

# Emission Impossible: privacy-preserving carbon emissions claims

Jessica Man

Department of Computer Science &  
Technology, Cambridge, UK  
psjm3@cam.ac.uk

Sadiq Jaffer

Department of Computer Science &  
Technology, Cambridge, UK  
sj514@cam.ac.uk

Patrick Ferris

Department of Computer Science &  
Technology, Cambridge, UK  
pf341@cam.ac.uk

Martin Kleppmann

Department of Computer Science &  
Technology, Cambridge, UK  
martin.kleppmann@cst.cam.ac.uk

Anil Madhavapeddy

Department of Computer Science &  
Technology, Cambridge, UK  
avsm2@cl.cam.ac.uk

## ABSTRACT

Information and Communication Technologies (ICT) have a significant climate impact, and data centres account for a large proportion of the carbon emissions from ICT. To achieve sustainability goals, it is important that all parties involved in ICT supply chains can track and share accurate carbon emissions data with their customers, investors, and the authorities. However, businesses have strong incentives to make their numbers look good, whilst less so to publish their accounting methods along with all the input data, due to the risk of revealing sensitive information. It would be uneconomical to use a trusted third party to verify the data for every report for each party in the chain. As a result, carbon emissions reporting in supply chains currently relies on unverified data. This paper proposes a methodology that applies cryptography and zero-knowledge proofs for carbon emissions claims that can be subsequently verified without the knowledge of the private input data. The proposed system is based on a zero-knowledge Succinct Non-interactive ARguments of Knowledge (zk-SNARK) protocol, which enables verifiable emissions reporting mechanisms across a chain of energy suppliers, cloud data centres, cloud services providers, and customers, without any company needing to disclose commercially sensitive information. This allows customers of cloud services to accurately account for the emissions generated by their activities, improving data quality for their own regulatory reporting. Cloud services providers would also be held accountable for producing accurate carbon emissions data.

## KEYWORDS

Carbon Emissions, Zero-Knowledge Proofs, zk-SNARK, Cloud Computing

## 1 INTRODUCTION

Regulatory requirements and sustainability initiatives mean that companies are increasingly having to report their carbon emissions. Looking at ICT companies in particular, customers of online services, who are also obligated to report their emissions data or who might want to take carbon emissions into account when deciding which service to use, are currently hindered by a lack of reliable emissions data that are 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 confidentiality-preserving emissions reporting across a chain of energy suppliers, cloud data centres, virtual machine hosting service providers and cloud services providers, which are ultimately passed through to their 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 [10].

Our technique centres around zero-knowledge proofs (ZKPs) [17]. When applying ZKPs to the issue of untrusted carbon emissions claims, a stakeholder in a supply chain proves to a verifier (who can be anyone, such as a customer, investor, or regulator) that the emissions calculations were performed accurately, without revealing commercially sensitive data about their business operations. The verifier decides whether the claim can be accepted using only public knowledge and the cryptographic proof provided by the stakeholder. The proposed system applies zero-knowledge Succinct Non-interactive ARguments of Knowledge (zk-SNARK) as the protocol to allow proofs to be generated and subsequently verified without the requirement of disclosing all of the input data.

In this paper, we present our argument on why conflicting incentives around carbon emissions reporting make existing systems unlikely to succeed (§2.2) and our contribution of applying ZKPs to allow more accurate reporting of carbon claims without compromising sensitive information (§4).

## 2 BACKGROUND

### 2.1 Carbon Emissions Reporting

In some countries, large companies must disclose their carbon emissions to comply with regulations. For example, the UK's Streamlined Energy and Carbon Reporting (SECR) regulations require all UK quoted companies and large limited liability partnerships to report on their global energy use in addition to greenhouse gas emissions [18]. The EU Corporate Sustainability Reporting Directive (CSRD) came into force on 5 January 2023, which requires large companies and listed SMEs to report on sustainability, highlighting the urgent need for the disclosure of 'relevant, comparable and reliable sustainability information' and the significant increase in demand for sustainability information [41]. Other than for regulatory reasons, companies publicising their plans towards net-zero could have a

positive impact on their businesses. On the one hand, transparency on climate actions taken as part of the manufacturing and distribution process behind commercial products can directly influence consumer behaviours [10]. On the other hand, if companies have to disclose their confidential business data as part of the reporting, this could benefit competitors.

Pressure from the media, investors, and customers could also have an effect on businesses. Amazon was given an ‘F’ grade (meaning no response) by CDP<sup>1</sup> until 2023 when they first submitted their report [14]. Shareholders of the company continue to request Amazon to provide additional information on climate-related impacts [3, 39].

## 2.2 Conflicting Incentives Behind CO<sub>2e</sub> Claims

An estimated 1.8% to 3.9% of global carbon emissions are attributable to Information and Communication Technology (ICT) [2]. Governments, investors and customers are therefore paying close attention to how cloud computing operators work towards net zero [5, 16, 35]. However, companies closely guard the metrics behind the computational resources used to provide particular services, as disclosing them would undermine their pricing and business model. 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’ [30]. Therefore, it is difficult for an outsider to know whether a company’s emissions claims are true.

