The carbon footprint of distributed cloud storage

Lorenzo Posani[a,b,c], Marco Moschettini[a], Alessio Paccoia[a], Giacomo Venezia[a]

[a] *Cubbit Research & Development*

[b] *Center for Computational Neuroscience, Columbia University, New York (NY)*

[c] to whom correspondence should be addressed, email: lorenzo.posani@cubbit.io

Abstract

The ICT (Information Communication Technologies) ecosystem is estimated to be responsible by 2040 for almost 15% of the total worldwide energy demand and data centers will account for about 30% it [1]. Cloud storage, mainly operated through large and densely-packed data centers, constitutes a non-negligible part of this energy demand. However, since the cloud is a fast-inflating market its carbon footprint shows no signs of slowing down. In this paper, we analyze a novel paradigm for cloud storage (implemented by cubbit.io) [2, 3], in which data are stored and distributed over a network of p2p-interacting ARM-based single-board devices. We compare Cubbit's distributed cloud to the traditional centralized solution in terms of environmental footprint and energy efficiency. We demonstrate that, compared to the centralized cloud, the distributed architecture of Cubbit reduces the carbon footprint of cloud storage by 75% and that of data transfer by 61%. These results provide an example of how a radical paradigm shift in a large-reach technology can benefit both the final consumer as well as society as a whole.

Keywords: Carbon footprint, Cloud Storage, Distributed, Peer-to-peer

1. Introduction

Over the last decades, the general acknowledgment of the climate crisis has driven the world towards an increasing awareness on environmental sustainability. In 2015, 196 Parties signed the Paris Agreement and committed to limiting global warming to below 2 degrees Celsius compared to pre-industrial levels [4]. Europe, the forerunner when it comes to environmental transition, put in place the Green New Deal, an ambitious plan to become the first climate-neutral continent with investments worth over 600 billion euro [5]. But the achievement of the goals of the Paris Agreement is still far away [6].

One of the most impactful sectors on climate will be ICT, that with its IT devices, data centers and communication networks will be responsible for almost 15% of global emissions by 2040 [1]. If emissions related to the electronics market are clearly visible and measurable, the impact of our online life is much less tangible and much more controversial. Every time a video is streamed from Youtube servers to an iPad, or a photo is accessed on Google Photos or Dropbox, the whole infrastructure that separates the final user from the corporate data center, as well as the data center itself, has to be powered to reliably transmit information in both directions. Also, the distance between the exchanging nodes has an influence on the energy needed to transmit information.

Name	Capacity	Wattage (peak)	PUE	Redundancy
Cubbit Cell	4 TB	4.7 W	1.0	2.0
DELL PowerVault ME484	806 TB	2200 W	1.57	2.0
HPE StorageEasy 1660	504 TB	1600 W	1.57	2.0
SuperMicro SuperStorage 6049P	380 TB	1200 W	1.57	2.0

Table 1: **Equipments** - power and capacity of storage equipments

Until a few years ago the impact of digital activities on the environment was largely underestimated, whilst the covid-19 pandemic shed light on it. During 2020, in fact, the increase of videoconferencing and gaming as well as the rise of content economy led global internet traffic to increase by 40%, 15 times bigger than ten years before [7]. Environmental sustainability concerns and awareness that data are stored in data centers put more pressure on the cloud storage market and researchers expect that, by 2025, data centers will be responsible for 3% of the CO2 global emissions [8].

This document aims to measure the carbon footprint of cloud storage services using an adaptation of the model of Baliga et al. [9], and compare it to an alternative setup where data is stored on peer-to-peer low-consumption devices located in users' houses and implemented by Cubbit [2, 3].

