Balloons in the Sky: Unveiling the Characteristics and Trade-offs of the Google Loon Service

Pablo Serrano, *Senior Member, IEEE,* Marco Gramaglia, Francesco Mancini, Luca Chiaraviglio, *Senior Member, IEEE,* and Giuseppe Bianchi

Abstract—The Google's Loon^{TM} initiative aims at covering rural or underdeveloped areas via fleets of high-altitude balloons supporting LTE connectivity. But how effective and stable can be the coverage provided by a network deployed via propulsion-free balloons, floating in the sky, and only loosely controllable through altitude variations? To provide some insights on the relevant performance and trade-offs, in this paper we gather real-world data from publicly available flight tracking services, and we analyze coverage and service stability in three past deployment scenarios. Besides employing a variety of metrics related to spatial and temporal coverage, we also assess service continuity, by also leveraging recently proposed "meaningful availability" metrics. While our analyses show that balloons are certainly a cost-effective way to provide a better-than-nothing and delay-tolerant service, there is yet no empirical evidence that an increase in the number of overlapping balloons may be rewarded with a substantial performance increase — in other words, we suspect that guaranteeing coverage and service stability levels comparable to that of a terrestrial cellular network is a challenging goal.

Index Terms—Digital divide, UAV, HAPs, Loon service

1 Introduction

TODAY, while most countries are heading towards the realization of the so-called Gigabit Society, more than half of the population of the planet has not yet been reached by *any* form of Internet connection. Many of those 4 billion people live in rural places where telecommunications companies do not have any convenience to build cellular towers or other network infrastructures. To face this situation, as early as 2011, Google conceived an initiative devised to ultimately bring Internet access to under-served areas via fleets of high-altitude *balloons*, or *loons* for short (in this paper we will use both terms interchangeably). The relevant project, called *Project Loon*, was officially announced as a Google R&D project in June 2013,[1] and then evolved in the form of an independent commercial company, Loon LLC, in July 2018.[2] While this paper was under review, the project has been shut down[3], apparently since the company could not succeed "to get the costs low enough to build a long-term, sustainable business." In this paper, we present a methodology driven by real-life data that unveils some intrinsic challenges that arise when providing a basic telecommunication service using loons, which might have triggered such decision. Still, while our analysis focused on the Google deployment and our results were obtained before the project was shut down, we remark that the technology *per-se* is still considered a valid solution to provide mobile network coverage in remote areas, as the recent acquisition of the Intellectual Properties involved in the project by Softbank[4] demonstrates.

Indeed, the technology is very advanced. The balloons are made from advanced materials such as special types of polyethylene and are powered by solar panels. They are not equipped with any propulsion engine: only their altitude can be controlled, by adjusting the volume and density of the internal lighter-than-air compound. However, carefully designed algorithms, based on machine learning and driven by weather forecasts, permit to change the balloon's location, by dynamically adjusting the flying altitude to opportunistically exploit *wind layers* characterized by different speeds and/or directions. Balloons can also be promptly targeted to cover on-demand a critical area — this was for instance the case in October 2017 where a fleet of balloons was deployed to provide emergency LTE coverage to Puerto Rico after Hurricane Maria wiped out the cellular infrastructure.

Once balloons reach a target area, they act as *floating cell towers*, deployed in the stratosphere at an altitude of around 20 km — about the double of other aircraft — and provide 4G LTE wireless connectivity to ground devices. To this aim, balloons are equipped with an e-NodeB and various antennas that are used to transmit LTE signals to the ground, and to communicate with other balloons and with ground gateway stations. Each balloon is expected to cover a region of tens of km in radius [1], thanks to the fact

- *P. Serrano and M. Gramaglia are with University Carlos III de Madrid. Corresponding e-mail: pablo@it.uc3m.es*
- *F. Mancini, L. Chiaraviglio, and G. Bianchi are with CNIT and Università di Roma Tor Vergata.*

Manuscript received XXX XX, XXX; revised YYY YY, YYY.

1. https://blog.google/topics/alphabet/introducing-project-loon/
2. https://blog.x.company/graduation-day-loon-and-wing-take-flight-e23a42620131
3. https://blog.x.company/loon-draft-c3fcebc11f3f
4. https://techcrunch.com/2021/09/30/alphabet-gives-some-loon-patents-to-softbank-open-sources-flight-data-and-makes-patent-non-assertion-pledge/

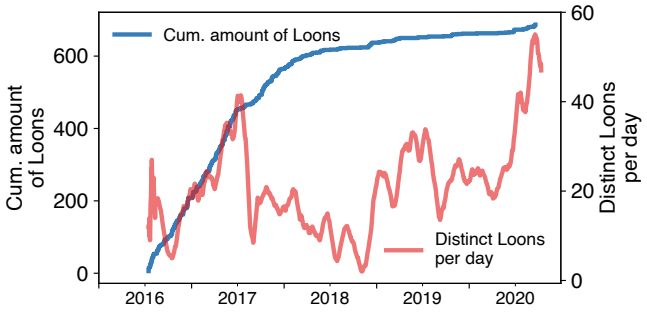

Fig. 1: Loon statistics from mid 2016.

that good propagation conditions are typically observed on the radio link between the balloon and each covered user. In a similar way, inter-balloon links can be very long, reaching up to 100 km [2].

While the first significant real-scale deployment was the above mentioned Puerto Rico's 2017 disaster scenario, the Loon technology is ultimately meant to be a cost-effective coverage solution to the difficult challenge of bringing Internet access to people in under-served remote areas. However, the floating nature of balloons and the difficulty in controlling their fine-grained position raises an obvious layman's question: can such a cost-effective solution be also a performance-effective solution, *i.e.*, can balloons guarantee some meaningful *availability and stability* of the coverage? In other words, while balloons may definitely help in providing an occasional emergency coverage, is there real-world evidence that they may also provide a stable and ubiquitous coverage service?

To the best of our knowledge, no prior work appears to comprehensively address this question upfront and using real-world evidence. On one side, several papers discuss the opportunity of using unmanned aerial vehicles (UAVs) [3], non-terrestrial networks [4] or high altitude platforms [5] to connect remote areas, or to use drones to improve the capacity of heterogeneous networks [6], [7], but either focus their analysis on system-level considerations and/or do not rely on real-world data in their evaluation. Furthermore, while the sheer majority of the works target propelled UAVs (namely, drones — see for instance the recent survey [8]), propulsion-free aerial devices such as balloons are largely neglected, arguably owing to their positioning complexity. As a matter of fact, the few Loon-specific works we could find in the literature address technical aspects related to propagation [1], [9], movement predictions [10], aviation-related details [11], or inter-loon communications [2], but do not tackle a broader analysis of the actual service level. And even the recent comprehensive network performance benchmark study performed by the Signals Research Group in September 2019 at Tarapoto (Peru) [12], commissioned by the Loon LLC company, also mainly focuses on assessing the (possible) interference impact of the additional Loon deployment on the existing terrestrial LTE network, and on measuring performance *during connectivity* to the Loon network, rather than focusing on service availability, *i.e.*, how effective and stable is the coverage.

The goal of our work is to fill this gap. However, we do not merely limit to report coverage statistics of Loon service, but we aim at providing a few more general messages. More specifically, we highlight the trade-offs between completeness of the coverage and service availability emerging for this kind of *floating cell towers*. Moreover, we focus on the investigation concerning not only service continuity, but also performance attainable when the Loon service is meant to be used as delay-tolerant network infrastructure. In addition, all the analyses documented in the remainder of this paper are performed via an independent point of view, i.e. they are exclusively based on *publicly available evidence* of the level and stability of the Loon service, in terrestrial areas and time windows where explicit coverage claims were made in the press (see details in Section 2.2). Indeed, a major hurdle that we addressed was the gathering of Loon-specific data out of publicly available flight tracking services.

Overall, the major contributions of this paper can be summarized as follows:

- We gathered real-world Loons positions over the last four years, post-processed and cast them into a tractable and easily downloadable form, hopefully useful to the research community (the details are provided in the Appendix);
- We design a wireless modeling methodology that, based on the gathered Loons' locations, do provide an estimate of the relevant coverage service under the best conditions (*e.g.*, maximum throughput);
- We demonstrate the use of this methodology by extensively characterizing the service that could have been provided in three case studies (Puerto Rico, Peru, and Kenya);
- We unveil the performance limits of the service, reveal its key trade-offs, and provide some hints of the material cost (i.e., no. of Loons) and their associated marginal improvement.

The rest of the paper is organised as follows. Section 2 describes the three case studies considered. Section 3 details the procedures to estimate the coverage and capacity based on our dataset, including the discretization of space and time. Section 4 presents average statistics about the performance provided by the service, while Section 5 presents "continuity" performance figures, to gather insight the nature of the interruptions of the service. Section 6 illustrates and discuss how our methodology can be extended to account for other channel models and network deployments. Finally, Section 7 summarizes the related work and Section 8 concludes the paper.

2 Loon: Data-sets & Case Studies

Since, to the best of our knowledge, no data-sets reporting the historical Loons' positions are publicly available, we needed to preliminary find a way to gather such data. We solved this problem by noting that Loons are *systematically* tracked using the same techniques as ordinary aircraft, and relied on crowd-sourced initiatives such as, e.g., FlightRadar24, RadarBox, or adsbScope, to collect the information about the positions of the Loons over the past several years (the details of the methodology and a link to our sanitized dataset are provided in the Appendix).