There are now tightened regulations to tackle greenwashing in some countries, for example, the EU has proposed the ‘Green Claims Directive’ [12] to prevent companies making claims without providing clear evidence. The verification, however, relies on accredited verifiers. The UK’s CMA has also developed the ‘Green Claims Code’ to combat greenwashing<sup>2</sup>, which focuses on a set of core principles based on existing consumer law, to protect businesses and consumers from misleading environmental claims. However, measurements of carbon emissions are continuous and frequent, and with input data varying at each measurement, it is not practical to have an independent auditor to verify every single claim for a supply chain with multiple companies involved.

Consider three of the biggest data centre providers: Amazon, Google and Microsoft. They have all reported their decarbonisation goals publicly, but have also been accused of using creative accounting to hide facts about their carbon emissions [40]. Both Microsoft and Google admitted that their carbon emissions had increased in recent years, despite their climate commitments [27, 34], and Amazon’s self-reporting did not include emissions data for products sold by third-party vendors [14]. The major risks involved in data centre emissions reporting are:

- (1) **Privacy and trade secrets concerns.** Reports on carbon emissions typically show only aggregated data at high levels. Validation of claims often requires details of carbon-emitting activities, which cannot be publicly disclosed because they are trade secrets of both suppliers and customers.

- (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 [30]. Yang et al. studied greenwashing behaviours and impact and found that greenwashing is often linked with scandals that occur at the supply chain level [42].
- (3) **Missing claims.** Companies can choose not to disclose anything or report claims that omit some of their emissions-generating activities. Amazon’s undercounting in their carbon footprint reports is a good example [14].

We can mitigate the first risk to protect businesses by ensuring that the verification method does not leak secrets. For the second risk, we can use verified data to provide trustworthy claims. Companies can provide proofs that their claims on the carbon emissions report are all true, and the proofs can be checked out by their customers, investors or auditors. ZKPs (§3.2) can be used for these mitigations. To ensure figures are comparable across different companies, the calculation methodology can be standardised, as discussed in §4. For the third risk, we cannot validate missing data. We can, however, bind carbon accounting to financial accounting to make it more difficult to cheat, as discussed in §5. Finally, we can analyse and benchmark the verified data across companies and look for discrepancies and anomalies through manual audits. Manual audits are carried out infrequently, typically annually, and it is a lengthy process. Therefore, while we cannot totally eliminate the need for manual audits, ZKPs provide a complementary process that can be automated for much faster and more frequent verifications.

## 3 RELATED WORK

### 3.1 Carbon Accounting and Reporting

The most commonly used approach to calculate carbon emissions is to follow the Greenhouse Gas (GHG) Protocol [32]. According to the GHG Protocol Corporate Standard, emissions are categorised in three scopes: Scope 1 emissions occur from sources that are owned or controlled by the company; Scope 2 emissions are generated from purchased electricity or heat consumed by the company; Scope 3 emissions are a consequence of the activities of the company, but occur from sources not owned or controlled by the company, primarily by its direct or indirect suppliers. In most companies, Scope 3 accounts for the majority of emissions by far [15], but is also the most complex to compute, given that the calculation of emissions requires data from the entire supply chain. For those companies that report their carbon emissions, most of them currently report on Scope 1 carbon emission, less so on Scope 2 and very little on Scope 3 [15].

When an organisation migrates their on-premises computing resources and IT workload to the cloud, emissions under scope 1 or scope 2 move to scope 3. Given that scope 3 emissions reporting is voluntary and data centres’ emissions are aggregated into the global reporting by large cloud providers, cloud service customers’ emissions become hidden [29]. Customers therefore rely on cloud providers to tell them their share of emissions arising from the data centre for the hosted services they use, but the lack of transparency has made this a big challenge for the customers. If the information is not provided by the cloud providers, customers have to estimate

<sup>1</sup>A non-profit carbon disclosure company, formally known as the Carbon Disclosure Project

<sup>2</sup><https://greenclaims.campaign.gov.uk>

using aggregated global data and other methods to obtain information, such as through Freedom of Information requests, or make in-house measurements to produce the metrics themselves [13, 37].

In recent years, the big cloud providers have started to provide their customers with more detailed reports about the emissions of the computing resources they use. For example, Microsoft offers an ‘Emissions Impact Dashboard’ tool to their customers [28], with the methodology verified by a third-party company [4]. Google also offers a similar tool to their cloud customers and has published their carbon accounting methods [36], and the methodology has also been reviewed by a third party company [1]. Neither of these tools provides a way for customers to independently verify the reported data. The review process by the cloud providers’ appointed trusted third-parties does not validate the reported data for all customers (only a sample during the review period).

WBCSD<sup>3</sup> (The World Business Council for Sustainable Development) was formed in 1995 and provides a platform for businesses around the globe to respond to sustainability challenges. WBCSD has over 200 leading business members globally across multiple industries, working together to create standards, policies, and best practices that drive the sustainability agenda. They published the ‘Partnership for Carbon Transparency’ (PACT) methodology to provide guidance on carbon accounting and exchange of emissions data, and highlighted that assurance and verification are key in ensuring the credibility and reliability of exchanged data [15]. Their methodology involves using a third-party provider for the verification; our proposal could complement that approach by providing an automatic approach to protecting confidential data with a high level of confidence.

