Emission Impossible: privacy-preserving carbon emissions claims

Jessica Man, Sadiq Jaffer, Patrick Ferris, Martin Kleppmann and Anil Madhavapeddy
Department of Computer Science & Technology, University of Cambridge

1 INTRODUCTION

Customers of online services may want to take carbon emissions into account when deciding which service to use, but are currently hindered by a lack of reliable emissions data that is comparable across services. Calculating accurate carbon emissions across a cloud computing pipeline involves a number of stakeholders, none of whom are incentivised to accurately report their emissions for competitive reasons. In this paper we explore mechanisms to support verifiable and privacy-preserving emissions reporting across a chain of energy suppliers, cloud data centres, virtual machine hosting services providers and cloud services providers, which are ultimately passed through to APIs used by customers. We believe that adding verifiable and composable emissions transparency to cloud computing architectures enables providers to compete on the basis of sustainability, resulting in demand-side pressure on cloud services to shift to renewable energy sources [6].

Our technique centres around zero-knowledge proofs (ZKPs) [10]. When applying ZKPs to our problem, a stakeholder in a supply chain proves to a verifier (who can be anyone) that the carbon emissions calculations were made accurately, without revealing commercially sensitive data about their business operations. The verifier decides if the claim can be accepted by using only public knowledge and the proof provided by the stakeholder, achieving verification with strong privacy guarantees. We propose the ZKP system to be based on a zero-knowledge Succinct Non-interactive ARguments of Knowledge (zk-SNARK) protocol, which provides "good enough" efficiency for computation of emissions claims and rich enough functionality for zero-knowledge proofs [3, 4].

Figure 1 (overleaf) shows the challenge of tracking carbon emissions from a cloud computing supply chain, and how ZKP can be applied to provide verifiable emissions data without revealing trade secrets.

In this talk we will explore: *(i)* how the conflicting incentives around carbon emissions reporting make existing systems unlikely to succeed; and *(ii)* how we could use ZKPs to allow for the accurate reporting of carbon claims *without* compromising on individual privacy and competitiveness requirements.

2 INCENTIVES TO CHEAT FOR CO2E CLAIMS

It was estimated that 1.8% to 3.9% of the global carbon emissions are attributable to Information and Communication Technology (ICT) [1]. Governments, investors and consumers are therefore watching how cloud computing operators are working towards net zero [2, 9, 16]. However, businesses have strong incentives to make only positive claims, which could involve hiding data or publishing misleading results with dubious evidence, a problem termed "greenwashing" [13]. Moreover, they are keen not to expose any private business data to preserve their competitive advantage. So we end up with the current state that many claims are disclosed to the public but it is difficult to know if any of the claims are true.

Consider three of the biggest data centre providers: Amazon, Google and Microsoft. They have all reported their emissions goals publicly, but have also used creative accounting to hide facts about their carbon emissions [18]. Both Microsoft and Google admitted that their carbon emissions had increased in recent years, despite their climate commitments [12, 15], and Amazon's self-reporting did not include the emissions data of the products sold by third-party vendors [7]. The major risks involved in data centre emissions reporting are:

(1) **Privacy and Trade secrets concerns.** Reports on carbon emissions typically show only high-level aggregated data. Validation of the claims often requires details of carbon activities, which could potentially expose trade secrets through disclosure.
(2) **Untrustworthy claims.** Companies make misleading claims based on dubious accounting methodologies to make it look like they are more environmentally friendly than they actually are [13]. Yang et al. studied greenwashing behaviours and impact and found that greenwashing often links with scandals that happen at the supply chain level [19].
(3) **Missing claims.** Companies can choose not to disclose anything, or report claims that omit some of their emissions-generating activities. Amazon's undercounting on their carbon footprint reports is a good example [7].

We can mitigate the first risk to protect businesses by ensuring that the verification method will not leak secrets. For the second risk we can use verified data to provide trustworthy claims. ZKPs can be used for these mitigations and provide a mechanism to standardise accounting methodologies. For the third risk, we cannot validate data that is missing. We can, however, bind carbon accounting to financial accounting to make it more difficult to cheat. For example, a ZKP can show that the total paid to all suppliers matches the number that is reported on the audited accounts, and that these supplier transactions are also reflected in the emissions calculation. Finally, we can analyse and benchmark the verified data across companies and look for discrepancies and anomalies through manual audits.

3 A ZKP EMISSIONS DISCLOSURE SCHEME

A zk-SNARK allows a prover to convince a verifier that the prover knows values (a *witness*) that satisfy a given set of equations (a *circuit*) without revealing any information about the witness [4, 5, 14]. Similarly to a digital signature, SNARKs do not require any interaction between the prover and verifier besides sending the proof. Proofs are short in length and can be verified quickly – in some cases, in constant time. SNARKs are actively used in real world applications such as the Zcash cryptocurrency, where transactions are executed without any sensitive information such as the origin or amount of the transaction being revealed [11].