2. Analysis of centralized cloud consumptions

The Green Houses Gas (GHG) emissions of a centralized data center are associated to 5 different phases: design and construction, mechanical and electrical fit out, IT fit out, operations, decommissioning. Each of these phases account for different levels of GHG emissions generated over different time spans. In this paper, we focus on the operations phase, which accounts for about 80% of the GHG emissions during the entire life cycle of a centralized data center of approximately 25 years [10]. We estimate the direct energy consumption of a cloud storage service, that can be divided into two main factors:

1. the cost of storing the data, i.e. powering and cooling the data center (Storage consumption)
2. the cost of sending the data from the user to the server and back (Transfer consumption)

While the first can be estimated from technical specifications of storage equipment, the second needs a more detailed analysis that takes into account the public internet infrastructure and the geographical distance between the user and the server. For both these estimations we refer to the model of Baliga et al. [9], where energy consumption is computed accounting for several factors, including the multiplicity of involved devices, redundancy, cooling, overbooking (see below). To delineate the calculation, we start from the storage consumption, i.e. the average power, expressed in W/TB, necessary to store the payload in *hot storage*. We updated the technical specifications with respect to [9], as hard-disk storage capacity has dramatically improved in the last years.

As a model for data center racks, we considered three of the most recent products from three leading companies in the sector of enterprise storage hardware, see Tab.1. For each storage appliance we take the peak consumption, as it is the information made available in the manufacturer specs sheet [11, 12, 13]. To estimate the capacity, we consider every rack to be filled with 10TB disks (estimated from the mean of HDD dimensions in a recent BackBlaze report [14], $\simeq$ 10.8 TB/Disk) that are fully employed to store customers' cloud files (with no empty space overhead).

Name	P storage
PowerVault ME484	8,57 W/TB
HPE StorageEasy 1660	9,97 W/TB
SuperMicro SuperStorage 6049P	9,92 W/TB
Mean	9,49 W/TB

Table 2: **Equipments** - storage consumption

	Equipment	Capacity	Power
Data Center gateway router	Juniper MX-960	660 Gb/s	5.1 kW
Ethernet Switch	Cisco 6509	160 Gb/s	3.8 kW
BNG	Juniper E320	60 Gb/s	3.3 kW
Provider Edge	Cisco 12816	160 Gb/s	4.21 kW
Core router	Cisco CRS-1	640 Gb/s	10.9 kW
WDM (800 km)	Fujitsu 7700	40 Gb/s	136 W/channel

Table 3: **Equipments** - power and capacity of routing equipments. Data from [9]

We consider an average of 1.57× of Power Usage Effectiveness [15], estimated by the Uptime Institute in the latest report Uptime Institute Global Data Center Survey 2021. To consider redundancy, we use a 2× factor [9] as if every file was mirrored two times in the same or in a different data center, as described in the Baliga Model. We consider this to be a good estimate since it is consistent with the average redundancy protocols used by hyperscale structures, that can rely upon very strict security levels and therefore require low redundancy overhead [16, 17]. These values are simply combined in the following formula to obtain the storage consumption per Terabyte:

$$\mathrm{P}_{\mathrm{dcenter}}^{\mathrm{storage}} = \mathrm{PUE} \times \mathrm{redundancy} \times \frac{\mathrm{peak\ Wattage}}{\mathrm{n\ disks} \times 10\ \mathrm{TB}} \tag{1}$$

Likewise, we compute the transfer energy, expressed in J/GB, following the public internet model of [9]. The analysis relies on the definition of the consumption per bit, which is computed by dividing the operating power (W) by the total transfer capacity (Gb/s), resulting in a Joule/bit measure, then converted in J/GB. These units are taken from the manufacturer specs sheet, shown in table 1. These quantities are then combined with a set of coefficients reflecting the redundancy of the packet transmission, the under-operating regime of the infrastructure and the cooling energy, and the multiplicity of some devices in a single transmission (e.g. two ethernet switches at entry points plus another one inside the data center). The average distance between core routers on the network is estimated to be c.a. 800Km. For a full description of coefficients and estimations we refer the reader to [9].