Heiss et al. [20] propose using zero-knowledge proofs to provide verifiable data in carbon emissions accounting, while protecting confidential business data. Their method applies to product carbon accounting, and they used the automotive industry as a practical use case to demonstrate at a high level how the approach would work. Their design extends the ‘Digital monitoring, reporting and verification (D-MRV)’ systems, which rely on blockchains to produce authenticity and integrity proofs.

### 3.2 Zero-Knowledge Proofs (ZKPs)

We propose using a zero-knowledge Succinct Non-interactive Arguments of Knowledge (zk-SNARK) protocol [8]. zk-SNARK protocols allow a prover to convince a verifier that the prover knows values (a *witness*) that satisfy a given set of equations (a *circuit*) in *zero knowledge*, i.e. without revealing any information about the witness [8, 9, 31]. zk-SNARKs are *succinct*: informally, this means that the proof is small compared to the witness and fast to verify [38]. The protocol is *non-interactive*, which is similar to digital signatures in that the prover only needs to send a single message to the verifier (the proof). The *ARK* part of zk-SNARK means that the prover can convince the verifier that it knows the witness. This property is formalised as *knowledge soundness*, which means that a computationally bounded prover cannot generate a proof of a false statement, or a statement for which it does not have a witness, except with negligible probability [38].

zk-SNARKs are typically applied in systems for managing digital assets, where transactions are executed without sensitive information such as the origin or amount of the transaction being revealed [7, 21, 22]. Zcash [7, 21] was one of the first widespread applications of zk-SNARKs, allowing transactions to be fully encrypted on a blockchain and still be verifiable. zk-SNARK has also been applied to protect medical data. Luong and Park [26] proposed a blockchain-based system with IoT devices, allowing data sharing without leaking patient records.

Blockchains incur large infrastructure costs due to their requirement to establish consensus on a global ledger of transactions in a trustless setting. However, our emissions reporting use case does not require such a global ledger, since emissions data can be exchanged directly between suppliers and customers. We can therefore avoid the costs of blockchains in our system.

## 4 A ZKP EMISSIONS DISCLOSURE SCHEME

The core challenge in applying ZKPs to carbon emissions reporting is ensuring that the data in the ZKP is an accurate reflection of reality. ZKPs can only check that the prover knows some input values, but those values could be made up. We approach this problem by identifying ways of cryptographically proving the correctness of every input to the emissions calculation, for example by assuming that it is signed by a trusted authority (whose public key becomes a root of trust), or by requiring the originator of an input value to also provide a SNARK proof of its validity, which can be checked recursively.

### 4.1 Data Centre Use Case

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.

We can model a simplified scenario for data centre emissions reporting as illustrated in Fig. 1. In this scenario, we only consider electricity as the source of emissions. In reality, there are other sources such as embodied emissions from hardware manufacturing, for example, but electricity supply is currently the dominant factor in cloud computing emissions [36].

To convert energy use (measured in kWh) into emissions (measured in kgCO<sub>2e</sub>) we also need to know the *carbon intensity* (measured in kgCO<sub>2e</sub>/kWh). The carbon intensity of the electricity grid varies by time and place (for example, it is lower when lots of renewables are available), and it is affected by the datacenter operator’s commercial arrangements (e.g. commitments to buy a certain amount of energy from a low-carbon supplier, or the operator may even generate their own power). For simplicity, we assume we can use grid average carbon intensity data provided by the data centre’s electricity supplier.

As a further simplification, we focus on the emissions of services that are the direct result of a particular customer’s usage of data centre resources. We leave for future work the question of how to deal with large fixed-cost emissions, such as those from training the AI models, which need to be amortised over the useful lifetime of the model.