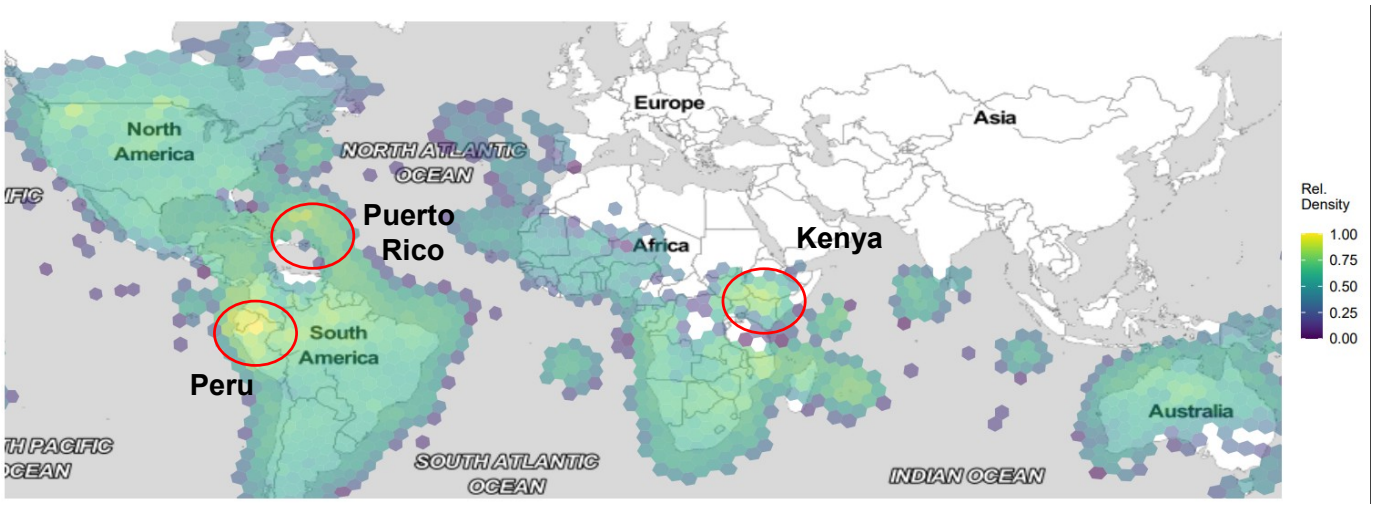

Fig. 2: Relative loon density over the Earth's terrestrial areas — as our data-set relies on data read from ground stations, activity over the Ocean could not be tracked.

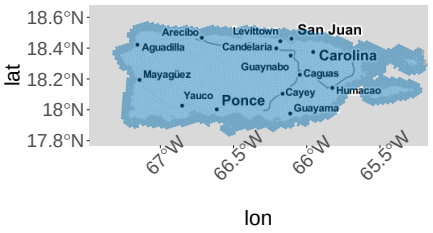

(a) Study surface over Puerto Rico

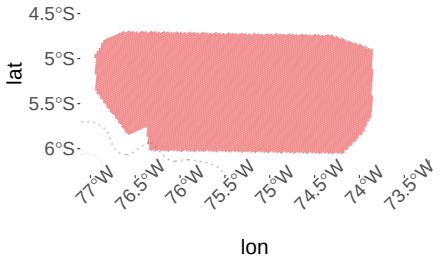

(b) Study surface over Peru

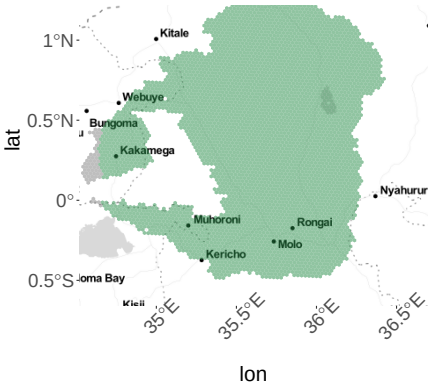

(c) Study surface over Kenya

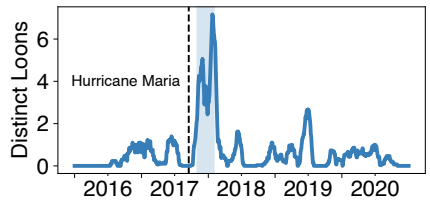

(d) Number of Loons over Puerto Rico

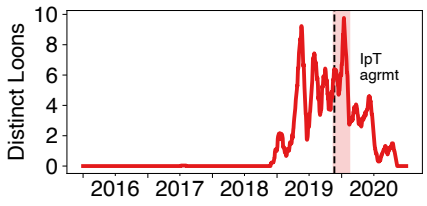

(e) Number of loons over Peru

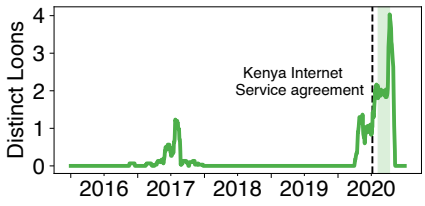

(f) Number of loons over Kenya.

Fig. 3: Presentation of the case studies. The three months with highest activity are marked with a shade.

2.1 Overall stats

The data-set developed for this work collects data starting from July 2016. Fig. 1 shows either the number of distinct balloons cumulatively detected over time (blue line) as well as the number of balloons simultaneously deployed in each day (red line). The number of operating loons has oscillated over the past years, with a peak in 2017 of more than 40 loons that was only recently reached and surpassed. This peak in 2017 roughly marks the end of the high pace of new loons identified (approx. 400 loons/year), which by the end of 2017 changed to a much smaller pace (approx. 70 loons/year in 2018 and 25 loons/year for 2019 and 2020). We note that loons typically operate for hundreds of days until they land, are recovered, refurbished (if possible) and put back in operation, so this total number of identifiers has to be over-dimensioned to account for the time until the transponder supporting the ADS-B system can be re-installed in a new loon.

To estimate the locations in which loons have operated, we tessellate the Earth into hexagonal pixels of approx. 86 000 km^2 (we used the Uber's `H3` library — see details in the next Section 3), and compute the average number of different loons per day tracked in each pixel. Fig. 2 reports a world map and the pixel in which at least one balloons was observed since July 2016. We then color each pixel with a proportional relative density-based color bar that is equal to the aforementioned number of different loons (we set this index to 1 for the pixel that experiences the largest

absolute value). By analyzing the figure, we can clearly see that there are some hot zones in **South and Central America** (with Peru and Puerto Rico being two of our case studies), as well as in **South West Africa** (with Kenya, our third case study, being the most recent deployment, in 2020). In addition, we remind that one of Loon's launching sites is located in Winnemucca, USA, thus justifying the hot area in **Nevada, North America**, Eventually, the activity in **Oceania** is primarily due to early tests performed by Google in New Zealand, indeed well known since they made the headlines — a balloon dropping in the sea was thought to be a crashing plane and created panic.[5]

Finally, for what concerns altitudes, our data suggest a roughly uniform distribution between 15 km and 20 km, while regarding loons movement, we measured a median ground speed of 25 km/h with some spikes above 80 km/h. When not providing coverage to a given area using relatively stationary patterns, loons move considerably faster to reach the intended destination, such as our considered case studies. For instance, from Ceiba (another launching site in Puerto Rico) it takes loon approx. 20 h to reach Peru , while it takes approx. 5 days to reach Kenya.

2.2 Case studies

To analyze the characteristics and trade-offs when deploying a mobile network with balloons, we focus on three main case studies located in Puerto Rico, Peru, and Kenya (depicted in Fig. 3, top), which have been extensively covered by media. Note that, for each case study, the areas under consideration do not exactly correspond to the whole country or the claimed provinces/regions appearing in the press releases, as we conservatively narrow each case study to the pixels with at least 30 days of "minimal coverage" ($>$ 2 hours per day) over the three months. For each case study, we count the number of unique loons per day over the area and represent its 30-day moving average (for ease of visualization) being depicted in Fig. 3, bottom. Unless otherwise noticed, our study focuses in the three months with the highest loon activity for each case, marked with a shade in the figure: from 2019-10-15 to 2020-01-15 for Puerto Rico, from 2017-11-01 to 2018-02-28 for Peru, and from 2020-07-15 to 2020-10-15 for Kenya.

Puerto Rico: Our first use case analyzes the loon performance over Puerto Rico, where the area under consideration mainly spans over the main island, as illustrated in Fig. 3a, and the number of loons per day are depicted in Fig. 3d. Obviously, the patterns in the figure can be mapped to the loon activity over the island that was mentioned in the press releases. More in depth, Hurricane Maria struck Puerto Rico on September 20, became extra-tropical on September 30, and dissipated by October 2 since 2017. Clearly, during this event, no loons have been observed in the area (as expected). Starting from Oct. 6 2017, Loon Inc. was authorized by the Federal Commission on Communications (FCC) to provide LTE coverage over Puerto Rico,[6] likely due to the loss of connectivity as a consequence of the damages that affected the terrestrial wireless infrastructure. On Oct. 21th 2017, Loon announced a partnership with AT&T and T-Mobile to deliver Internet service to the hardest hit parts of the island,[7] later tweeting that more than 100K people were provided basic Internet connectivity by Nov. 9th 2017.[8] Although the company did not report how many balloons it took to serve Puerto Rico,[9] our data shows that more than seven loons were simultaneously over the country during different periods of time. On March 2nd 2018, the company tweeted that the Loon service over the island would start to "wind down", since the connectivity was largely been restored.[10] This claim is also confirmed by the notch appearing in Fig. 3d after the shade. However, our data reveals that Loon activity over the country is sometimes experienced, likely because Ceiba (on the north-east coast) serves as a launch site for loons that then travel to other world destinations.

Peru: Loon and Telefónica started collaboration in 2014 with an early test of the service, and gained notable media attention with the emergency deployments after the flooding due to El Niño in May 2017[11] and the 8.0 magnitude earthquake in May 2019.[12] On Nov. 21st 2019, Loon announced an agreement with the neutral-host Internet Para Todos Peru (IpT) operator (owned by Telefónica, Facebook, IDB Invest and CAF) to serve parts of the Loreto Region in Peru.[13] According to our recorded data, the aim is to provide a sustained service over certain parts of the Loreto region, depicted in Fig. 3b. Like in the previous use case, we count the unique number of loons per day over the considered region and illustrate the 30-day moving average in Fig. 3e. Interestingly, some tests have been done before the earthquake in May 2019, which triggered a lot of loon activity over the area that seems to spike by the end of 2019. However, after the collaboration agreement was signed, the Loon activity seems to have decreased during 2020.

