saving operations. These services encompass, for instance, voice communication (e.g., Voice over IP, VoIP) and text-based collaboration (e.g., Extensible Messaging and Presence Protocol, XMPP).

Authorities or organizations must quickly establish services during urgent situations, like natural disasters, which damage infrastructure, including communications. Rescue personnel and humanitarian organizations need immediate access to these services to coordinate operations effectively. The critical period, known as the *golden* 72 hours, greatly increases survival chances if rescue efforts are initiated [1]. Similarly, military operations, such as United Nations (UN) peacekeeping missions [2], require rapid service deployment. These situations can arise suddenly, like earthquakes, or be part of preplanned operations with extensive preparations.

Several papers have been presented utilizing cloud computing to assist citizens, such as providing instructions to mobile phones for navigating out of disaster areas [3], [4]. d’Oro et al. [5] propose an architecture that integrates Edge and Cloud computing for small-area emergency management, specifically tailored for emergency workers; however, the paper neglects to address the organization of telecommunications and its boundary conditions. Kuwata et al. [6] concentrate on local emergencies in contexts where the local mobile network is accessible. Stodola et al. present an Advanced Command and Control System [7] designed for military operations; nonetheless, the services are localized in this instance as well. Mohan et al. [8] analyzed Starlink’s suitability for real-time applications and found it comparable to cellular networks for tools like Zoom and Luna cloud gaming, though proximity to ground infrastructure affects this comparison.

This paper introduces a novel approach to the provision of C2 services in cloud environments, emphasizing rapid and secure operations irrespective of geographical distance. It’s especially relevant for authoritative services in various applications, like emergency response. The framework was tested with real-world latency measurements across two C2 services in three case studies, revealing significant constraints in cloud service design and implementation.

For example, latency between Finland and Australia is high, with a 314 ms Round-trip time (RTT) over a Virtual Private Network (VPN). This affects call quality for Sydney users when the service is in Otaniemi, Finland. If security allows, moving the service to a nearby public cloud would improve Quality of Experience (QoE). Relocating classified data to a public cloud is difficult if it requires a private cloud in Finland. In contrast, latency is not a major issue for communication from Stockholm to Sydney when using a service in Otaniemi, even if it relies on a stationary private cloud.

This article is organized as follows. We start with an overview of operating environment in the context of this paper in Section II. After that we present our measurement setup in Section III and results in Section IV. Then we discuss findings related to the topic of this paper in Section V, and conclude the article with future directions in Section VI.

979-8-3315-7896-1/26/$31.00 ©2026 IEEE

## II. Critical Communications and Clouds

Earthquakes and tsunamis are unpredictable disasters. Unlike floods, which can be anticipated, earthquakes strike suddenly. For instance, a 7.7 magnitude earthquake in Myanmar on March 28, 2025, led to over 3,300 deaths and 4,500 injuries, impacting more than three million people [9]. In January 2010, a major earthquake in Haiti caused an estimated 230,000 fatalities. The tsunami from the December 2004 Indian Ocean earthquake resulted in 184,167 deaths across 14 countries, with a total death toll of 227,898 [10]. Deadly earthquakes often occur in areas with weak building infrastructures, leading to collapses. They also frequently happen in challenging terrains, like mountains. Furthermore, factors such as conflict can obstruct rescue efforts, highlighting the urgent need for effective rescue operations.

### A. Communication

Communication is vital in rescue and military operations. Initially, communication relies on voice communications, like unencrypted Citizens Band (CB) radio, for local use. However, modern C2 services, including chat, video, and web-based applications, require data connections via 4G, 5G, or satellite links. Wired telecommunications should also be considered where possible.

However, if the infrastructure, including data networks, is compromised (e.g., during a major earthquake) or the relief area is underdeveloped, the relief team will initially lack C2 services, relying only on verbal communication from CB equipment.

Data network nodes and C2 services in shipping containers are gradually deployed to the assistance area. They connect to the home country via satellite links if needed, but data transfers are minimal, and the containers work autonomously. However, this approach isn't a zero-day solution and doesn't adequately meet the critical 72-hour time frame for C2 services post-deployment, even if prepared at the origin. Conversely, a notable advantage of this containerized approach is the gradual establishment of a local network around the containers.