<sup>3</sup><https://www.wbcsd.org>

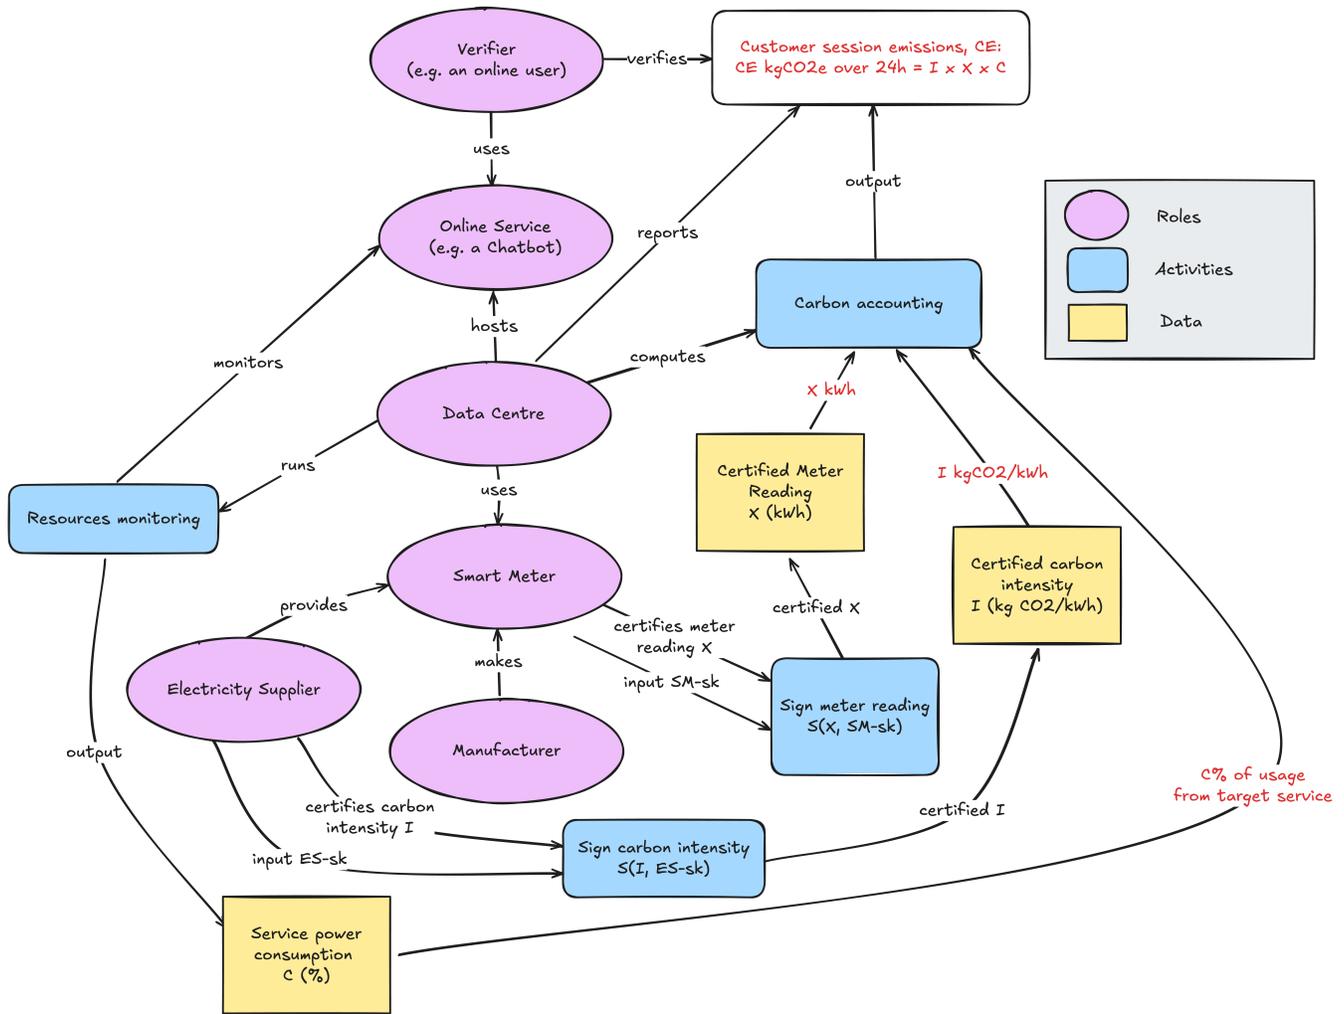


Figure 1: A simplified carbon accounting flow overview

We can thus define the carbon emissions calculation for a user of cloud computing services as follows:

$$CustomerEmissions = Intensity \cdot TotalEnergy \cdot Share$$

or as presented using one-letter symbols on Fig. 1:

$$CE = I \cdot X \cdot C$$

The data centre operator should be able to send each of their customers the value *CustomerEmissions*, the carbon emissions resulting from that customer’s resource usage, without that customer learning anything about other customers’ emissions or the data centre’s total emissions. Let *Share* be the fraction of the data centre’s total energy consumption allocated to a particular customer over the reporting period. The carbon intensity could be public knowledge if a grid average is used, but it might be sensitive if it reflects commercial arrangements with electricity suppliers. The data centre’s total energy consumption *TotalEnergy* and a particular customer’s fraction of it *Share* are commercially sensitive because they reflect the data centre operator’s business volume and profitability.

All the figures are averages over the same fixed period of time, e.g. over 24 hours.

## 4.2 Verifiable Inputs

To make it more difficult to cheat on the electricity usage figures, we can assume that the data centre operator uses smart meters with a secure hardware element that signs meter readings with a private key configured by the meter manufacturer. To ensure that the public key used to verify the meter reading signatures represents a genuine meter, we need the smart meter’s public key to be signed by the meter manufacturer. The meter manufacturer’s public key can in turn be signed by a trusted Certificate Authority (CA) that has checked that the public key belongs to a reputable manufacturer of smart meters. Similarly to a TLS certificate chain, this approach can be used to verify signed meter readings and prevent tampering and forgery. We have to assume that the smart meter is correctly installed and not bypassed, but since the meter

readings are also needed for billing, it is in the electricity supplier's interest to check the correct installation of the meter.