1st International Workshop on Low Carbon Computing (LOCO 2024), Dec 03, 2024, Glasgow, Scotland, UK
2024. https://doi.org/10.1145/nnnnnnn.nnnnnnn

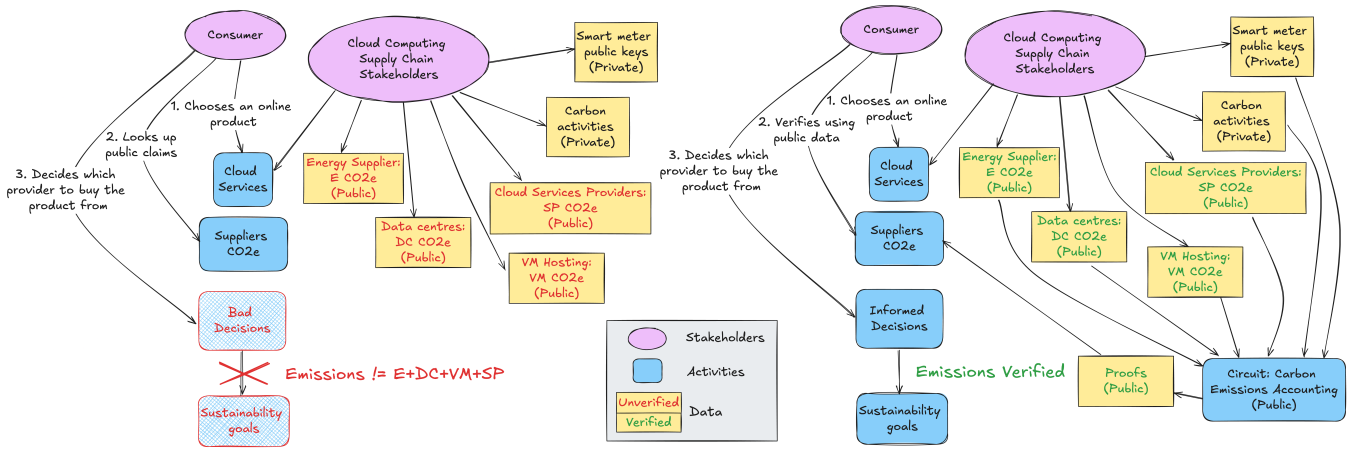

Figure 1: Unverified emissions claims can affect sustainability goals negatively *(left)* but ZKP claims can be verified directly by consumers to make informed choices *(right)* via transparent emissions declarations across the cloud stakeholder chain.

To illustrate, consider a scenario where a user wants to compare the carbon emissions of AI chatbots such as ChatGPT (OpenAI), Gemini (Google) and Claude (Anthropic). To ensure a fair comparison, they must be computed according to the same methodology, and each emissions claim must be verified end-to-end. One aspect of such a claim would be to verify that the disclosed electricity consumption figures come from a trusted source.

First we assume that the data centre operator uses smart meters to track power consumption, and that the smart meters have a secure hardware element that signs meter readings with a private key configured by the meter manufacturer. Then we apply a zk-SNARK in three stages:

Stage 1: Generate circuit C. The circuit encodes the accounting methodology used to calculate the emissions. The constraints and relationships among the input parameters are reduced to a set of polynomial equations. Cryptography is applied such that it is not feasible for anyone to calculate the input data by knowing the output proof generated by the circuit, hence protecting the secrets. There are many ways to calculate emissions and the accounting is non-trivial because of many variables. For example, to calculate the power consumption of a ChatGPT session, we cannot simply divide the total consumption by the number of users, as different users may have very different usage levels. We also need to account for fixed costs (such as model training) and waste (such as models that were trained but never used because they performed poorly). There are emerging carbon accounting methods (e.g. PACT [8]) and methodologies for attributing load to individual customers within a data centre [17], which could inform circuit construction.

Stage 2: Prove. The prover proves that the emissions were calculated using only trusted and accurate data, such as emissions figures, smart meter readings, carbon intensity, GPU power consumption per tenant, AI token metrics and many more. Some of these parameters are only known to the prover. To prove that the smart meter readings came from trusted smart meters, the public keys used by the smart meters can be revealed as public data used to generate the proofs.

Stage 3: Verify. The verifier uses the proof and public data to determine if they can believe that the prover has the knowledge of the private data, such that when both the private and public data are used along with the emissions figures, they satisfy all the equations encoded in the circuit. That is,

$$\exists \textit{private witness}.\ C(\{\textit{public}, \textit{private}\}\ \textit{witness}) = \textit{True}$$

In our example, the verifier would only believe a smart meter reading came from a trusted smart meter if the reading has been signed by a private key and the signature can be verified using the public key available to them, without knowing the private key.