C2 connections to external networks have historically depended on slow, low-capacity, high-latency, and costly satellite links. These links, reliant on geostationary satellites located approximately 35,786 kilometers away, yield an average RTT of around 600 ms and possess limited capacity. The situation has changed significantly with new communications satellites deployed into Low Earth Orbit (LEO), 160 to 2000 kilometers above sea level [11]. LEO satellites have reduced latency, with RTT as low as 30–60 ms [12], making them ideal for latency-sensitive applications. Maintaining continuous coverage requires a large constellation of satellites. Commercial satellites have increased [13], like Amazon's Kuiper[1] (3236 satellites), SpaceX's Starlink[2] (4000 satellites), and OneWeb[3] (650 satellites). This has led to lower data connection prices for consumers. Currently, about 7,100 satellites are in LEO [14], expected to rise to 50,000 in the next decade [15].

Commercial entities have created new uncertainties in market utilization. Many countries aim to establish connections using affordable Internet from LEO satellites. However, the Ukraine conflict has revealed vulnerabilities not rooted in technology. For example, Ukrainian forces tried to use Starlink for drone operations, but access was sometimes restricted [16]. The European Union is working to expedite Eutelsat satellite deployment as an alternative to commercial operators [17].

LEO satellite communications have a key advantage over fiber networks: lower latency [18]. LEO networks achieve shorter RTT due to the faster speed of radio waves (300,000 km/s) compared to optical signals in silica fiber (200,000 km/s). This allows satellites to sometimes outperform fiber in propagation delay, especially over long distances, despite their higher altitude [19].

Ma *et al.* [20] studied latency across nine Amazon Web Services (AWS) regions, comparing Starlink to terrestrial networks. Results show Starlink's latency is about 10% higher, with more instability. Contributing factors include physical obstructions, satellite movement, and Internet Service Providers (ISP) routing decisions. Mohan et al. [8] noted that Starlink shows performance variations likely due to internal network reconfigurations every 15 seconds. The findings indicate these reconfigurations are globally synchronized, not caused by satellite handovers.

### B. C2 Services: Voice and Instant Messaging

As previously stated, this paper focuses on two C2 communication services: voice and instant messaging (IM). Voice remains crucial in communication, defying predictions of becoming obsolete. In IP networks, Voice over Internet Protocol (VoIP) enables real-time voice transmission. Once an end-to-end connection is established, data is typically sent using the Real-time Transport Protocol (RTP) [21] over the User Datagram Protocol (UDP). VoIP networks mainly use a Client-to-Server (C2S) architecture, with the server called a Private Branch Exchange (PBX) [22].

Extensible Messaging and Presence Protocol (XMPP) is an IM protocol that uses TCP, mainly for government applications. It's based on five Request for Comments (RFC) [23]. Many commercial messaging apps like WhatsApp and Facebook Messenger originally built on XMPP. It also acts as a C2S protocol, like VoIP.

### C. Quality of the Communication Connection

Round-trip time (RTT) [24], [25] is a fundamental metric for evaluating network status and user experiences pertaining to Quality of Experience (QoE). A reduced RTT signifies a seamless, low-latency experience for users. RTT may vary from a few milliseconds within a building to hundreds of milliseconds across continents. Measurements can be categorized as active, utilizing probes such as Internet Control Message Protocol (ICMP) [26], or passive, derived from packet capture timestamps.

Jitter refers to the variation in packet arrival times, impacting end-user QoE in real-time applications and gaming.

[1] https://www.aboutamazon.com/what-we-do/devices-services/project-kuiper

[2] https://www.spacex.com

[3] https://oneweb.net

Unlike RTT, jitter measures delivery inconsistency. High jitter means packets arrive irregularly, often due to congestion, route changes, and variable queuing delays.

Latency and jitter significantly impact protocols, especially VoIP. These applications rely on the timely arrival of voice packets for real-time speech. Jitter can disrupt this, degrading voice quality with dropouts, echoes, or lost content. If a voice packet is delayed, it may be discarded by the receiving device, leading to lost information.