Kenya: On July 7th 2020, it was announced that Loon was already providing service in certain regions of the country, mainly over the area depicted in Fig. 3c. According to the announcement, the plan is to add loons to achieve a target of around 35 or more flight vehicles.[14] While some sources claimed that the capital Nairobi was also targeted,[15] according to our coverage study this is not the case. The 30-day moving average of the number of distinct loons over the region, illustrated in Fig. 3f, shows that it has been a long way for a loon to provide this service. Some activity was already taking place back in 2017, followed by no

5. https://www.theverge.com/2014/6/20/5826988/google-loon-balloon-crash-new-zealand

6. https://apps.fcc.gov/oetcf/els/reports/STA_Print.cfm?mode=initial&application_seq=80734

7. https://medium.com/loon-for-all/turning-on-project-loon-in-puerto-rico-f3aa41ad2d7f

8. https://twitter.com/Theteamatx/status/928729756534583296

9. https://spectrum.ieee.org/tech-talk/telecom/internet/how-project-loon-built-the-navigation-system-that-kept-its-balloons-over-puerto-rico

10. https://twitter.com/Theteamatx/status/969604066136375296

11. https://medium.com/loon-for-all/helping-out-in-peru-9e5a84839fd2

12. https://medium.com/loon-for-all/working-with-at-t-to-offer-a-global-connectivity-solution-in-times-of-disaster-450d8cb9a448

13. https://medium.com/loon-for-all/loon-signs-deal-to-bring-balloon-powered-internet-to-amazon-rainforest-region-in-peru-34696714976c

14. https://medium.com/loon-for-all/loon-is-live-in-kenya-259d81c75a7a

15. https://www.dezeen.com/2020/07/09/loon-balloon-powered-internet-service-kenya/

activity during the 2018 and 2019 years, despite on July 2018 it was announced that service would start in 2019.[16] However, it was not until March 2020 when the Kenyan government gave formal approval for balloons to operate in the country,[17] and then, a month later, two weeks of network testing were reported.[18] This activity is also confirmed by the spike in Fig. 3f.

3 METHODOLOGY

Our goal is to analyze the service characteristics provided by the loons over the different geographical regions introduced in the previous section. To estimate the performance provided by loons, we tessellate the territories using hexagons, which we name *pixels*, and assume that the service over each pixel is represented by the service provided at its center. Based on this methodology, we need to: (*i*) define how coverage and service levels are computed, (*ii*) select a proper pixel granularity to evaluate the coverage and service levels, and (*iii*) discretize the loon trajectory over time to enable the evaluation of the service over meaningful time periods. Clearly, we need to identify the proper resolution levels for both (*ii*) and (*iii*) to make the computation feasible while capturing all the significant changes of the service in both space and time.

In the following, we first detail the model that we use to evaluate the coverage and capacity at a given pixel and then select the proper space and time granularity based on this model.

3.1 Coverage and capacity evaluation

To evaluate the levels of coverage and service, we adopt a set of optimistic/simplifying assumptions in line with [1] (we discuss in Section 6 how to relax some of these assumptions), namely:

1) we focus on the downlink direction, as this is the first necessary condition to provide coverage,
2) loons always transmit at the maximum power to outdoor receivers,
3) given the loon altitude, no buildings or obstacles affect the path loss, so we have free-space propagation (thanks to the fact that loons are located at a non-negligible elevation w.r.t. the terrain);
4) given the operating frequency (800 MHz), there is no rain attenuation nor Doppler effect caused by relative movements, and
5) the signal from one loon do not interfere with the one emitted by other loons and/or with other terrestrial cellular network deployments. In other words, we assume a noise-limited system.

We note that this noise-limiting assumption is not only optimistic (which guarantees that our performance bounds constitute an upper bound on performance), but also sensible, since interference between loon traffic can be handled with standard mobile procedures (scheduling, handovers), while interference with other terrestrial networks should be sporadic (since loons are deployed over unserved areas) and can be circumvented with e.g. *ad hoc* negotiations of spectrum.[19] or freeing bandwidth for an emergency deployment[20]. With these assumptions, we evaluate the *best* coverage and capacity levels that can be achieved over a given pixel (the radio resource management policies applied to single users are intentionally omitted, since our work is tailored to the coverage and capacity evaluations over the set of pixels, rather than the set of users).

More formally, we consider a generic pixel p that is served by a single loon l at time t. Let us denote with $D^{\mathrm{V}}_{l,p,t}$ the altitude above ground of loon l at time t w.r.t. pixel p. In addition, the projection of the loon position over the ground level of the pixel p generates a point, whose distance from p is denoted as $D^{\mathrm{H}}_{l,p,t}$. The 3D distance $D^{\mathrm{3D}}_{l,p,t}$ is (obviously) equal to:

$$D^{\mathrm{3D}}_{l,p,t} = \sqrt{\left(D^{\mathrm{V}}_{l,p,t}\right)^2 + \left(D^{\mathrm{H}}_{l,p,t}\right)^2} \tag{1}$$

Moreover, let us introduce the elevation angle $\theta_{l,p,t}$, formally expressed as:

$$\theta_{l,p,t} = \tan^{-1}\left(\frac{D^{\mathrm{V}}_{l,p,t}}{D^{\mathrm{H}}_{l,p,t}}\right) \tag{2}$$

Fig. 4 reports an example of $D^{\mathrm{3D}}_{l,p,t}$ and $\theta_{l,p,t}$ over an hexagonal pixel p and two time slots. Given that we consider ground and altitude distances in terms of kilometers, we assume that the Earth curvature ($\sim$ 80 cm/km) does not affect our calculations. Clearly, since the position of loon l varies over time, both $D^{\mathrm{3D}}_{l,p,t}$ and $\theta_{l,p,t}$ also notably change, thus affecting the propagation conditions experienced on p.

The free-space based path loss $F_{l,p,t}$ is expressed as:

$$F_{l,p,t} = 20\log_{10}\left(D^{\mathrm{3D}}_{l,p,t}\right) + 20\log_{10}(f) + 20\log_{10}\left(\frac{4\pi}{c}\right) \text{[dB]} \tag{3}$$

where $f = 800$ MHz is the operating frequency (in accordance with [1]), and c the speed of the light.

Given $F_{l,p,t}$, the received power $P^{\mathrm{RX}}_{l,p,t}$ is then computed with a simple noise-limited link budget. More formally, we have:

$$P^{\mathrm{RX}}_{l,p,t} = P^{\mathrm{TX}}_{l} + G^{\mathrm{UE}} + G^{V}(\theta_{l,p,t}) - F_{l,p,t} \tag{4}$$

where P^{TX}_{l} is the loon transmission power, G^{UE} is the User Equipment gain, $G^{V}(\theta_{l,p,t})$ is the antenna elevation gain and $F_{l,p,t}$ is the already-introduced path loss term. In line with [1], we set: $P^{\mathrm{TX}}_{l} = 37$ dBm and $G^{\mathrm{UE}} =$ -10 dB (i.e, a UE loss is considered). On the other hand, $G^{V}(\theta_{l,p,t})$ is generated via an interpolation function that approximates the antenna pattern depicted in [1, Fig. 1].

We then introduce a minimum sensitivity to verify if a pixel p is covered by any loon at time t. More formally, let us introduce the parameter P^{SENS}, representing the lowest

16. https://medium.com/loon-for-all/bringing-loon-to-kenya-c1e7f65fefe4
17. https://medium.com/loon-for-all/loons-internet-balloons-approved-to-fly-in-kenya-fa36de6ad53a
18. https://medium.com/loon-for-all/take-me-to-kenya-23221c3f80d
19. https://indianexpress.com/article/technology/tech-news-technology/expanding-internet-reach-in-india-google-asked-to-revise-frequency-for-loon-white-spaces-to-connect-schools/
20. https://www.satcom.guru/2017/10/google-loon-floating-nest-of-radio.html

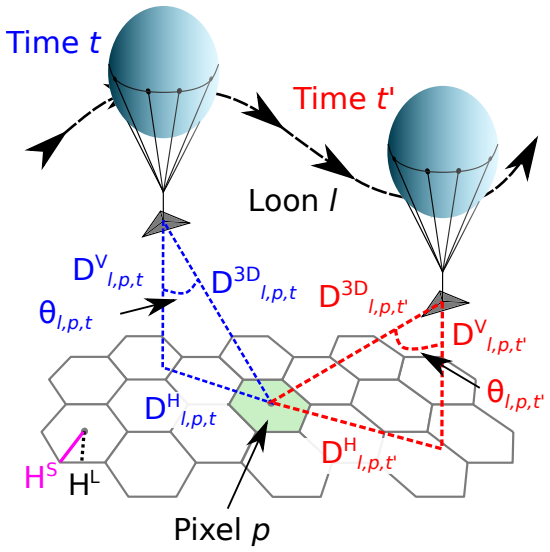

Fig. 4: Computation of the loon distance $D^{3D}_{l,p,t}$ and loon elevation angle $\theta_{l,p,t}$ for a single pixel as a function of the loon trajectory.

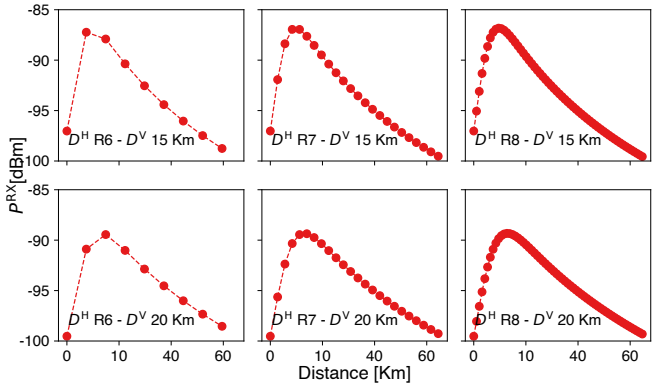

Fig. 5: $P^{RX}_{l,p,t}$ vs. the two selected values for $D^{V}_{l,p,t}$ and the sampled values of $D^{H}_{l,p,t}$.

amount of received power that can allow a communication from the loon to the pixel. Let us then denote with $C_{p,t}$ a coverage parameter, taking value 1 if p is covered by (at least) one loon at t and 0 otherwise. We then set the $C_{p,t}$ as:

$$C_{p,t} = \begin{cases} 1 & \text{if} \quad \max\limits_{l} \{P^{RX}_{l,p,t}\} \geq P^{SENS} \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

where the max operation selects the highest received power from all the loons. To set P^{SENS}, we note that authors in [1] considered two values for the *maximum coupling loss* between transmitter and receiver, namely, 140 dB and 132 dB. Such values correspond to two possible receiver sensibilities, i.e., -103 dBm and -95 dBm, respectively. For simplicity, we set P^{SENS} to -100 dBm, due to the following reasons: (i) this number falls inside the admissible interval selected by [1], and (ii) the same value was also measured as a good connectivity threshold in terrestrial LTE network deployments [13].