$$
\begin{aligned}
\mathrm{E}^{\text{transfer}}_{\text{dcenter}} &= \\
&= 6.28\times\left(3\frac{P_{es}}{C_{es}}+\frac{P_{bg}}{C_{bg}}+\frac{P_g}{C_g}+2\frac{P_{pe}}{C_{pe}}+18\frac{P_c}{C_c}+4\frac{P_w}{C_w}\right) \\
&\simeq 30.6\ \frac{\text{kJ}}{\text{GB}}\ ,
\end{aligned} \tag{2}
$$

where the prefactor of 6.28 accounts for redundancy ($\times$ 2), cooling and other overheads ($\times$ 1.57), and the fact that today's network typically operate at under 50% utilization ($\times$ 2); the addends represent, in order, the ethernet switch, the broadband gateway, the data center gateway, the provider edge router, the core network, and the relay optical fiber transmission. In particular, the pre-factors are: 3 for the ethernet switch, accounting for the two routers involved in the access to the public internet plus the router located inside the data center; 2 for the provider edge router, considering the edge router in the edge network and the gateway router in the data center; 18 for the core network accounts for an average of 9 hops (2 baseline + 7 for the 800km distance between core nodes) of internet packets from source to destination, times 2 for the redundancy; 4 for the relay optical fiber transmission, since traffic must traverse an estimate of four core hops. A more detailed analysis of pre-factors can be found in [9].

3. Analysis of Distributed Cloud consumptions

To assess the GHG emissions of a distributed data center we must consider the operations phase, as well as the production, the shipping, the replacement and the end-of-life of the hardware. As for centralized data centers, in this chapter we focus on the operational consumption. The distributed architecture of the Cubbit network relies on the same public internet infrastructure delineated in the previous chapter. However, the distributed paradigm has three key differences from server-based cloud storage:

1. The low energy consumption of storage devices (based on the ARM Architecture [18]);
2. The absence of cooling overhead;
3. The geo-redundant architecture, meaning a geographical proximity between the users and their stored data without the need of any additional geo-redundancy policy.

We consider a network of Cubbit Cells [19], each composed by an ARM-based SBC and a HDD Seagate of 4 TB. Each Cell is located in a user's house and connected to internet by an internet service provider. Files on Cubbit's cloud are stored with a redundancy factor of 2 (Reed Solomon erasure coding with 12 + 12 redundancy shards [2, 20]). In addition, we estimate Cubbit to have a PUE of 1.0, since the Cells do not need additional cooling. Finally, we need to take into consideration the consumption of the Coordinator, a suite of machine learning algorithms that optimize the payload distribution on the network, while also taking care of security and metadata and in charge of triggering the recovery procedure for files on the Swarm.

3.1. Storage consumption

In the same way we did for the centralized cloud, we analyze the consumption of both storage (W/GB) and transfer (J/GB), as well as the impact of the Coordinator.

We assume that the network is composed by 4 TB devices. The ARM Architecture has a single-core peak consumption of $\sim$ 1W [18], while the embedded HDD Seagate has an operating consumption of 3.7 W. The storage energy consumption of the Cubbit network is therefore computed as:

$$\mathrm{P}_{\mathrm{cubbit}}^{\mathrm{storage}} = 2 \times \frac{4.7\ \mathrm{W}}{4\ \mathrm{TB}} \simeq 2.35\ \frac{\mathrm{W}}{\mathrm{TB}} \quad . \tag{3}$$

3.2. Transfer consumption

In Cubbit, each chunk of uploaded data is divided in a given number of shards that are distributed in an equal number of Cubbit Cells. Since the distribution of these shards is regulated by the AI Coordinator, and since the geographical region of each Cell can be estimated by its IP address, the Coordinator can choose to locate the uploaded data in the geographical proximity of the user, minimizing the need for data transfer on the public internet network [2]. To study the energy needed to transfer data on Cubbit, we used the geographical distribution of Cubbit Cells taken from the current (Jul 2022) batch of users of the public Cubbit network (5000 connected Cells). For each user, we computed the maximum physical distance that the data needs to travel from the user's own Cell to find 24 other connected Cells - enough to store a chunk of data with the 12+12 redundancy protocol. We called this measure Distance to Local Swarm (DLS). In Fig. 1, left panel, we show the distribution of DLS for Cells in the current public network, along with its median value, which is estimated to be 72.7 km. The median DLS depends on the redundancy protocol through the size of the local swarm (24, in the present case): the larger the number of Cells required to store a chunk of data, the longer the data will need to travel to find enough storage nodes. In Fig. 1, right panel, we show that the median DLS increases almost linearly on the local swarm size.