The same applies to the carbon intensity metric used in the accounting. We assume that the carbon intensity has been signed by the electricity supplier, and the public key of the electricity supplier has been signed by a trusted CA, which confirms that the key belongs to a genuine electricity supplier. Fig. 2 shows these certificate chains. In principle, electricity suppliers could use another zk-SNARK to prove that the carbon intensity figures have been calculated accurately, and the data centre operator's proof can recursively attest to the validity of the carbon intensity proof. We plan to explore recursive ZKPs in future work.

A zk-SNARK then allows a cloud provider to prove that there is a valid signature chain for every meter reading and every carbon intensity figure used in the emissions accounting, but without disclosing the value of those meter readings, the identity of the electricity supplier, or any other internals that might be sensitive. Only the public key of the CA needs to be disclosed, since it serves as the trust anchor for the computation.

A customer's share of total energy consumption is internal and sensitive data, therefore should not be disclosed. One way to check that each customer's share of usage is accurate is to apply the 'Completeness Principle' defined in the Greenhouse Gas Protocol [32], namely that the entire emissions of the data centre must be allocated to customers. That means that if we can prove that all customers' shares from the same data centre over the same period of time add up to 100% of the power consumption, and that the verifier's share is one of them, then we have reasonable confidence that the input metric can be trusted. This extended proof is beyond the scope of this paper, but it is important to consider it as future work (see §5).

### 4.3 Applying zk-SNARKs

Now that we have defined what the witnesses are (input data) and what we need to prove (signature verification and carbon emission calculation), we can then apply a zk-SNARK in three stages, as illustrated in Fig. 3.

*Stage 1: Define computation.* In traditional zk-SNARKs, we must define a circuit that encodes the accounting methodology used to calculate the emissions and any verification that needs to occur, such as checking the signatures in a certificate chain. The relationships among the input parameters are reduced to a set of polynomial equations, or *constraints*. Domain-specific languages such as Circom [6] can be used to define the circuit in code, and the circuit is made public. This is illustrated in Fig. 3 as the centre box showing each proof generation on the left, and the public output, labelled "Emission Claim Proof" on the right.

Alternatively, there are also zero-knowledge virtual machines (zkVMs), such as SP1 [25] and RISC Zero [33], that offer a more developer-friendly alternative to writing circuits by executing programs written in a conventional programming language such as Rust. These frameworks are widely used in the cryptocurrency community. We are exploring both the use of zkVMs and circuit-based techniques.

*Stage 2: Generate proof.* The prover proves that it knows a set of witnesses that satisfy all the constraints, and it can choose which of the witnesses to disclose and which to keep private. Public keys and signatures from smart meters, meter manufacturers, and electricity suppliers do not need to be disclosed to the verifier.

The prover in zk-SNARK requires more computational resources than the verifier. The proof can be constructed using one of the several zk-SNARK protocols, such as Groth16 [19], which is a pairing-based zk-SNARK widely used for its succinctness. Groth16 requires a trusted setup before a proof can be generated, but there are also alternative constructions that do not require any trusted setup. The prover then takes the proving key and witness to compute a proof.

In our example scenario, the cloud provider acts as the prover. They input the carbon emissions claim, carbon intensity for the electricity used by the data centre, the total power consumption, the signatures, public keys, and any other required data to create a proof that the witnesses satisfy the equations defined in the circuit.

An example of input, assuming that the public key of the smart meter manufacturer is signed by a CA using the EdDSA scheme [23]:

```
{
  "customer_emission": "xxxxx", (in kgCO2e)
  "carbon_intensity": "xxxxx", (in kgCO2e/kWh)
  "total_consumption": "xxxxx", (in kWh)
  "customer_share": "xxxxx", (in %)

  "ca_pk_x": "xxxxx",
  "ca_pk_y": "xxxxx",
  "manufacturer_pk_signature_r_x": "xxxxx",
  "manufacturer_pk_signature_r_y": "xxxxx",
  "manufacturer_pk_signature_s": "xxxxx",
  "hashed_manufacturer_pk": "xxxxx"
}
```

*Stage 3: Verify.* The verifier uses the proof and the knowledge they have to determine whether they can believe that the prover has knowledge of all input data, such that the private and public witnesses satisfy all the equations encoded in the circuit. That is,

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

The customer in our example, as verifier, can use the generated proof and the shared knowledge, namely the carbon emissions claim and the CA's public key, to verify the truthfulness of the claim, as shown at the bottom of Fig. 3. The public witness sent to the verifier to the verification therefore contains only three numbers in our example:

```
{
  "customer_emission": "xxxxx", (in kgCO2e)
  "ca_pk_x": "xxxxx",
  "ca_pk_y": "xxxxx"
}
```

A more fully-fledged example would need to contain several more fields, such as the start and end timestamps of the reporting period, the identity of the data centre operator, and the ID of the customer for which the report has been generated.