Finally, we compute the maximum capacity $R_{p,t}$ provided by the loon with the strongest signal by applying the well-known Shannon-Hartley model:

$$R_{p,t} = B_l \log_2 \left(1 + \frac{\max\limits_{l} \{P^{RX}_{l,p,t}\}}{N_0 \cdot B_l} \right) \tag{6}$$

where B_l is the loon operating bandwidth and N_0 is the noise floor, which are set to 10 MHz and -174 dBm/Hz, respectively. Clearly, we stress the fact that Eq. (6) represents an upper bound, which may bring to an over-estimation of coverage and capacity levels. Still, this methodology allows us to discuss the Loon deployment without making too stringent assumptions on the underlying wireless technology, and focusing on the distinctive characteristics of technology: their ability to move and provide coverage *on demand*.

3.2 Pixel tessellation and time discretization

Given that the evaluation of coverage/service levels has to be performed over a huge set of pixels and a very large number of time instants, it is of mandatory importance to properly discretize each region under consideration and also the periods of time.

Focusing on the territory discretization, we rely on Uber's `H3` library[21] to tessellate each region with a set of hexagonal pixels. H3 supports up to sixteen resolutions, ranging from a coarse resolution in which each hexagon has a side $H^S = 1\ 279$ km (index 0) to a very detailed tessellation in which $H^S = 0.58$ m (index 15).[22] To select the proper tessellation resolution, we compute the received power $P^{RX}_{l,p,t}$ with (4) for two different altitudes $D^{V}_{l,p,t} = \{15, 20\}$ km (in accordance with our measurements in Sec. 2) and for $D^{H}_{l,p,t} \in [0, 50]$ km, with the distance sampled in multiples of the hexagon diameter $2 \cdot H^S$ (bottom left part of Fig. 4). More in depth, we evaluate three different resolution levels: index 6 ($H^S = 3.72$ km), index 7 ($H^S = 1.4$ km), and index 8 ($H^S = 533$ m).

Fig.5 reports the obtained results over the different indexes and the different altitudes. Several considerations hold by analyzing the figure. First, $P^{RX}_{l,p,t}$ initially increases up to around 10 km and then decreases, due to the combination of the variation of the elevation gain $G^V(\theta_{l,p,t})$ and the contribution of the horizontal distance $D^{H}_{l,p,t}$ on the path loss. Interestingly, the resolution with index 6 appears to be too coarse, because the relative difference in the received power between one pixel and the neighboring one can be up to 5 dB. On the other hand, index 8 introduces a non-negligible number of pixels, with steps in the received power well below 1 dB. As a result, we select index 7 in order to balance between the variations of the received power among neighboring pixels and the pixel size.

In the following, we concentrate on the selection of the proper time resolution $\Delta_t = t - t'$ among two consecutive time slots. To this aim, we assume that during Δ_t the loon should not travel for more than $2 \cdot H^S$, with $H^S = 1.4$ km in

21. https://h3geo.org

22. The resolution table provided by Uber differentiates the levels based on the radius H^L of the circle inscribed to the hexagon. In this work, we adopt the side hexagon H^S rather than the radius H^L because we believe that the former is more intuitive than the latter. Clearly, it holds that $H^S = 2\ H^L/\sqrt{3}$.

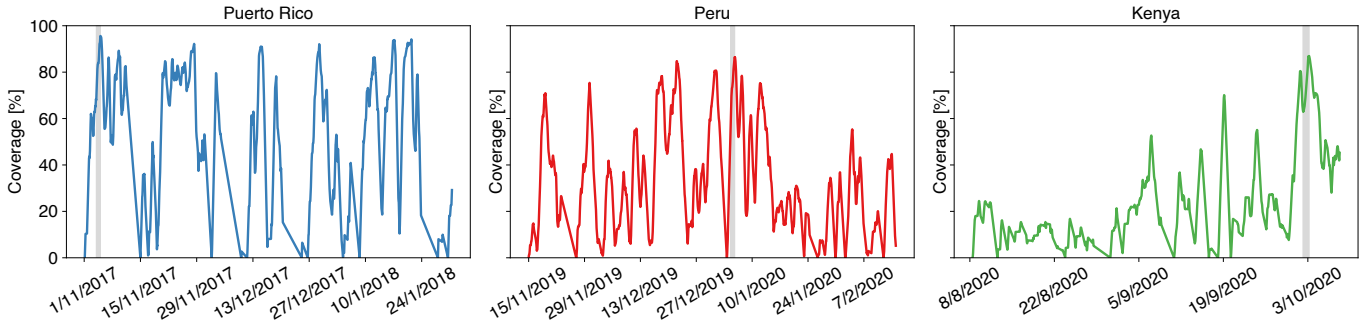

Fig. 6: Daily average coverage. Best day marked in gray.

accordance to the already introduced pixel tessellation. Following this, and taking into account that the 99$^{\text{th}}$ percentile of measured ground speed is approx. $V^{\text{99-PCT}} = 100$ km/h, Δ_t is bounded by:

$$\Delta_t \leq \frac{2 \cdot H^S}{V^{\text{99-PCT}}} = \frac{2.8 \text{ [km]}}{100 \text{ [km/h]}} \approx 1 \text{ min } 26 \text{ s} \qquad (7)$$

As a result, we conservatively fix $\Delta t = 1$ minute.

With the above, we perform the following time discretization of the trajectory of each loon: first, we round the timestamp of the raw data to minutes (since the accuracy is in seconds), averaging the latitude (lat), longitude (long), and altitude (alt) in case of multiple samples with the same timestamp. Then, we process the results to identify potential inter-sample gaps longer than a minute, to perform a linear interpolation of the lat, long and alt values. Note that we do not perform this interpolation for the gaps that are too long, i.e., more than 120 minutes between two consecutive loon positions (less than 0.5% of the cases). After the time discretization, we leverage on the pixel tessellation to compute, for each pixel and minute, the power received from each loon. Based on these, we determine if the pixel had coverage or not via (5) and the corresponding capacity via (6).

4 Performance Analysis

We next analyze the service provided by Loon in the three case studies presented in Section 2.2. For each of them, we report here the total areas and the corresponding number of pixels: Puerto Rico (9 052 km^2 and 3 604 pixels), Peru (46 157 km^2 and 7 716 pixels), and Kenya (27 800 km^2 and 4 457 pixels). Their representation in a map is shown in the top row of Fig. 3.

4.1 Coverage over time

Leveraging on our fine-grained dataset, we compute the % of pixels with coverage (i.e., the conditional expectation of $C_{p,t}$ over p) per minute for each case study, and depict the results in Fig. 6 using a daily moving average. The results illustrate that providing continuous coverage is a non-trivial task, to say the least. The scenario with the best coverage during the three months is **Puerto Rico**, with an average of 59.9% and frequent spikes above 90%, although it suffers some relatively long disconnections (e.g., the notch in Dec. 2017). The situation seems worse for the case of **Peru**, not only in terms of average coverage (38.8%), but also in terms of best performing days (i.e., the spikes), as in these the coverage barely reaches 80%. However, as compared vs. the Puerto Rico case, it seems that the periods of poor coverage or even complete disconnections are shorter, a result that we will formally analyze in Section 5. Finally, the case of **Kenya** shows an increasing improvement of the coverage over time (e.g., 57% during the last month), with very few and short periods of zero coverage. According to these results, then, providing continuous and complete coverage using loons seems a very challenging task. In the next section, we try to unveil some of the reasons behind this result.

4.2 Cost of coverage

We next analyze the means required to cover the use cases, i.e., the number of loons deployed over the considered areas. Note that since a loon covers an area of approx. 7 500 km^2, the absolute minimum amount of loons required for each scenario would be: 2 loons for Puerto Rico, 7 loons for Peru, and 4 loons for Kenya. To perform our analysis, we compute for every Δt the number of loons over the area and the % of pixels with coverage. We then compute the estimated density function of the coverage, conditioned to this number of loons, in Fig. 7, along with the marginal distribution of the number of loons.

For the case of **Puerto Rico**, the figure confirms the above "back of the envelope" calculation: with just one loon it is impossible to provide complete coverage, but with two loons the density function reaches 100%, although the median is approx. 75% and the coverage values spread over a wide span of results. With three loons, the median groups up to 85%, but since this point on the addition of more loons results in minor increases of (that slowly grows to 100%), with the main effect being the reduction of the dispersion of the results (the estimated density function gets narrower). We note that achieving $>$ 90% coverage requires up to 6 loons, while ideally two perfectly placed Loons could provide 100% coverage. We conjecture that this required "overprovision" the service (i.e., a factor of 3$\times$) could have impacted the sustainability of the project.