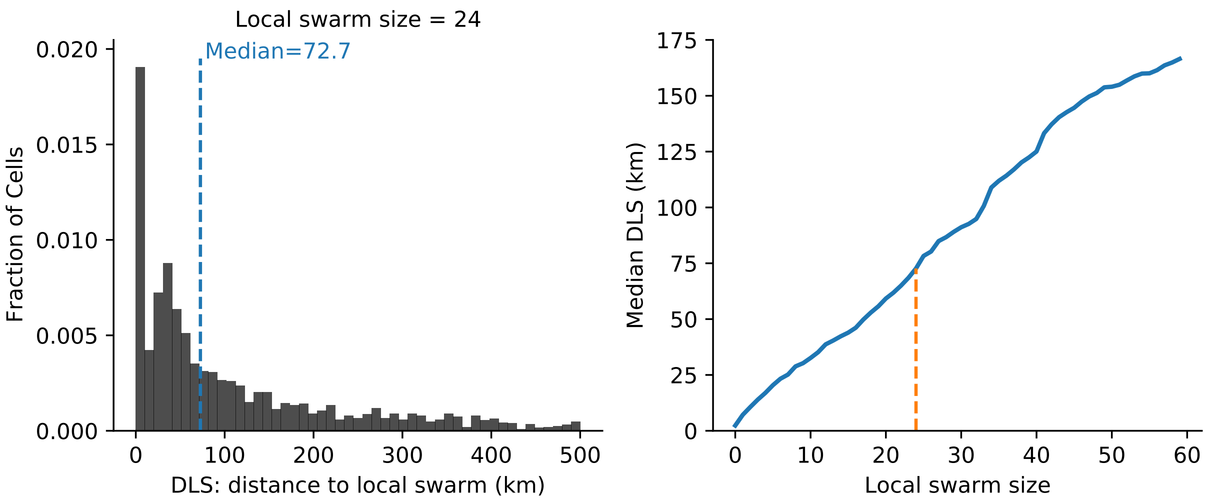

Figure 1: **Left**: distribution of Distance to Local Swam (DLS), defined as the maximum physical distance that the data needs to travel from the user's own Cell to find 24 other connected Cells, for Cells in the public Cubbit network computed as in Jul 2020. For graphical reasons, only Cells with DLS$<$ 500km are shown in the distribution, while all Cells are used to compute the median DLS. **Right**: median DLS (dashed blue line in the left panel) computed as in the left panel, as a function of the local swarm size. The dashed orange line represents the value of median DLS for local swarm size $= 24$, as in the $12 + 12$ redundancy protocol used in the public Cubbit network.

For the present computation, we can assume that an average of 2 packet hops in the core network routers are required to store the data locally in the Cubbit network. This lowers the corresponding factor 18 in Eq. 4 to a factor 4, accounting for two core hops and the redundancy of the packets on the network (factor 2). For the same reason, the 800km-relay consumption P_w is not taken into account. With respect to Eq. 4 we also ignore all data-center specific terms: one Ethernet switch and the data center gateway. However, we need to consider an additional BNG, since transfers are performed through p2p connections between endpoints located within an ISP network. The transfer energy per GB is therefore computed as

$$\begin{aligned} \mathrm{E}_{\text{cubbit}}^{\text{transfer}} &= 6 \times \left(2\frac{P_{es}}{C_{es}} + 2\frac{P_{bg}}{C_{bg}} + 2\frac{P_{pe}}{C_{pe}} + 4\frac{P_c}{C_c} \right) \\ &\simeq 11.9\ \frac{\text{kJ}}{\text{GB}} \quad . \end{aligned} \tag{4}$$