## 5 CAVEATS AND FUTURE WORK

The main challenge with the proposed approach is the accuracy and quality of the source data. A ZKP alone cannot guarantee

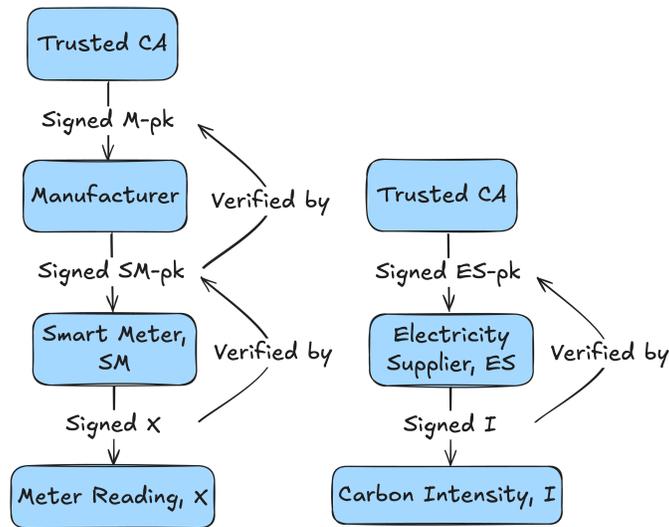


Figure 2: Signature chains for verifiable meter readings and carbon intensity

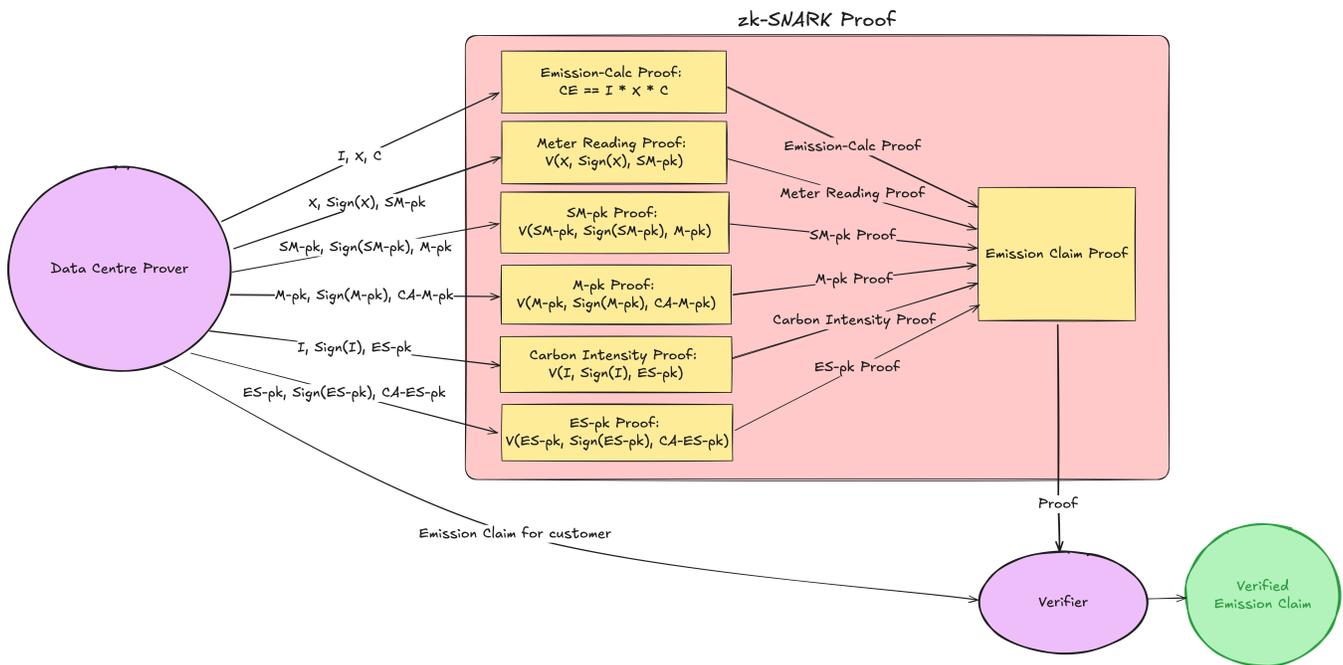


Figure 3: zk-SNARK proof system for private data protected carbon emissions claims. Private data:  $I$ : Carbon intensity,  $X$ : Data centre’s total electricity consumption,  $C$ : Customer’s share of the total electricity consumption,  $Sign(n)$ : signature of  $n$ ,  $SM-pk$ : Smart meter’s public key,  $M-pk$ : Manufacturer’s public key,  $CA-M-pk$ : CA’s public key for verifying the signature of manufacturer’s public key,  $ES-pk$ : Electricity supplier’s public key,  $CA-ES-pk$ : CA’s public key for verifying the signature of electricity supplier’s public key. Knowledge of verifier:  $CE$ : Emission claim

everything. For example, if a cloud provider gets meter readings from multiple data centres, there is a chance that they might input the wrong ones or leave out the reading from some meters. If the meters are incorrectly connected, the readings would be wrong. The carbon intensity figures could also be applied inaccurately if they

were meant for a different region. These problems would require manual checks outside of the ZKP. Separate validations to check that the data is consistent across multiple data centres would be useful, and it would make the validation stronger if the carbon-related data is linked with financial data, such that discrepancies

would arise if incorrect data has been used for one side of the equation. For example, a ZKP can show that the total paid to all suppliers matches the number reported on the audited accounts and that these supplier transactions are also reflected in the emissions calculation.