If we now consider the case of **Peru**, the results also confirm an improvement of the coverage with the number of loons. Given the larger area and its shape, the coverage reaches only 100% only when there are 8 or more loons over the area. Furthermore, it seems that there is little if any gain in bringing more loons into the area, as the median

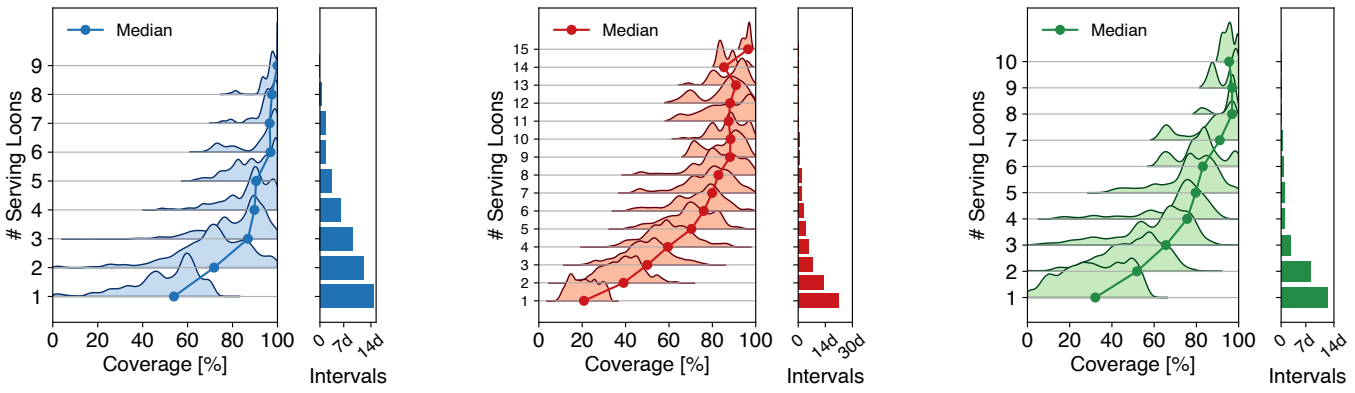

(a) Puerto Rico (b) Peru (c) Kenya

Fig. 7: Coverage as a function of the number of service loons during the best three months

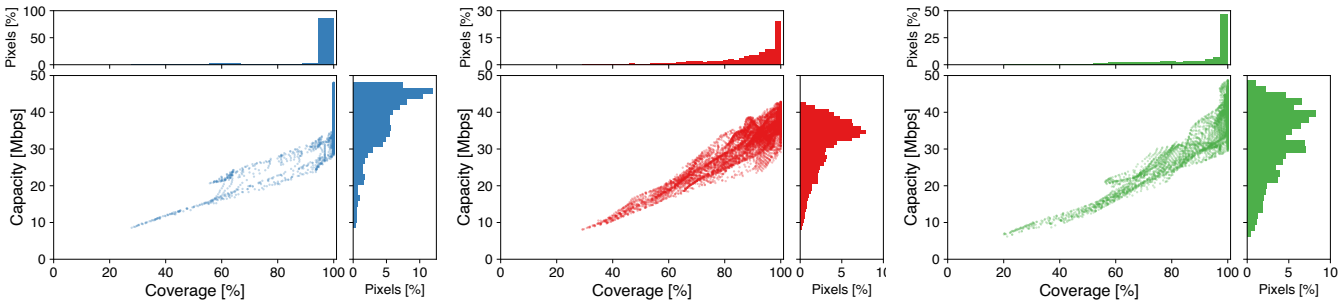

(a) Puerto Rico (b) Peru (c) Kenya

Fig. 8: Capacity $C_{p,t}$ vs. coverage $R_{p,t}$, best day

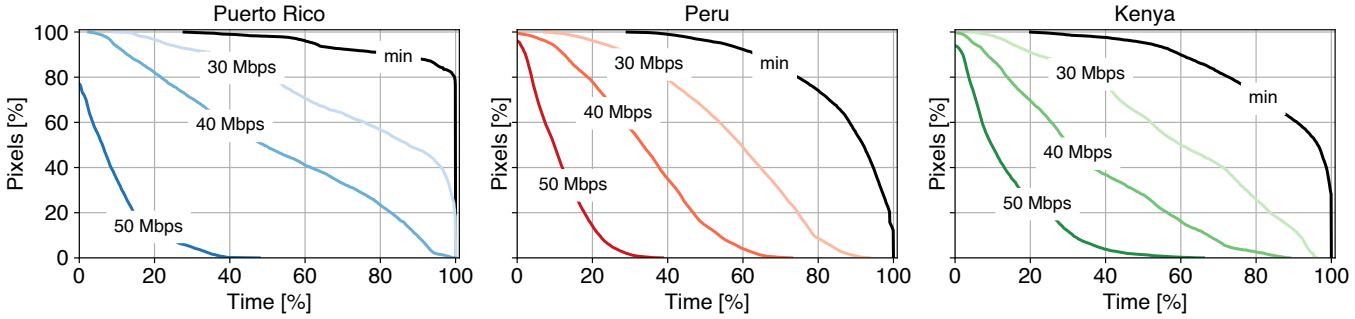

Fig. 9: Service guarantees for the best day

of the coverage and its span remains practically constant between 9 and 13 loons (note that according to the marginal distribution, these are very infrequent cases). These results suggest that the incremental value of adding extra loons is very small for more than 8 loons and that providing good coverage in this area is very challenging.

Finally, providing good coverage in **Kenya** seems more feasible than in Peru, despite its less regular shape. Starting from 5 loons, which is the minimum value to reach 100% coverage, the addition of new loons serves to improve the median coverage and narrows the dispersion of the results (still, note that these cases are relatively infrequent). However, like in the previous cases, the improvement is way more substantial before the complete coverage is first reached.

According to these results, coverage suffers from some sort of "law of diminishing returns," as the incremental value of an additional loon significantly reduces beyond the point where complete coverage is first reached. This effect may jeopardize deployments leveraging exclusively on this technology; on the other hand, the good coverage improvements brought by a few loons suggests that the Loon service could be used in conjunction with other technologies (e.g., terrestrial deployments in hotspots) to provide *umbrella cell*-alike service over very large areas. In any case, the ability to provide coverage over time will be revisited in Section 5,

where we will analyze the service (dis)continuity both over the whole area and its pixels.

4.3 Capacity distribution

We next look at the other metric, namely, the obtained capacity per pixel $R_{p,t}$. According to our results, there are significant periods with poor if any coverage during the considered three months periods. To obtain meaningful results, we decide to analyze a time window short enough that guarantees that the coverage is provided to at least 80% of the area (otherwise any performance figure would be significantly biased towards zero due to the poor coverage). As we will see in the next section, this time window is one single day.

We start our analysis by computing, for each considered pixel in the area, the average coverage and the average capacity during this best day (Nov. 4th 2017 for Puerto Rico, Jan. 4th, 2020 for Peru, and Oct. 2nd, 2020 for Kenya), which results in a `(coverage, capacity)` pair. We plot these pairs for all the pixels in each of the case studies in Fig. 8. Regarding **Puerto Rico**, the results confirm the remarkably good coverage obtained, as most of the points are located in the 100% coverage line (note that, according to Fig. 10b, the coverage during the best day is above 95%). For these pixels, the capacity ranges from a minimum of 28 Mbps to almost 48 Mbps, a result caused by the relatively high number of loons over the island (and therefore relatively good links), as we saw before.

The situation is remarkably different for the case of **Peru**: here, both the dispersion in terms of coverage is larger (note that the average coverage during the best day is approx. 85%) and, given a coverage value, the capacity distribution has a notable span, e.g., for 70% it varies between 20 Mbps and 30 Mbps, while for 90% it ranges between 22 Mbps and 38 Mbps. Furthermore, the maximum capacity is approx. 10% smaller than the one obtained in the Puerto Rico case. These results confirm that the coverage in this case is way more challenging, with several pixels experiencing relatively short periods of very good link quality (hence the large capacity for small periods of coverage).

The case of **Kenya** has a coverage similar to that of Puerto Rico, but with more dispersion in terms of capacity – note that the area is notably larger. Still, the capacity figures are remarkable, with an apparent multi-modal distribution. In fact, this is one of the few cases where the company recently reported user throughput of approx. 18.9 Mbps in downlink,[23] a result quite feasible (depending on the configured resource allocation) according to our estimations of the received power levels.

4.4 Service guarantees

We next analyze for how long a pixel is provided with a certain quality of service (i.e., basic coverage or a given capacity), which we refer to as *service guarantee*. To gather more insight into the distribution of resources across the areas, we analyze the number of pixels that are provided a given quality of service during the best day. More specifically, given a pixel and target capacity R, we compute the % of time that the pixel has a capacity larger than R. We then compute the % of pixels that were provided at least the above `(time, capacity)` service, and represent the results in Fig. 9 for R=$\{30, 40, 50\}$ Mbps. We also depict in black the results corresponding to just having coverage (`min`).

For instance, for the case of **Puerto Rico**, the black line illustrates that 100% of the pixels have coverage for at least 30% of the time, 90% of the pixels have at least 80% coverage, and 80% of the pixels have complete coverage (of course, the closer the black line to the top right corner, the better). As the service improves, the % of time reduces (i.e., lines get closer to the bottom left corner): the $R = 40$ Mbps shows an almost linear inverse relationship between time and area (e.g., 80% of the pixels have 40 Mbps for 20% of the time, and 20% of the pixels have that capacity for 80% of the time), while the $R = 50$ Mbps results very challenging and for very few pixels.

All pixels in **Peru** are also provided with coverage for at least 30% of the time. However, the minimum coverage line rapidly decreases with the time (i.e., only 15% of the pixels have coverage for the whole day). It is worth remarking, though, that despite the $R = \{30, 40\}$ Mbps lines are "Pareto better" in the case of Puerto Rico than in Peru, for the 50 Mbps line the performance in the case of Peru is very similar and certainly better for time guarantees smaller than 20%. In other words: more pixels in Peru experienced short bursts of very good capacity than in Puerto Rico. Finally, the case of **Kenya** shows a worse coverage than Puerto Rico and better capacity than Peru, with the same phenomena of short bursts of very good performance.

5 Continuity Analysis

Before we have have provided statistics about the Loon service over fixed period of times. We next analyze the *continuity* of the service, i.e., its availability (or lack of) over time.