3.3. Coordinator consumption

To conclude our estimate of the total energy consumption of Cubbit, we need to take into account the consumption of the Coordinator. As the Coordinator is hosted on the public cloud, we used the production data gathered throughout the year 2021 to compute its energy consumption using the centralized cloud consumption estimated above [21]. The impact on storage was computed as the energy required for a centralized cloud infrastructure to store all the metadata produced by users in the public Cubbit cloud, divided by the total amount of stored TB in the public Cubbit cloud. The resulting energy consumption per TB was found to be $< 0.01\ W/TB$, therefore negligible compared to the storage consumption ($2.35\ W/TB$).

Similarly, the impact on transfer energy (kJ/GB) was estimated as the energy required by a centralized cloud to transfer the total amount of meta-data that the Coordinator exchanged with the users and Cells divided by the total mass of data uploaded on and downloaded from the public Cubbit network. Again, the final consumption was found to be negligible compared to the energy consumption computed above ($< 0.03\ kJ/GB$ compared to $11.9\ kJ/GB$).

Therefore, the impact of the Coordinator was not taken into consideration for the final calculations.

4. Comparison between centralized cloud and Cubbit distributed cloud

As a first analysis, we show in Fig. 2 the comparison between Cubbit and the centralized cloud in terms of pure storage consumption. Cubbit achieves a reduction of about 75% with respect to data centers' racks. By taking the average over racks as a mid-range estimation (although we have no information on the relative distribution of these racks in the data center market), we obtain

$$\Delta P^{\text{storage}} = \left\langle P_{\text{dcenter}}^{\text{storage}} \right\rangle - P_{\text{cubbit}}^{\text{storage}} \simeq 7.2\ \frac{\text{W}}{\text{TB}} \quad , \tag{5}$$

corresponding to an overall 75% reduction:

$$\frac{\Delta P^{\text{storage}}}{\left\langle P_{\text{dcenter}}^{\text{storage}} \right\rangle} = \simeq 0.75 \quad . \tag{6}$$