With regard to smart meters themselves, previous work has looked at secure hardware in more detail. For example, Karakashev et al. looked at using secure hardware for trustworthy renewable energy certificates [24]. Human auditing or random sampling verification can also be used to provide extra reassurance.

Another potential concern is the energy consumption allocation. There are shared resources, such as lighting, cooling, networking equipment and other ancillary energy consumers, that cannot be allocated directly to customers based on resource usage. One way to deal with this is to separate the emissions from customers' usage and shared or fixed resources, and to allocate shared resources across all customers proportionally to usage. For example, Google's carbon accounting methods first allocate emissions from electricity consumption per customer, and then augment the figures with proportional allocations of emissions arising from the non-electricity sources [11].

Companies could also make up customer shares, allocate shares dishonestly, or create fictitious customers and allocate emissions to them. Having a requirement for a public customer-based transparency log would help to make it difficult for companies to cheat. With a transparency log, customers could look up their entry to validate that they have been accounted for, and the data could be encrypted to protect confidentiality.

There are still open questions about the proposed approach. First, for verifiable data exchange to work, it relies on standardised accounting schemes being adopted by the stakeholders, but how do we encourage companies to commit to these schemes in the first place? The good news is that there are emerging standards for exchanging emissions data between companies. WBCSD (as mentioned in §3.1) is leading the effort and has produced a set of standards for emissions data exchange. However, the data exchange methodology does not currently involve cryptographic verification. The proposal presented in this paper could be extended and integrated into their framework to achieve data quality and reliability.

Second, the example used in this paper only considers one meter reading, one signature from each party on the chain, for one customer's carbon emissions claim. The proof system has to be able to scale when we consider many customers and frequent meter readings. Performing proofs involving many customers and many meter readings in a single zk-SNARK circuit quickly runs into scalability limitations; we are therefore investigating methods for breaking down such large proofs into many smaller proofs, which are more computationally feasible.

Third, as mentioned in §4, the proposed approach needs to be extended to include a way to verify that all the customer's share of power consumption add up to 100%. This can be achieved in multiple possible ways, but the details are beyond the scope of this paper.

## 6 CONCLUSION

It seems almost impossible to balance privacy and competitiveness needs with our urgent sustainability goals to reduce emissions, particularly in the ICT sector. The approach outlined in this paper is an example of how we can achieve privacy-preserving and trustworthy carbon emissions claims for data centres. To adopt the ZKP system, companies can apply carbon accounting alongside their financial accounting, which already needs to attribute the use of computing resources to individual customers for billing purposes. zk-SNARK proofs are small enough (a few kilobytes) 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 CO<sub>2e</sub> 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.

The approach can also be used in other industries. For instance, in an automobile manufacturing supply chain, motor part manufacturers would not want to disclose where their factories are or some of the metrics regarding the production of their products. Equally, in a coffee making supply chain, coffee beans suppliers would not want to make public where they source the beans. Companies typically have a lot of confidential data that are part of the carbon emissions calculations, and it is important that they can disclose verifiable Scope 1, 2 and 3 emissions to earn trust and for regulatory reasons, and still be able to protect sensitive business data.

## REFERENCES

- [1] 3Degrees. 2022. *3Degrees Cloud Services Review Statement*. Google. [https://services.google.com/fh/files/misc/3degrees\\_cloud\\_services\\_review\\_statement\\_final.pdf](https://services.google.com/fh/files/misc/3degrees_cloud_services_review_statement_final.pdf)
- [2] ACM. 2021. *ACM Technology Policy Council Releases TechBrief on Computing and Carbon Emissions*. ACM. <https://www-acm-org.ezp.lib.cam.ac.uk/media-center/2021/october/tpc-tech-brief-climate-change>
- [3] Amazon. 2024. *Notice of 2024 Annual Meeting of Shareholders & Proxy Statement*. [https://s2.q4cdn.com/299287126/files/doc\\_financials/2024/ar/Amazon-com-Inc-2024-Proxy-Statement.pdf](https://s2.q4cdn.com/299287126/files/doc_financials/2024/ar/Amazon-com-Inc-2024-Proxy-Statement.pdf)
- [4] LLC Apex Companies. 2021. *Microsoft Sustainability Calculator 2 GHG Methodology Assurance Statement*. [https://download.microsoft.com/download/9/d/e/9de2b0ba-45f1-4f6c-bb82-1cddb4fc54e5/Microsoft\\_Sustainability\\_Calculator\\_2\\_GHG\\_Methodology\\_Assurance\\_Statement.pdf](https://download.microsoft.com/download/9/d/e/9de2b0ba-45f1-4f6c-bb82-1cddb4fc54e5/Microsoft_Sustainability_Calculator_2_GHG_Methodology_Assurance_Statement.pdf)
- [5] 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.
- [6] Marta Bellés-Muñoz, Miguel Isabel, Jose Luis Muñoz-Tapia, Albert Rubio, and Jordi Baylina. 2022. Circom: A circuit description language for building zero-knowledge applications. *IEEE Transactions on Dependable and Secure Computing* 20, 6 (2022), 4733–4751.
- [7] 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*. IEEE, 459–474.
- [8] 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.
- [9] 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.