5.1 Windowed availability

We start our analysis by borrowing the *minimal cumulative ratio (MCR)* metric (or windowed user-uptime) from a recent paper on "meaningful availability" from Google [14]. Given a period of interest $[T_1, T_2]$ and a time window x, the MCR_x is given as the least of the availabilities of all windows of size w, i.e.,

$$\text{MCR}_x \equiv \min_{T_1<t_1<t_2<T_2} \{A_{t_1,t_2} \mid t_2 - t_1 = x\} \tag{8}$$

where A_{t_1,t_2} denotes the availability during the interval $[t_1, t_2]$, which is defined as the ratio between the "good service" and the total service during that interval. The key idea behind this so-called "actionable" metric is to provide some insights, via the *knees* in the MCR_x curve, into the longest incidents and their distribution for the service. To analyze the coverage disruptions, we define the windowed minimal coverage $W^{\text{C-MIN}}(x)$ as

$$W^{\text{C-MIN}}(x) = \min_{T_1<t_1<t_2<T_2} \{N^{\text{COV}}_{t_1,t_2} \mid t_2 - t_1 = x\} \tag{9}$$

where $N^{\text{COV}}_{t_1,t_2}$ represents the sum of the pixels that had coverage during each of the Δt time intervals, divided by the total

23. https://medium.com/loon-for-all/loon-is-live-in-kenya-259d81c75a7a

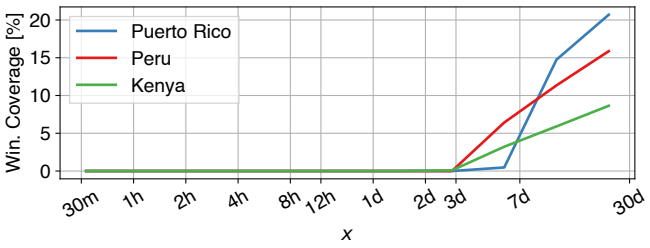

(a) Windowed minimal coverage $W^{\text{C-MIN}}(x)$

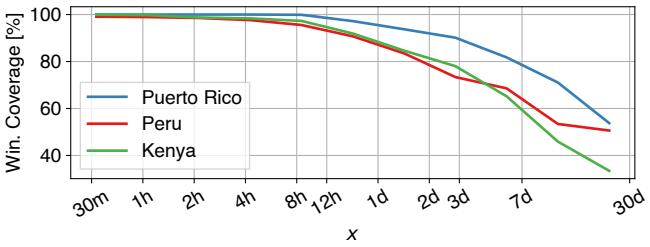

(b) Windowed maximal coverage $W^{\text{C-MAX}}(x)$

Fig. 10: Windowed coverage

number of intervals. If we denote these as $k = (t_2 - t_1)/\Delta t$, this is given by

$$N^{\text{COV}}_{t_1,t_2} = \frac{\sum_{i=0}^{k} \sum_p C_{p,t_1+i\Delta t}}{(t_2 - t_1)\,|p|} \tag{10}$$

where $|p|$ represents the cardinality of the pixel set.

We represent the windowed minimal coverage for the three case studies in Fig. 10a. The figure shows that all case studies experience at least one period of zero coverage for 1 day, and in Puerto Rico this lasted for several days (see the period 20–27/11/2017 in Fig. 6). However, for this case, the coverage rapidly increases with w, which illustrates that this long period is the main period of disconnection during the three months we analyze. In contrast, the slow increase in the Kenya case indicates the presence of other periods of poor coverage and of similar length (as Fig. 6 illustrates). The case of Peru falls between these two cases, with its corner in w happening before that of Puerto Rico (the longest service disruption lasts for one day), and steeper increase with w than Kenya (so its service disruptions are relatively shorted and less frequent).

The $W^{\text{C-MIN}}(x)$ metric, inspired by Google's "meaningful availability," provide us with some insights on the length and relative frequency of the periods with no service, i.e., *bad* performance. Since we are also interested in the *best* achievable performance with the Loon service, we define in a similar manner the windowed maximal coverage $W^{\text{C-MAX}}(x)$ as

$$W^{\text{C-MAX}}(x) \equiv \max_{T_1<t_1<t_2<T_2} \{N^{\text{COV}}_{t_1,t_2} \mid t_2 - t_1 = x\} \tag{11}$$

with the corresponding results depicted in Fig. 10b. The figure confirms the relatively good coverage service provided in Puerto Rico: not only in this case the window of continuous and complete coverage is the longest (8 h), but also the decay of $W^{\text{C-MAX}}(x)$ with x is not as steep as in the Peru or Kenya cases. In these two cases, the longest period of

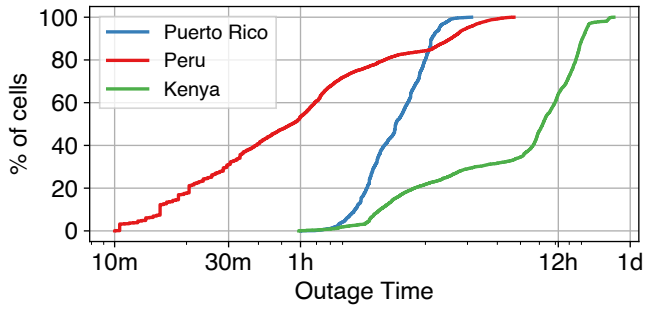

Fig. 11: Outage time

complete coverage approx. is 1 h, and decreases to approx. 90% for 12 h and to 80% for 2 days (which is the reason why we decided to limit the analysis in the previous section to the best day). From 2 days on, the decay of $W^{\text{C-MAX}}(x)$ is very steep in both cases, which illustrates how difficult it is to provide continuous coverage to these relatively large areas over "long" periods of time.

5.2 Outage time

The windowed minimal coverage defined above provides an overview of the downtime of the service for the whole surface analyzed, which is a meaningful metric from the perspective of a network operator, i.e., a macroscopic analysis. We next deep into the downtime by performing a microscopic analysis, i.e., characterizing the outage time of a pixel, which can be considered as a user-perspective metric. To this aim, we measure for each pixel all the downtimes, i.e., the length of the intervals without network connectivity, and compute the median of such values per pixel. We depict the empirical CDF of these medians in Fig. 11.

The results illustrate a relatively good service over Peru, as the median outage times range from tens of minutes to several hours. The cases of Kenya and Puerto Rico exhibit similar starting points of the empirical CDF (at approx. 4 hours), but Kenya has both longer and more disperse median outage times, with peaks that approximate one entire day (vs. the few hours in Puerto Rico). The case of Peru seems to have the shortest median times, but its heterogeneity is remarkable, as it spans over one order of magnitude; in contrast, the Puerto Rico case seems to be the fairest (thus the most predictable) as all the median outage times fall in the [1h,6h] range.

The results suggest that in the most optimistic conditions registered in Peru, even some media player application which relies on buffering to provide continuous media playback could be supported, as a non-negligible share of the outage times are below half an hour. This does not hold for the other two use cases, whose minimum median outage times are well above one hour (as shown in Fig. 10). In these cases, the service may reach larger continuity, but it does it in a more exacerbated *on-off* fashion, with long service up-times followed by long service disruptions.

6 Extensions and Discussions

As discussed until now, the main goal of this paper was to understand whether moving loons, irrespective of the

specific wireless technology involved on board, can provide a reasonable coverage and service continuity. In this section we challenge our results with selected examples of different configurations and deployment parameters, so as to assess the sensitivity of our results to different channel modeling assumptions and Loon mobility patterns.

6.1 Different channel models

Since our motivation is to assess if moving loons could provide a meaningful service, we have assumed that path losses correspond to those of free-space (3). But our methodology can be easily extended to account for more sophisticated channel models, as we exemplify in this section.

For instance, by following the design steps in [15, Section 5] we can add to (3) the "excessive path loss component" (η) caused by Line of Sight (LoS) or Non Line of Sight (NLoS) propagation in different environments. In our case, we assume a suburban scenario and the parametrization in [15, Table II] for $f = 700$ MHz (i.e., the closest to the Loon's operating frequency), and generate a random value of η for each pixel, which is then added to the path loss $F_{l,p,t}$. To assess the impact of this channel model, we compare the new median coverage with N loons (denoted as "Suburban") against the previous one ("Original") in Fig. 12. According to the results, there is a small decrease in coverage for the Puerto Rico and Kenya cases, while the performance for the Peru case remains very similar. The small impact on performance was expected, as the elevation angle is typically very large since the loon altitude is high (for smaller elevation angles, where NLoS is more likely, the path is already very long). Only in Puerto Rico and Kenya, where it is more likely to have loons closer to the borders of the coverage area (see the coverage with $N = 1$ in Fig. 7, where in both Puerto Rico and Kenya the coverage reaches 0, but not in Peru), the elevation angles are smaller and the impact of the extra losses is more noticeable.

Similarly, we could follow [16, Section 2.2.1.1] to assess the impact of rain on performance, by computing the specific attenuation factor γ_R and the effective path length L_E for each pixel, which would result in an extra attenuation factor $\gamma_R L_E$. However, given the (relatively low) operation frequency of 800 MHz, γ_R is smaller than 0.002 dB/km while L_E is at most several kms, which results in a negligible total attenuation –this is not the case for higher operation frequencies, as discussed in [17].

6.2 Comparison vs. a random deployment

Since the actual Loons' movements are the result of the combination of wind conditions and complex proprietary strategies (changes of altitudes so as to exploit different wind conditions), a thorough and realistic Loos' mobility model seems an hardly attainable challenge. Still, it is fair to ask how *different* would our results become, if Loons were *randomly* deployed on the territory - indeed a random Loon placement model (e.g. a Poisson Point Process) would in fact yield easy analytical tractability. To answer such question, in this section we compare the actual real world coverage with that of a random deployment over the same area. More specifically, given a number of loons N, we run 1000 experiments consisting of (1) deploying each of the N loons over a randomly selected pixel of the area, at a randomly selected altitude between 15 km and 20 km, and (2) computing the resulting coverage over the area, following our methodology. We plot the median coverage with N loons from this random deployment, which mimics a Poisson Point Process, vs. the median coverage of the Loon service with N loons in Fig. 13.

According to the results, in practically all cases the loon coverage consistently outperforms that of a random deployment with the same number of loons, since most of the points are below the $y = x$ line –in particular, when N is relatively low. Although we do not know the exact areas that were targeted and therefore we might have introduced some bias that favors loon (we only consider pixels with > 2 hours of coverage), we are confident that these results confirm the ability of Alphabet's engineers to place and maintain HAPs over the geographical region of interest despite the changing weather conditions and lack of propulsion engine. They also confirm, however, that the improvements are moderate (10% at most), which also challenges the efficiency of such type of deployments to provide a network service. This similarity in performance suggest that tools from stochastic geometry [18] could be used to characterise the performance of the service.