Humans can detect call latency as low as 150 milliseconds. The International Telecommunication Union (ITU) establishes one-way latency limits at 150 ms and recommends a maximum of 300 ms for round-trip latency [27]. Exceeding these thresholds can adversely affect call quality, resulting in distorted conversations. Latency exceeding 400 milliseconds is deemed unacceptable; however, VoIP users generally perceive lag at approximately 300 ms. XMPP over TCP is less sensitive to latency than VoIP but encounters challenges related to high jitter in real-time updates.

### *D. Different Cloud Environments*

The National Institute of Standards and Technology (NIST) identifies five characteristics of cloud services: 1) Self-service in accordance with requirements, 2) Extensive access to resources via the network, 3) Centralized resources, 4) Rapid elasticity, and 5) Metered service. These criteria outline the creation of a cloud computing environment. Simply connecting ten computers with an operating service doesn't constitute a cloud service; management software and infrastructure are essential for flexible user interaction with resources. Cloud environments can be categorized into three distinct groups: public, private, and hybrid.

*1) Public Cloud:* A public cloud is a variant of cloud service accessible via the Internet. Prominent global cloud providers develop optimal public cloud environments, facilitating the migration of services in proximity to customer requirements, thereby reducing latency and enhancing response time. Nevertheless, public clouds present challenges to data security, as operators possess the capability to access service data, which constrains governmental services and influences aid organizations that manage classified information.

*2) Private Cloud:* Private clouds limit users to a specific group, usually within a single organization, resembling an internal network [28]. Enforcement methods for this limitation vary. Unlike public clouds, private clouds, especially for governments, prioritize national infrastructure for security. Data sharing across borders is often resisted, even among the European Union (EU, an economic and political union) or North Atlantic Treaty Organization (NATO, a military alliance). Public clouds generally avoid handling classified information, and increased data classification leads to restricted traffic. A VPN from an isolated network may be inadequate, as some connections are limited to one building. Sometimes, private clouds appear as a few virtualized servers but have higher maintenance costs due to separation, starting with the maintenance staff.

*3) Hybrid Cloud:* Public cloud environments may not always be ideal for processing sensitive data. Private clouds have their challenges, especially with connectivity. The hybrid cloud model combines the benefits of both, balancing public cloud flexibility with data privacy. One way to implement a hybrid solution is to use private clouds for sensitive data and public clouds for less critical data, requiring an Information Exchange Gateway (IEG) for data flow management. Another approach is using Homomorphic Encryption (HE), which allows operations on encrypted data in the public cloud [29]. Despite advancements since its successful implementation in 2009 [30], HE remains complex and costly for widespread use [31].

## III. Measurement Setup

Latency measurements have a historical foundation and are not a novel concept. However, the objective of these measurements is to establish a foundation of empirically obtained results for the discussion section of the subsequent part, where various use cases and their usability will be analyzed through two C2 services informed by these measured delays. While numerous continuous latency measurement platforms exist globally, from which values could be extracted between measurement points, our measurements facilitate the use of a precisely tailored measurement platform.

Measurements were conducted in two distinct scenarios. In the first scenario, the measurements were executed with the user's endpoint situated in close proximity to the declared physical location of the cloud service. In the second scenario, the distance between the two was maximized.

The user endpoint was situated in Otaniemi, Finland at Aalto University campus. Measurements were conducted utilizing AWS cloud infrastructure at two distinct locations (regions). The first location was the nearest AWS site to Otaniemi, Stockholm, Sweden (approximately 400 km, *eu-north-1* region), while the second location was one of the most distant from Otaniemi, Sydney (approximately 15200 km, *ap-southeast-2* region) in Australia. Wired connections were employed for the measurements, with the user endpoint in Otaniemi connected to the Finnish Universities Network (FUNET).

ICMP latency measurements were conducted on three routes: Otaniemi $\leftrightarrow$ Stockholm, Otaniemi $\leftrightarrow$ Sydney, and Stockholm $\leftrightarrow$ Sydney. The Virtual Private Network (VPN) configuration involved a site-to-site IPsec tunnel established between VyOS (local) and the AWS Virtual Private Gateway. This VPN was configured utilizing IKEv1, employing AES-128 encryption and SHA-1 hashing algorithms. An open-source VyOS[4] router operated on a VMware ESXi[5] platform. The network interfaces were organized into distinct public and private zones. The Debian virtual machine, situated behind the VyOS router, utilized the tunnel to access private AWS IP addresses. The measurement setup is illustrated in Fig. 1.