4 UNLOCKING EMISSIONS IMPOSSIBLE

It seems almost impossible to balance privacy and competitiveness needs with our urgent sustainability goals to reduce emissions, particularly in ICT sector. The use case outlined in this paper is an example of how we can achieve privacy-preserving and trustworthy carbon emissions claims for data centres. The approach can be used in other industries as well. To adopt the ZKP system companies can apply carbon accounting alongside their financial accounting, which already needs to attribute use of computing resources to individual customers. A benefit of using a SNARK is that the proofs are small enough to be bundled along with emissions reporting.

We believe that this proposal is a step forward for carbon emissions accounting to be public and explicit, making emissions tracking more accurate and comparable across companies. We are exploring how this could be exposed directly to end users via browser plugins, providing an end-to-end verifiable CO2e cost alongside conventional costs used by users to make their buying decisions (such as price, delivery time, or distance to the service). Our overall aim is to drive demand-side pressure to reduce unnecessary emissions from data centre use by informing consumers about the environmental cost of their actions online.

REFERENCES

[1] ACM. 2021. *ACM Technology Policy Council Releases TechBrief on Computing and Carbon Emissions.* https://www-acm-org.ezp.lib.cam.ac.uk/media-center/2021/october/tpc-tech-brief-climate-change

[2] José Azar, Miguel Duro, Igor Kadach, and Gaizka Ormazabal. 2021. The big three and corporate carbon emissions around the world. *Journal of Financial Economics* 142, 2 (2021), 674–696.

[3] Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. 2014. Zerocash: Decentralized Anonymous Payments from Bitcoin. In *2014 IEEE Symposium on Security and Privacy*. 459–474.

[4] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. 2014. Succinct Non-Interactive Zero Knowledge for a von Neumann Architecture. In *23rd USENIX Security Symposium (USENIX Security 14)*. 781–796.

[5] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. 2012. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *Proceedings of the 3rd innovations in theoretical computer science conference*. 326–349.

[6] Esther Calderon-Monge, Ivan Pastor-Sanz, and F Javier Sendra Garcia. 2020. Analysis of sustainable consumer behavior as a business opportunity. *Journal of Business Research* 120 (2020), 74–81.

[7] Will Evans. 2022. *Private Report Shows How Amazon Drastically Undercounts Its Carbon Footprint.* https://revealnews.org/article/private-report-shows-how-amazon-drastically-undercounts-its-carbon-footprint/

[8] Partnership for Carbon Transparency. 2023. *Pathfinder Framework: Guidance for the Accounting and Exchange of Product Life Cycle Emissions.* https://www.wbcsd.org/wp-content/uploads/2023/01/PACT-Pathfinder-Framework-WBCSD.pdf

[9] Erol Gelenbe and Yves Caseau. 2015. The impact of information technology on energy consumption and carbon emissions. *ubiquity* 2015, June (2015), 1–15.

[10] Shafi Goldwasser, Silvio Micali, and Chales Rackoff. 2019. The knowledge complexity of interactive proof-systems. In *Providing sound foundations for cryptography: On the work of shafi goldwasser and silvio micali*. 203–225.

[11] Daira Hopwood, Sean Bowe, Taylor Hornby, Nathan Wilcox, et al. 2016. Zcash protocol specification. *GitHub: San Francisco, CA, USA* 4, 220 (2016), 32.

[12] Tobias Mann. 2024. *So much for green Google ... Emissions up 48% since 2019.* https://www.theregister.com/2024/07/02/google_datacenter_emissions/

[13] United Nations. 2024. *Greenwashing – the deceptive tactics behind environmental claims.* https://www.un.org/en/climatechange/science/climate-issues/greenwashing

[14] Assa Naveh and Eran Tromer. 2016. PhotoProof: Cryptographic Image Authentication for Any Set of Permissible Transformations. In *2016 IEEE Symposium on Security and Privacy (SP)*. 255–271. https://doi.org/10.1109/SP.2016.23

[15] Dan Robinson. 2024. *Microsoft's carbon emissions up nearly 30% thanks to AI.* https://www.theregister.com/2024/05/16/microsoft_co2_emissions/

[16] Goldman Sachs. 2024. *AI is poised to drive 160% increase in data center power demand.* https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand

[17] Ian Schneider and Taylor Mattia. 2024. Carbon accounting in the Cloud: a methodology for allocating emissions across data center users. *arXiv preprint arXiv:2406.09645* (2024).

[18] James Temple. 2024. *Google, Amazon and the problem with Big Tech's climate claims.* https://www.technologyreview.com/2024/07/17/1095019/google-amazon-and-the-problem-with-big-techs-climate-claims/

[19] Zhi Yang, Thi Thu Huong Nguyen, Hoang Nam Nguyen, Thi Thuy Nga Nguyen, and Thi Thanh Cao. 2020. Greenwashing behaviours: Causes, taxonomy and consequences based on a systematic literature review. *Journal of business economics and management* 21, 5 (2020), 1486–1507.