6.3 Other extensions

Uplink traffic: applications generating significant uplink traffic are becoming dramatically popular (e.g., user generated video). Our model can be extended to analyze the performance bounds in this scenario, by assuming a similar channel model but taking into account the lower transmission power of phones (a typical maximum transmission power of 23 dBm). Provided with the (downlink,uplink) pairs over time, it would be possible then to use the tools presented in our methodology (e.g., windowed availability, outage time) to determine the type of applications that could be supported by the service. This constitutes part of our future work.

Heterogeneous networks: Finally, our results illustrate that a deployment based exclusively on loons might not provide a sustainable service. Still, given the unique properties of this technology (endurance, coverage, etc.), they could be integrated in *heterogeneous* networks composed by different types of UAVs (drones, tethered ballons, etc.), such as the ones as discussed in [17], [19]. To this aim, our dataset could be used, for instance, to characterize more accurately the links between the different types of UAVs [19], or to assess the accuracy of the "quasi stationary" position of the loons [17]. On the other hand, the apparent closeness of loon coverage to that of a PPP (Fig. 13) suggests that these could be used as a first approximation to model the performance of moving macrocells, while e.g. rotary-wings UAVs serve to deploy targeted microcells.

7 Related Work

The use of UAVs to provide cost-effective connectivity has been receiving growing attention from the research community due to its opportunities and challenges [20]. In general, low-altitude platforms (LAPs) consisting of drones

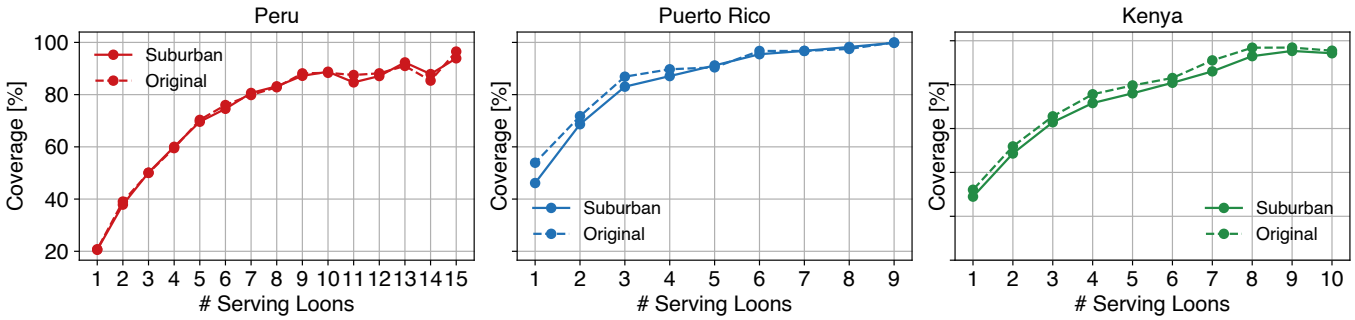

Fig. 12: Impact of non line of sight propagation on performance

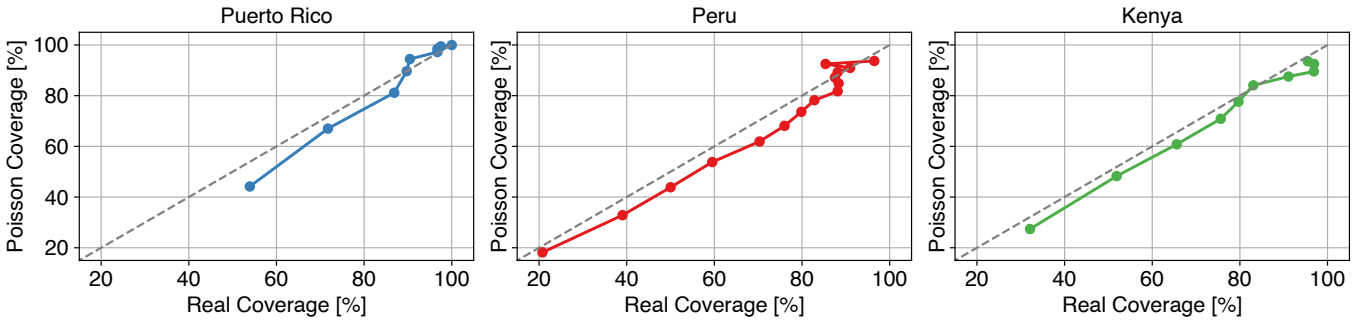

Fig. 13: Performance comparison of Loon coverage vs. a random deployment.

are preferred over high-altitude platforms (HAPs), since the deployment of these results much more challenging (including the potential interference with existing networks) despite their wider coverage and longer service lifetimes. To the best of our knowledge, the scientific literature has almost exclusively focused on LAPs, while the deployment of HAPs has been pursued by Internet companies such as Google and Facebook (see, e.g., the recent survey [8]). Some recent proposals advocate for the use of heterogeneous networks (e.g. drones and loons) [21], while others propose the use of *tethered* platforms [22] that are provided with a continuous supply of power and data.

In fact, most of the existing work on the Loon project has been produced or motivated by Loon itself. For instance, the design of its physical layer is described in [1], along with some experiments that show good coexistence figures with other networks. This good coexistence is also the main focus of the recent performance study performed by SRG, which is commissioned by Loon [12]. The inter-loon link, based on free-space optical communications, is analyzed in [2], demonstrating full-duplex 130 Mbps over 100 km distances. Finally, Loon researchers have participated in studies comparing the actual UAVs trajectories vs. those obtained using different weather forecasting models [10].

Finally, we can highlight two additional contributions: a simulation study of passive antennae to improve Loon's indoor coverage [9], although the numerical values of key parameters differ from those listed in [1] (e.g., an operation frequency of 2.6 GHz instead of 700 MHz); and a student's final project presenting an economic analysis of the feasibility of the service [23], based on a number of questionable assumptions (e.g., overly small coverage areas).

8 Conclusions

In this paper, we have carried out an analysis, entirely based on empirical evidence (data from publicly available flight tracking services), of the extent to which balloons might provide a reliable connectivity service. Even if it is fair to say that no specific claims have been so far made by the company providing the service, we were curious to see whether a spatially and temporally stable coverage was reached in some of the previous deployments, and — if not — what were the so-far accomplished service continuity trade-offs.

Since balloons are obviously non-stationary, and since our goal was to understand how *good* the service could eventually be, we carefully selected the terrestrial areas under investigation among the most covered ones, and we restricted our analysis on time windows exhibiting the highest number of transiently deployed balloons. Our results bring about mixed feelings. On one side, the results confirm that balloons are certainly a cost-effective way to provide a better-than-nothing, delay-tolerant, connectivity service, with outage periods that may grow to more than a few hours. On the other side, there is yet no empirical evidence that an increase in the number of overlapping balloons may be rewarded with a substantial increase in the service performance — in other words, we suspect that a coverage and service stability level comparable to that of a planned network might not be attainable. We cannot but conjecture whether this result could be one of the main reasons for the shutting down of the Loon project (as mentioned in the Peru case, a $3\times$ over provision). In any case, it is clear that given their unique properties in terms of coverage and lifetime,

loons could enrich the wireless communication networks when used in combination with e.g. satellites or drones, as discussed above.

Finally, it is worth to remark that our paper brings about a substantial contribution in terms of methodology. Indeed, being the first (to the best of our knowledge) to attempt to provide such an analysis, we needed to overcome a significant number of hurdles in terms of data gathering, spatio-temporal tessellation, coverage and radio connectivity modeling, and of course identification of key performance indicators and service availability metrics which could make sense in such rapidly changing deployment scenarios. We believe that such work may be useful to our scientific community, and perhaps also for different scopes.

ACKNOWLEDGEMENTS

This work has been partly funded by the Spanish State Research Agency (TRUE5G project, PID2019-108713RB-C52/AEI/10.13039/501100011033) and by the Madrid Government (Comunidad de Madrid-Spain) under the Multiannual Agreement with UC3M in the line of Excellence of University Professors (EPUC3M21), and in the context of the V PRICIT (Regional Programme of Research and Technological Innovation). We thank Yan Grunenberger for his extremely valuable feedback.

REFERENCES

[1] S. Ananth, B. Wojtowicz, A. Cohen, N. Gulia, A. Bhattacharya, and B. Fox, "System design of the physical layer for loon's high-altitude platform," *EURASIP Journal on Wireless Communications and Networking*, vol. 2019, no. 1, Jun. 2019. [Online]. Available: https://doi.org/10.1186/s13638-019-1461-x

[2] B. Moision *et al.*, "Demonstration of free-space optical communication for long-range data links between balloons on Project Loon," in *Free-Space Laser Communication and Atmospheric Propagation XXIX*, vol. 10096. SPIE, 2017. [Online]. Available: https://doi.org/10.1117/12.2253099

[3] X. Huang, J. A. Zhang, R. P. Liu, Y. J. Guo, and L. Hanzo, "Airplane-aided integrated networking for 6g wireless: Will it work?" *IEEE Vehicular Technology Magazine*, vol. 14, no. 3, pp. 84–91, 2019.

[4] H. Saarnisaari *et al.*, "A 6G White Paper on Connectivity for Remote Areas," Tech. Rep., April 2020. [Online]. Available: https://arxiv.org/pdf/2004.14699.pdf

[5] Gunes Kurt et al., "A Vision and Framework for the High Altitude Platform Station (HAPS) Networks of the Future," Tech. Rep., July 2020. [Online]. Available: https://arxiv.org/pdf/2007.15088.pdf

[6] I. Bor-Yaliniz and H. Yanikomeroglu, "The new frontier in ran heterogeneity: Multi-tier drone-cells," *IEEE Communications Magazine*, vol. 54, no. 11, pp. 48–55, 2016.