[4] https://vyos.io

[5] https://www.vmware.com

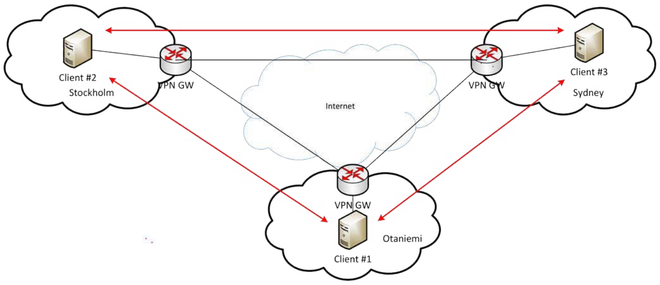

Fig. 1: Measurement setup.

## IV. Results

The measurements are conducted based on the Round Trip Time (RTT) principle, wherein the round trip delay of the packet is aggregated. While it cannot be definitively asserted that the routes taken by the round trip and one-way packets are identical, it can be posited that the One-Way Delay (OWD) in this context is approximately RTT/2.

The RTT measurements conducted between the on-premise end-host device in Otaniemi and the AWS instance in Stockholm reveal a small and reasonable differentiation between public Internet routing and site-to-site VPN path, which is illustrated in a time-series plot in Fig. 2 and a Cumulative Distribution Function (CDF) plot in Fig. 5a. On average, the public routing exhibited lower latency (mean RTT = 11.44 ms) relative to the VPN path (mean RTT = 15.09 ms), which can be attributed to the encryption and decryption process at VPN gateways at both locations. More significantly, the uniformity of the public path merits attention: both the standard deviation and jitter were markedly lower (0.62 ms and 0.50 ms, respectively) compared to those recorded in the VPN measurements (1.28 ms and 0.75 ms).

In Fig. 2 we can see that the public path consistently maintained a close clustering around the mean, whereas the VPN path exhibited intermittent spikes that surpassed 30 ms. These spikes may be attributable to tunnel re-keying delays or transient congestion along the encrypted path.

In Fig. 3 and Fig. 5a, the Otaniemi-Sydney connection effectively illustrates the geographical distance involved: public IP routing achieved a mean RTT of approximately 311 ms, whereas VPN routing increased this metric to 314 ms. The overhead of 3 ms introduced by the VPN tunnel is likely attributable to the encryption and decryption processes, as demonstrated in the Otaniemi-Stockholm case.

The distance does not affect on stability, both routes exhibited relatively stable latency. Public routing demonstrated lower jitter (0.535 ms) and standard deviation (0.736 ms), indicating more consistent packet delivery timing. In contrast, VPN routing displayed higher variability, with standard deviation of 1.251 ms and jitter of 0.820 ms and several noticeable latency spikes exceeding 330 ms in the time-series plot as illustrated in Fig. 3. Correspondingly, a CDF plot is presented in Fig. 5b.

The Otaniemi-Stockholm and Otaniemi-Sydney graphs illustrate that VPN usage inherently contributes to increased latency between points, as it necessitates additional processing at both endpoints for the encryption and decryption of IP packets. However, the observed increase in latency is minimal and remains consistent (approximately 3 ms) when comparing these two scenarios, indicating that it does not vary as a function of distance. Any deviation from this outcome would be unexpected, given that routers and similar devices in the network treat encrypted packets in the same manner as other IP packets.