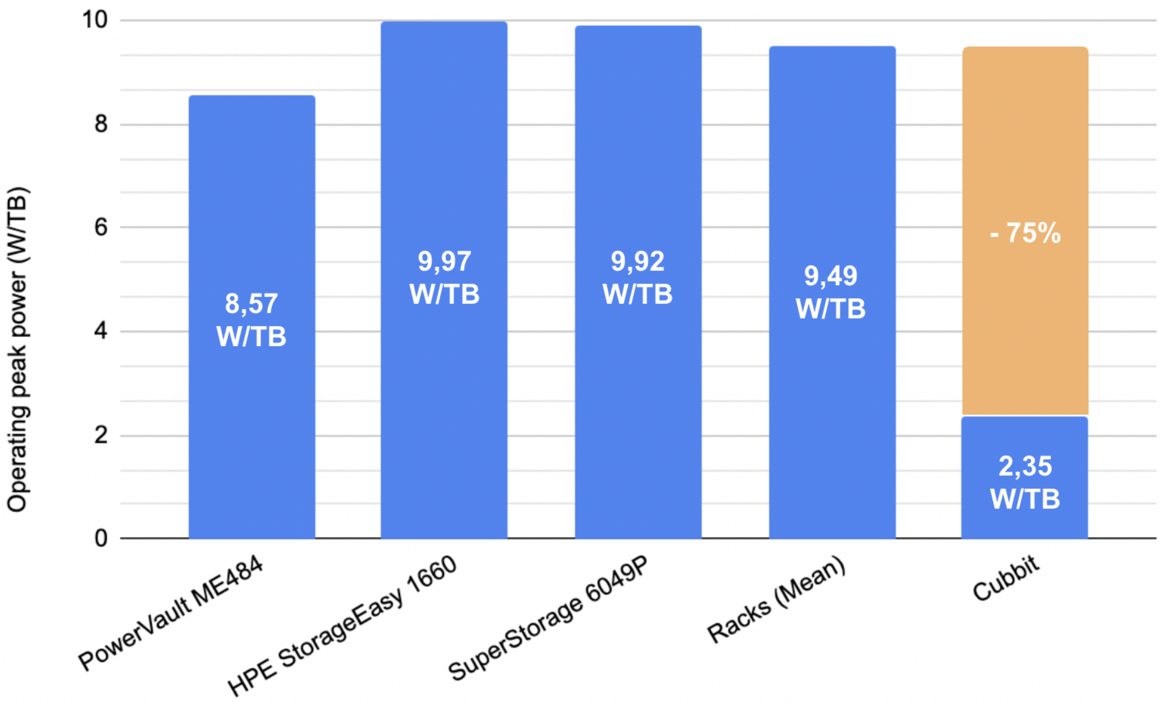

Figure 2: Comparison between centralized and distributed clouds in terms of peak consumption per TB of stored data.

Similarly, the difference in terms of transfer energy per GB is

$$\Delta E^{\mathrm{transfer}} = E^{\mathrm{transfer}}_{\mathrm{dcenter}} - E^{\mathrm{transfer}}_{\mathrm{cubbit}} \simeq 18.7\,\frac{\mathrm{kJ}}{\mathrm{GB}} = 5.2\,\frac{\mathrm{kWh}}{\mathrm{TB}} \quad , \tag{7}$$

which corresponds to a 61% reduction of the energy needed to transfer data from the cloud to the user, and back.

The reduction of carbon footprint of Cubbit compared to centralized solutions can be computed by comparing the storage power and the transfer energy for typical use case, such as backup plans and frequent access of, for instance, a web-hosted video.

5. Use cases

5.1. Backup

A backup service hosted on the cloud is characterized by large volumes that are not frequently accessed. In the context of the carbon footprint, the consumption of a backup plan will, therefore, be dominated by the storage term. If we consider a storage plan for a professional backup of 25 TB with very small daily access, we find that the total energy saved in a year is

Cloud — Storage	25 TB	250 TB	1 PB	100 PB
Cubbit	206	2.057	8.234	823.400
Average Centralized Cloud	831	8.313	33.253	3.325.300
Saved emissions	-625	-6.255	-25.019	-2.501.900

Table 4: **Backup comparison** - Kg CO2 emissions saved per year by storing data on a decentralized cloud

$$\Delta E(25 \text{ TB backup}) = \\ 25 \text{ TB} \times \Delta P^{\text{storage}} \times 365 \times 24h \simeq 1565 \text{ kWh} . \tag{8}$$

By considering a rough factor of 0.4 KgCO2 for each kWh of consumed energy[22], the use of a distributed cloud over a centralized one would correspond, for such a backup plan, to a reduced carbon emission of c.a. -625 kgCO2/year. If we consider that data centers usually operate on the

Petabyte scale, we easily see that the yearly reduction in emissions scales up to hundreds of tons of CO2

$$\Delta \text{CO2 emissions(backup)} \simeq -25\ 000 \text{ kgCO2/year/PB} \tag{9}$$

Tab.4 illustrates the savings in terms of CO2 emissions per year by storing data on a decentralized cloud in different scenarios.

5.2. Streaming on a local scale

The reduction in consumed energy and, consequently, in carbon emission is significantly larger when considering large volumes of data transfers. For example, if we consider a medium news website hosting 25 TB of data and streaming, on average, 10 TB of data per day (e.g. 10,000 visualizations of 100 MB each) the saved energy per year would be

$$\begin{aligned} \Delta E(10 \text{ TB streaming}) = \\ & 25 \text{ TB} \times \Delta P^{\text{storage}} \times 365 \times 24h \\ & + 365 \times \Delta E^{\text{transfer}} \times 10 \text{ TB} \\ & = 20676 \text{ kWh} , \end{aligned} \tag{10}$$

which roughly corresponds to -8,360 kgCO2 emitted per year. Note that these computations assume that data are broadcasted to a local audience. While this might be the case for university data, local news, or targeted marketing, it has a limited range of applicability that has to be taken into account.