[7] V. Sharma, M. Bennis, and R. Kumar, "Uav-assisted heterogeneous networks for capacity enhancement," *IEEE Communications Letters*, vol. 20, no. 6, pp. 1207–1210, 2016.

[8] A. Fotouhi *et al.*, "Survey on UAV Cellular Communications: Practical Aspects, Standardization Advancements, Regulation, and Security Challenges," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3417–3442, 2019.

[9] P. Lynggaard, "Improve Google Loon´ s Indoor LTE Coverage in Rural Africa by Using Passive Repeaters," *Nordic and Baltic Journal of Information and Communications Technologies*, vol. 2017, pp. 91–106, 01 2017.

[10] L. S. Friedrich *et al.*, "A comparison of loon balloon observations and stratospheric reanalysis products," *Atmospheric Chem. Phys*, vol. 17, no. 2, pp. 855–866, Jan. 2017. [Online]. Available: https://doi.org/10.5194/acp-17-855-2017

[11] M. Mulhern, "Project loon," GREPECAS PPRC/4, Tech. Rep., July 2016. [Online]. Available: https://www.icao.int/SAM/Documents/2016-CRPP4/Peru%20GREPECAS%20PPRC-4%20brief%20Jul%2016.pdf

[12] "White paper: Loon and a terrestrial LTE network," Signals Research Group, Tech. Rep., February 2020. [Online]. Available: https://signalsresearch.com/issue/loon-and-a-terrestrial-lte-network/

[13] L. Chiaraviglio, J. Galán-Jiménez, M. Fiore, and N. Blefari-Melazzi, "Not in my neighborhood: A user equipment perspective of cellular planning under restrictive emf limits," *IEEE Access*, vol. 7, pp. 6161–6185, 2019.

[14] T. Hauer, P. Hoffmann, J. Lunney, D. Ardelean, and A. Diwan, "Meaningful availability," in *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. USENIX Association, Feb. 2020, pp. 545–557. [Online]. Available: https://www.usenix.org/conference/nsdi20/presentation/hauer

[15] A. Al-Hourani, S. Kandeepan, and A. Jamalipour, "Modeling air-to-ground path loss for low altitude platforms in urban environments," in *2014 IEEE Global Communications Conference*, 2014, pp. 2898–2904.

[16] *Propagation Data and Prediction Methods Required for the Design of Earth-Space Telecommunication Systems*, ITU-R Std. P.618-13, December 2017.

[17] Y. Huo, X. Dong, T. Lu, W. Xu, and M. Yuen, "Distributed and multilayer uav networks for next-generation wireless communication and power transfer: A feasibility study," *IEEE Internet of Things Journal*, vol. 6, no. 4, pp. 7103–7115, 2019.

[18] F. Baccelli and S. Zuyev, "Stochastic geometry models of mobile communication networks," in *Frontiers in queueing: models and applications in science and engineering*. CRC Press, 1996, pp. 227–243.

[19] P. G. Sudheesh, M. Mozaffari, M. Magarini, W. Saad, and P. Muthuchidambaranathan, "Sum-rate analysis for high altitude platform (hap) drones with tethered balloon relay," *IEEE Communications Letters*, vol. 22, no. 6, pp. 1240–1243, 2018.

[20] Y. Zeng, R. Zhang, and T. J. Lim, "Wireless communications with unmanned aerial vehicles: opportunities and challenges," *IEEE Communications Magazine*, vol. 54, no. 5, pp. 36–42, 2016.

[21] J. Qiu, D. Grace, G. Ding, M. D. Zakaria, and Q. Wu, "Air-ground heterogeneous networks for 5g and beyond via integrating high and low altitude platforms," *IEEE Wireless Communications*, vol. 26, no. 6, pp. 140–148, 2019.

[22] B. E. Y. Belmekki and M.-S. Alouini, "Unleashing the potential of networked tethered flying platforms for b5g/6g: Prospects, challenges, and applications," 2021.

[23] J. Burr, "The feasibility of Google's project Loon," ENGN2226 Systems Engineering, Student Work, Australian National University, Tech. Rep., 2015. [Online]. Available: https://users.cecs.anu.edu.au/~u3951377/student_work/example_work/15_2226_lp_jamesb.pdf

ANNEX: METHODOLOGY

Loon identification

To comply with the aircraft tracking standards, each Loon is equipped with a Mode-S transponder, a secondary radar system used to support the Automatic Dependent Surveillance–Broadcast (ADS-B), a satellite based surveillance system. The transponder is provided with a unique International Civil Aviation Organization (ICAO) address. Thanks to ADS-B, an aircraft can constantly broadcast messages that include, along with its 24-bit ICAO address, an estimation of its position, speed, altitude, as well as the *callsign* corresponding to the flight. This callsign, also known as the flight number, is intended to uniquely identify an airline service (e.g., a regular flight between two airports). For the case of the Loon service, it has been reported that flight numbers start with `HBAL`.[24] ADS-B messages can also be easily received with cheap ground receivers, such as, e.g., DVB-T USB sticks, which has fostered a number of crowd-sourced initiatives such as, e.g., FlightRadar24, RadarBox, or adsbScope.

24. See, e.g., https://www.flightradar24.com/blog/keep-your-eye-on-the-hbal-tracking-project-loon-balloons/

Data gathering

As discussed in Section 2.2, we identified relevant areas by looking at the news related to the Loon's deployment, which are corroborated *a posteriori* by the high number of flying loons in the area during the considered period. We initially defined the area using publicly available shapefiles of the considered regions, extending them approx. 50 km to ensure that relevant loons are captured. We gathered the data used in this work by harvesting one of the major aircraft tracking websites, for which we hold a gold subscription that allows us to browse data from the last three years. By using their search capabilities, we identify the flights that match a given aircraft ID and download the data in JSON format.

We identified a number of issues in the resulting "raw" dataset, such as missing flight identifiers or aircraft identifiers reporting almost-simultaneous flights in locations very far apart. To *sanitize* the data base, we filter out all registers outside the stratosphere, as well as locations associated to aircrafts flying at too high speeds (> 100 km/h), and rely on the uniqueness of the `ICAO address` to distinguish between Loons (as aircraft identifiers are not 100% reliable). We collected a total flight time of 305.343 days, yielding more than 8 million km, with flight spans that reach almost half of the circumference of the earth (the longest recorded flight in our database was 18.000km long).

Data processing

Following Section 3.2, we process the loon data to have an accuracy of one minute in the timestamps. For each loon, we convert the units of the timestamp value to minutes, removing the number of seconds. In case several entries have the same timestamp, we substitute them with a single entry consisting of the corresponding arithmetic mean of the latitude, longitude, and altitude. If case there is a time gap between two consecutive positions of the same loon longer that is longer than a minute, we linearly interpolate its position, but only for those time gaps shorter than two hours –in this case, we assume that the loon was not providing service during the time gap.

Due to the size of the dataset, in the order of hundreds of millions of rows, we performed the computations using HPC clusters and Spark. To foster the replicability of our findings, and to allow researchers to perform other analyses, the whole dataset consisting of all the $C_{p,t}$ and $R_{p,t}$ values per pixel and minute during the three months (6 GB of data) is published at https://zenodo.org/record/5676213.

Pablo Serrano (M'04, SM'15) has an Associate Professor with the Telematics Department of Univ. Carlos III de Madrid since 2002. He has held visiting positions at University of Massachusetts Amherst, University of Edinburgh, Trinity College Dublin, University of Rome Tor Vergata, University of Brescia, and Telefónica R+D Barcelona. He has over 100 scientific papers in peer-reviewed international journal and conferences. He regularly serves TPC member of a number of conferences and workshops. He serves as an Associate Editor of the IEEE Open Journal of the Communications Society. He is a Senior Member of the IEEE.

Marco Gramaglia received the M.Sc. and Ph.D. degrees in telematics engineering from the University Carlos III of Madrid (UC3M), in 2009 and 2012, respectively. He held post-doctoral research positions at ISMB, Italy, CNR-IEIIT, Italy, and IMDEA Networks, Spain. He is currently a PostDoctoral Researcher at UC3M. He was involved in EU projects and has authored more than 50 papers appeared in international conference and journals.

Francesco Mancini received his Master's Degree in Computer Engineering from the University of Rome Tor Vergata (Italy), in October 2021. He was a researcher at CNIT (National Inter-University Consortium for Telecommunications), in the research unit of the University of Rome Tor Vergata, from 2019 to 2021. He is currently a PhD student in Electronics Engineering at the same university. His research interests include network and computer security.

Luca Chiaraviglio (SM) received the M.Sc. degree in Computer Engineering and the Ph.D. degree in Telecommunication and Electronics Engineering from the Polytechnic of Turin, Italy. He is Associate Professor with the University of Rome Tor Vergata, Italy. In the past years, he held research/visiting positions with Boston University (USA), INRIA (France), Auckland University of Technology (New Zealand) and University of Rome Sapienza (Italy). He has co-authored over 150 papers published in international journals, books and conferences, resulting from the collaboration with over 200 coauthors. His current research topics cover 5G and B5G networks, electromagnetic fields, optimization applied to telecommunication networks, and health risks assessment of 5G communications. He has received the Best Paper Award at IEEE VTC (twice) and ICIN conferences. Some of his papers are listed as Best Readings on Green Communications by IEEE. He has been recognized as an author in the top 1% most highly cited papers in the ICT field worldwide. He participates in the TPC of IEEE INFOCOM, IEEE GLOBECOM, IEEE ICC, IEEE VTC. He serves on the editorial board of IEEE Communications Magazine. He is also the founding Specialty Chief Editor for the Networks Section of Frontiers in Communications and Networking.

Giuseppe Bianchi has been a Full Professor of networking with the School of Engineering, University of Roma Tor Vergata, since January 2007, and he further collaborates with CNIT. His research interests include programmable network systems, wireless networks, privacy and security, traffic control, and is documented in about 300 peer-reviewed international journal articles and conference papers, having received more than 20 000 citations to date. He has held a general or technical coordination roles in six European projects, has served as an Editor for the IEEE/ACM Transactions on Networking, the IEEE Transactions on Wireless Communications, the IEEE Transactions on Network and Service Management, and Computer Communications (Elsevier).