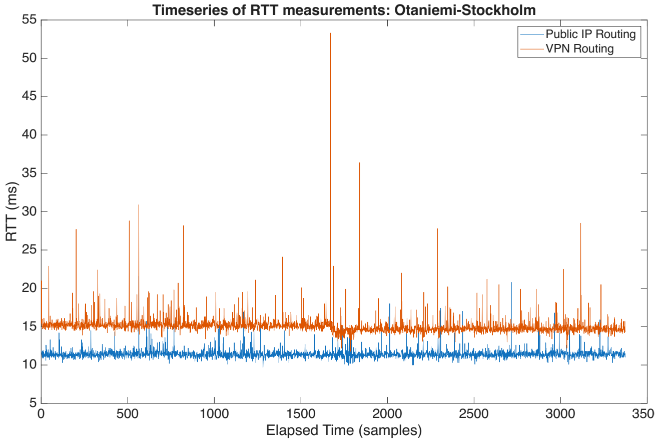

Fig. 2: Timeseries of RTT measurements: Otaniemi - Stockholm. The blue color signifies traffic in public Internet routing, while the orange color denotes traffic over site-to-site VPN routing/VPC Peering.

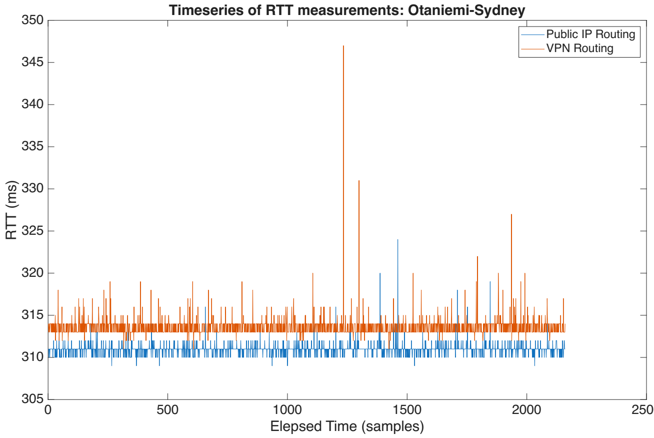

Fig. 3: Otaniemi - Sydney.

In comparison to the initial two cases in this study, the Stockholm-Sydney path demonstrated the most stable performance despite its intercontinental span, see Fig. 4 and Fig. 5c. Public IP routing achieved a mean RTT of 269.63 ms, whereas VPN path added less than 1 ms on average, resulting in a mean RTT of 270.44 ms. Both routing types exhibited narrow RTT distributions: the standard deviation for public routing was 0.605 ms, whereas VPN routing demonstrated a standard deviation of 0.913 ms. Additionally, jitter remained exceptionally low, at 0.183 ms for the public route and 0.950 ms for the VPN route. In each of three measurement cases, the packet drop rate was recorded at 0 %.

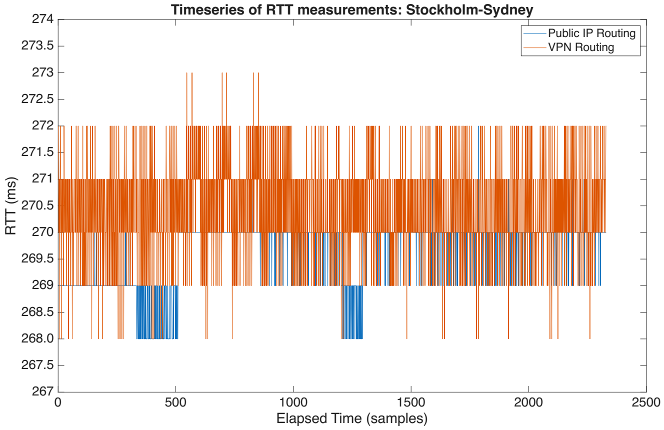

Fig. 4: Stockholm - Sydney.

The technology utilized between Stockholm and Sydney differs from that employed by the other two locations, which utilized site-to-site VPN. This observation elucidates the comparatively smaller increase in latency associated with the use of a ”VPN” during this interval, in contrast to the other two cases. An internal AWS connection—VPC Peering—was implemented between these two locations. AWS provides limited details regarding the implementation of this connection, stating only that [32] ”A VPC peering connection is a networking feature that enables secure and direct communication between two virtual VPCs within the AWS infrastructure.”

## V. Discussion

What impact do the measurement results from the previous section have on service topology design? The most significant findings are discussed below based on the topology illustrated in Fig. 6. In addition to the latency measurement, two exemplary protocols, VoIP and XMPP, were analyzed to evaluate the impact of latency on these distinct types of protocols.