Tab.5 illustrates the savings in terms of emissions by streaming data on a decentralized cloud in different scenarios.

Cloud — Storage; transfer	25 TB; 10 TB	250 TB; 100 TB	1 PB; 1 PB	100 PB; 100 PB
Cubbit	5.024	50.240	490.034	4.900.340
Average Centralized Cloud	13.387	133.873	1.288.853	12.888.530
Saved emissions	-8.363	-83.635	-798.819	-7.988.190

Table 5: **Streaming comparison** - Kg CO2 emissions saved per year by streaming data on a decentralized cloud

5.3. Large scale: the global cloud industry

Finally, if we speculate about the overall data volume of a global consumer cloud storage service like, for example, Dropbox or Google, values rise dramatically. Such interpolations have to be taken with due caution, but researches and press releases show how, in 2020, Dropbox and Google Suite together had over 21 millions paying customers, with minimum 2 TB premium cloud storage each [23, 24]. Considering the presence of at least other 5 big players with similar numbers in the market, we can estimate a x2.5 cloud storage market size [25]. This results in a theoretical data volume of ca. $\sim 1 \cdot 10^8$ TB of storage. Considering a factor 5 due to overbooking, it gives an estimation of $\sim 2 \cdot 10^7$ TB of effective cloud storage. We can make a conservative estimation that each user transfers, on average, 50 MB of files from/to the cloud, which implies a daily transfer volume of c.a. 190 TB.

If we plug these estimations in our model, we obtain a total saved annual energy, using a distributed architecture rather than a centralized one, of $\sim 1.3 \cdot 10^9$ kWh, equivalent to saving carbon emissions in the order of 500 million kgCO2 per year.

6. Conclusions

To better compare centralized and decentralized clouds in terms of CO2 emissions, it is important that further research takes into consideration not only the operation phase but also the set up, maintenance and dismissal of the infrastructure.

In particular, this research does not take into account the externalities generated by the construction of centralized data centers, that pose a big threat to local communities and biodiversity. If it is clear the trend for centralized data center to use renewable energy, it is important to consider the externalities generated by this industry. Also, it is important to consider the externalities generated by the cooling system, that use Flourinated gases or liquids that need to be disposed of.

In addition, this study does not take into consideration the production, maintenance, substitution and end of life of the hardware used in centralized data centers, which we know is on average being replaced every 3-4 years [10]. Finally, also the geo-redundancy should be further investigated, since geography has a relevant influence on energy consumption.

For what concerns the emissions related to Cubbit, different aspects need to be taken into consideration to estimate the entire CO2 value generated by its distributed cloud. In particular, production, shipping, replacement, end-of-life of the Cells and the recovery procedure of the Coordinator have not been included in the research.

One element to note is that the efficiency of manufacturing small board computers as in the case of Cubbit, as well as the repairability and hardware longevity associated with such devices, should be considered when comparing it with rack-mounted server boxes.

On top of committing in the near future to measuring and updating this study with further information, Cubbit also intends to expand its infrastructure with already existing hardware and virtualized Cells. For example, if a company or a private owns a that is left, entirely or partially, unused, this can be plugged into the infrastructure of Cubbit, adding new space to its cloud. By doing this, Cubbit intends to 'recycle the internet': the production of new hardware Cells will no longer be necessary, thus reducing its impact in terms of production and shipping.

To conclude, a lot of information in the cloud storage sector is not publicly available, though it is necessary to improve the energy consumption standards of the market. More transparency in the market would allow positive competition in the race for reducing the CO2 impact of the cloud.