In the initial measurement, two users in Finland connected to an AWS server in Stockholm (see Fig. 6). Each user had a separate Client-to-Server connection, but they could not communicate via Client-to-Client, despite being close and the service allowing it. Thus, the delay between the users matches the RTT value from earlier. This delay measurement excludes processing delays from the service, but all services, including ping (ICMP echo), involve processing, affecting overall delays. Processing delays depend on factors like CPU, memory, and software [33].

The 150 ms delay between users delineated by ITU [27] is markedly lower in the Otaniemi-Stockholm scenario, where the mean RTT is 15.09 ms and jitter is 0.75 when employing a VPN connection. Consequently, it would not present an issue that the service itself is geographically situated approximately 400 kilometers from the users.

The Otaniemi-Sydney case shows concerns with VoIP traffic. The OWD was 314 ms over a VPN, twice the ITU threshold. This delay can lead to ”speak-and-wait” discussions. While instant messaging is fine, XMPP-based services, like fire control apps, may struggle. Other larger factors likely contribute to the communication delay. Jitter is minimal and consistent; the main issue is latency.

Delays change significantly when communication shifts from user A and user B, as depicted in Fig. 6, to user A and user C, with C as a command center. User A is in Sydney, and the command center is in Stockholm, while services are in Otaniemi. The OWD between user A and user C is about 150 ms, which is acceptable for normal speech communication.

This clearly indicates that the technical plan and the measurements derived from it, in isolation, do not unequivocally determine the efficacy of the plan; they need evaluation against the intended use case. For instance, the Sydney-Otaniemi connection showed that speech delay hindered normal conversation in one context but was not an issue in another.

The third observation focuses on hybrid clouds using a public-private pairing. Users A and B in Otaniemi use public services in Stockholm. Some internal services are on a private cloud in Sydney (VPC Peering) for security. Latency issues are minor for chat applications but problematic for VoIP, especially when the PBX is in Sydney and accessed via VPN by users.

## VI. Conclusions

C2 services in cloud environments with LEO satellite connections provide quick command and coordination support for urgent needs, like after a major earthquake, unlike traditional container solutions. VoIP and XMPP show varying latency in different scenarios since service locations can be adjusted near the area of need.

When linking a security class to data, consider additional use cases to ensure a standard service experience, especially for VoIP. Classified data usually stays in a static location within a private cloud, while assistance may be needed globally. Planning can be improved by deploying latency-sensitive services in public clouds closer to the point of use.

The next phase involves the establishment of a comprehensive measurement platform that incorporates LEO satellite connections, through which QoE associated with VoIP and XMPP can be assessed, analyzed, and validated in a real-world setting.

## References

[1] E. B. Hsu, M. Ma, F. Y. Lin, M. J. VanRooyen, and F. M. Burkle, “Emergency medical assistance team response following taiwan chi-chi earthquake,” *Prehospital and Disaster Medicine*, vol. 17, no. 1, p. 17–22, 2002.

[2] “United nations: What is peacekeeping,” https://peacekeeping.un.org/en/what-is-peacekeeping, accessed: 2025-08-08.

[3] D. Facchinetti, G. Psaila, and P. Scandurra, “Mobile cloud computing for indoor emergency response: the ipsos assistant case study,” *Journal of Reliable Intelligent Environments*, vol. 5, 09 2019.

[4] M. Qiu, Z. Ming, J. Wang, L. T. Yang, and Y. Xiang, “Enabling cloud computing in emergency management systems,” *IEEE Cloud Computing*, vol. 1, no. 4, pp. 60–67, 2014.

[5] E. Cavalieri d’Oro, S. Colombo, M. Gribaudo, M. Iacono, D. Manca, and P. Piazzolla, “Modeling and evaluating a complex edge computing based systems: An emergency management support system case study,” *Internet of Things*, vol. 6, p. 100054, 2019.

[6] Y. Kuwata, Y. Ishikawa, and H. Ohtani, “An architecture for command and control in disaster response systems,” in *2000 26th Annual Conference of the IEEE Industrial Electronics Society. IECON 2000.*, vol. 1, 2000, pp. 120–125 vol.1.

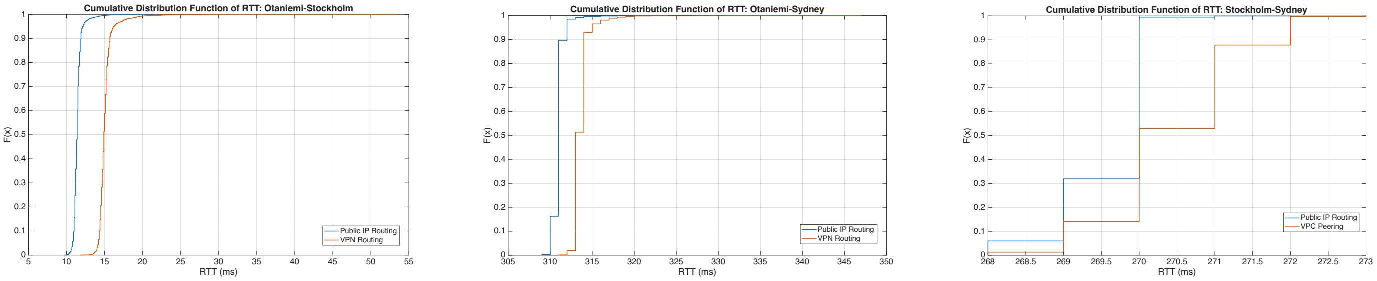

(a) Otaniemi - Stockholm.

(b) Otaniemi - Sydney.

(c) Stockholm - Sydney.

Fig. 5: CDF plots of RTT measurements. The blue color signifies traffic in public Internet routing, while the orange color denotes traffic over site-to-site VPN routing/VPC Peering.

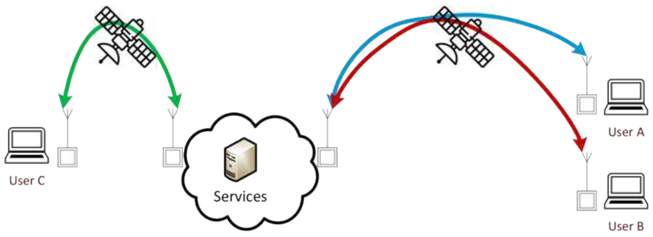

Fig. 6: The figure illustrates the topology relevant to the discussion of the measurement results of the various distances measured in the previous section.

[7] P. Stodola and J. Mazal, “Architecture of the advanced command and control system,” in *2017 International Conference on Military Technologies (ICMT)*, 2017, pp. 340–343.

[8] N. Mohan, A. E. Ferguson, H. Cech, R. Bose, P. R. Renatin, M. K. Marina, and J. Ott, “A multifaceted look at starlink performance,” in *Proceedings of the ACM Web Conference 2024*, ser. WWW ’24. New York, NY, USA: Association for Computing Machinery, 2024, p. 2723–2734.

[9] “United Nations: Myanmar quake: Ongoing aftershocks spread fear,” https://news.un.org/en/story/2025/04/1162711, accessed: 2025-05-05.

[10] “Wikipedia: 2004 indian ocean earthquake and tsunami,” https://en.wikipedia.org/wiki/2004_Indian_Ocean_earthquake_and_tsunami, accessed: 2025-05-06.

[11] N. U. Hassan, C. Huang, C. Yuen, A. Ahmad, and Y. Zhang, “Dense small satellite networks for modern terrestrial communication systems: Benefits, infrastructure, and technologies,” *IEEE Wireless Communications*, vol. 27, no. 5, pp. 96–103, 2020.

[12] L. Izhikevich, M. Tran, K. Izhikevich, G. Akiwate, and Z. Durumeric, “Democratizing leo satellite network measurement,” *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 8, no. 1, Feb. 2024. [Online]. Available: https://doi.org/10.1145/3639039

[13] B. Al Homssi, A. Al-Hourani, K. Wang, P. Conder, S. Kandeepan, J. Choi, B. Allen, and B. Moores, “Next generation mega satellite networks for access equality: Opportunities, challenges, and performance,” *IEEE Communications Magazine*, vol. 60, no. 4, pp. 18–24, 2022.