Acknowledgements

The authors are grateful to S. Baldi, A. Albani, and A. Rovai for valuable comments on the manuscript and fruitful discussion.

References

[1] . E. A. Belkhir, L., Assessing ict global emissions footprint: Trends to 2040 recommendations (2018).
[2] L. Posani, M. Moschettini, A. Paccoia, Cubbit: a distributed, crowd-sourced cloud storage service, Invited Talk at CERN Workshop on Cloud Storage Services (CS3 2018) (2018).
[3] L. Posani, M. Moschettini, A. Paccoia, Cubbit, the distributed cloud, Invited Talk at CERN, CNR Workshop on Cloud Storage Services (CS3 2019) (2019).
[4] UNFCCC, The paris agreement. retrieved 2022, june, https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement, 2022, February.
[5] E. Commission, A european green deal — striving to be the first climate-neutral continent. retrieved 2022, june, https://ec.europa.eu/info/strategy/priorities-2019-2024/european-green-deal_en, 2022.
[6] IPCC, Sixth Assessment Report - Impacts, Adaptation and Vulnerability. Retrieved 2022, June, Technical Report, 2022.
[7] I. E. Agency, Data centres and data transmission networks – analysis. retrieved 2022, june, https://www.iea.org/reports/data-centres-and-data-transmission-networks, 2021, November.
[8] B. Golub, The green case for storj, https://www.storj.io/blog/the-green-case-for-storj, 2022.
[9] J. Baliga, R. W. Ayre, K. Hinton, R. S. Tucker, Green cloud computing: Balancing energy in processing, storage, and transport, Proceedings of the IEEE 99 (2011) 149–167.
[10] J. Booth, Data center sustainability. data center forum, helsinki 2021, 2021.
[11] Dell, Powervault me4 series san/das storage specification sheet. retrieved 2022, june, https://i.dell.com/sites/csdocuments/Product_Docs/en/powervault-me484-jbod-spec-sheet.pdf, 2022.
[12] HPE, Hpe storeeasy 1000 storage with the hpe storeeasy management console, https://www.hpe.com/psnow/doc/a00041171enw?section=Document%20Types&r=none, 2022.
[13] I. Super Micro Computer, 6049p-e1cr24h — superstorage — products — super micro computer, inc.. retrieved 2022, june, https://www.supermicro.com/en/products/system/4U/6049/SSG-6049P-E1CR24H.cfm, ????
[14] Backblaze, 2021.
[15] U. Institute, Uptime institute global data center survey 2021 - growth stretches an evolving sector. retrieved 2022, june, 2021.
[16] Azure, Data redundancy - azure storage, "https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy", 2022.
[17] Dropbox, Dropbox business security, a dropbox whitepaper. retrieved 2022, june, https://aem.dropbox.com/cms/content/dam/dropbox/www/en-us/business/solutions/solutions/dfb_security_whitepaper.pdf, 2022.
[18] Marvell espressobin specs sheet, https://www.http://espressobin.net/, web page.
[19] Cubbit website, https://www.cubbit.io/technology, web page.
[20] M. Monti, S. Rasmussen, M. Moschettini, L. Posani, An alternative information plan, Technical Report, Working paper Santa Fe Institute, 2017.
[21] Cubbit, Cloud consumption - usage of the coordinator 2021, 2021.
[22] U. E. I. Administration, How much carbon dioxide is produced per kilowatthour of u.s. electricity generation? retrieved 2022, june, https://www.eia.gov/tools/faqs/faq.php?id=74&t=11, 2021, November.
[23] Statista, Dropbox number of paying users 2015-2020. Retrieved 2022, June, Technical Report, 2022.
[24] CNBC, Google g suite passes 6 million customers., "https://www.cnbc.com/2020/04/07/google-g-suite-passes-6-million-customers.html, 2020.
[25] F. B. Insight, Cloud storage market size, share, growth forecast [2029] (2022, April).