[14] “Low earth orbit,” https://orbit.ing-now.com/low-earth-orbit/, accessed: 2025-07-26.

[15] C. Daehnick, I. Klinghoffer, B. Maritz, and B. Wiseman, “Large leo satellite constellations: Will it be different this time?” Tech. Rep., May 2020. [Online]. Available: https://www.mckinsey.com/industries/aerospace-and-defense/our-insights/large-leo-satellite-constellations-will-it-be-different-this-time

[16] “Wikipedia: Starlink in the russian-ukrainian war,” https://en.wikipedia.org/wiki/Starlink_in_the_Russian-Ukrainian_War, accessed: 2025-05-08.

[17] “The register: Eutelsat in talks with euro leaders as they mull starlink replacement in ukraine,” https://www.theregister.com/2025/03/07/eutelsat_starlink_ukraine, accessed: 2025-05-08.

[18] M. Handley, “Delay is not an option: Low latency routing in space,” in *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*, ser. HotNets ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 85–91.

[19] J. Conde, G. Martínez, P. Reviriego, and J. A. Hernández, “Round trip times (rtts): Comparing terrestrial and leo satellite networks,” in *2024 27th Conference on Innovation in Clouds, Internet and Networks (ICIN)*, 2024, pp. 42–46.

[20] S. Ma, Y. C. Chou, H. Zhao, L. Chen, X. Ma, and J. Liu, “Network characteristics of leo satellite constellations: A starlink-based measurement from end users,” in *IEEE INFOCOM 2023 - IEEE Conference on Computer Communications*, 2023, pp. 1–10.

[21] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, “RTP: A Transport Protocol for Real-Time Applications,” RFC 3550 (Internet Standard), RFC Editor, Fremont, CA, USA, Jul. 2003.

[22] A. Roach, “Registration for Multiple Phone Numbers in the Session Initiation Protocol (SIP),” RFC 6140 (Proposed Standard), RFC Editor, Fremont, CA, USA, Mar. 2011.

[23] “XMPP RFCs,” https://xmpp.org/rfcs/, accessed: 2025-08-12.

[24] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, “Framework for IP Performance Metrics,” RFC 2330 (Informational), RFC Editor, Fremont, CA, USA, May 1998, updated by RFCs 7312, 8468.

[25] G. Almes, S. Kalidindi, and M. Zekauskas, “A Round-trip Delay Metric for IPPM,” RFC 2681 (Proposed Standard), RFC Editor, Fremont, CA, USA, Sep. 1999.

[26] J. Postel, “Internet Control Message Protocol,” RFC 792 (Internet Standard), RFC Editor, Fremont, CA, USA, Sep. 1981, updated by RFCs 950, 4884, 6633, 6918.

[27] Telecommunication Standardization Sector of ITU., *ITU-T Recommendation G.114: Transmission Systems and Media : General Recommendations on the Transmission Quality for an Entire International Telephone Connection : One-Way Transmission Time*. International Telecommunication Union, 1994.

[28] N. K. Sehgal and P. C. P. Bhatt, *Cloud computing : concepts and practices*. Springer, 2018.

[29] R. L. Rivest, L. Adleman, M. L. Dertouzos *et al.*, “On data banks and privacy homomorphisms,” *Foundations of secure computation*, vol. 4, no. 11, pp. 169–180, 1978.

[30] C. Gentry, “Fully homomorphic encryption using ideal lattices,” in *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, ser. STOC ’09. New York, NY, USA: Association for Computing Machinery, 2009, p. 169–178.

[31] C. Dobraunig, L. Grassi, L. Helminger, C. Rechberger, M. Schofnegger, and R. Walch, “Pasta: A case for hybrid homomorphic encryption,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2023, no. 3, p. 30–73, Jun. 2023.

[32] “AWS: What is VPC peering?” https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html, accessed: 2025-07-16.

[33] R. Ramaswamy, N. Weng, and T. Wolf, “Characterizing network processing delay,” in *IEEE Global Telecommunications Conference, 2004. GLOBECOM ’04.*, vol. 3, 2004, pp. 1629–1634 Vol.3.