- [10] 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.
- [11] Google Cloud. 2025. *Carbon Footprint reporting methodology*. Google. <https://cloud.google.com/carbon-footprint/docs/methodology>
- [12] European Commission. 2023. *Proposal for a Directive on Green Claims*. European Union. [https://environment.ec.europa.eu/publications/proposal-directive-green-claims\\_en](https://environment.ec.europa.eu/publications/proposal-directive-green-claims_en)
- [13] Benjamin Davy. 2021. *Should estimating the power consumption of AWS EC2 instances really be quite this hard?* <https://www.thestack.technology/power-consumption-of-aws-ec2-instances/>
- [14] Will Evans. 2022. *Private Report Shows How Amazon Drastically Undercounts Its Carbon Footprint*. <https://revealnews.org/article/private-report-shows-how-amazon-dramatically-undercounts-its-carbon-footprint/>
- [15] 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>
- [16] Erol Gelenbe and Yves Caseau. 2015. The impact of information technology on energy consumption and carbon emissions. *ubiquity* 2015, June (2015), 1–15.
- [17] 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.
- [18] Gov.uk. 2018. *The Companies (Directors' Report) and Limited Liability Partnerships (Energy and Carbon Report) Regulations 2018*. Gov.uk. <https://www.legislation.gov.uk/uksi/2018/1155/contents/made>
- [19] Jens Groth. 2016. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology—EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8–12, 2016, Proceedings, Part II* 35. Springer, 305–326.
- [20] Jonathan Heiss, Tahir Oegel, Mehran Shakeri, and Stefan Tai. 2024. Verifiable Carbon Accounting in Supply Chains. *IEEE Transactions on Services Computing* 17, 4 (2024), 1861–1874. <https://doi.org/10.1109/TSC.2023.3332831>
- [21] Daira Hopwood, Sean Bowe, Taylor Hornby, Nathan Wilcox, et al. 2016. Zcash protocol specification. *GitHub: San Francisco, CA, USA* 4, 220 (2016), 32.
- [22] Jiahui Huang, Teng Huang, and Jiehua Zhang. 2022. zkChain: An Efficient Blockchain Privacy Protection Scheme Based on zk-SNARKs. In *International Conference on Machine Learning for Cyber Security*. Springer, 400–410.
- [23] S. Josefsson. 2017. *RFC 8032: Edwards-Curve Digital Signature Algorithm (EdDSA)*. <https://datatracker.ietf.org/doc/html/rfc8032>
- [24] Dimcho Karakashev, Sergey Gorbunov, and Srinivasan Keshav. 2020. Making renewable energy certificates efficient, trustworthy, and anonymous. In *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 1–7.
- [25] Succinct Labs. 2025. *Succinct Processor 1*. <https://docs.succinct.xyz/docs/sp1/introduction>
- [26] Duc Anh Luong and Jong Hwan Park. 2022. Privacy-preserving blockchain-based healthcare system for IoT devices using zk-SNARK. *IEEE Access* 10 (2022), 55739–55752.
- [27] Tobias Mann. 2024. *So much for green Google ... Emissions up 48% since 2019*. [https://www.theregister.com/2024/07/02/google\\_datacenter\\_emissions/](https://www.theregister.com/2024/07/02/google_datacenter_emissions/)
- [28] Microsoft. 2025. *Emissions Impact Dashboard*. <https://www.microsoft.com/en-us/sustainability/emissions-impact-dashboard>
- [29] David Mytton. 2020. Hiding greenhouse gas emissions in the cloud. *Nature Climate Change* 10, 8 (2020), 701–701.
- [30] United Nations. 2024. *Greenwashing – the deceptive tactics behind environmental claims*. <https://www.un.org/en/climatechange/science/climate-issues/greenwashing>
- [31] 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>
- [32] GreenHouse Gas Protocol. 2025. *Standards and Guidance*. <https://ghgprotocol.org/standards-guidance>
- [33] RiscZero. 2025. *RiscZero*. <https://dev.risczero.com/api/zkvm/>
- [34] Dan Robinson. 2024. *Microsoft's carbon emissions up nearly 30% thanks to AI*. [https://www.theregister.com/2024/05/16/microsoft\\_co2\\_emissions/](https://www.theregister.com/2024/05/16/microsoft_co2_emissions/)
- [35] 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>
- [36] 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).
- [37] Emily Sommer, Mike Adler, John Perkins, Joshua Thiel, Hilary Young, Chelsea Mozen, Dany Daya, and Katherine Sundstrom. 2020. *Cloud Jewels: Estimating kWh in the Cloud*. <https://www.etsy.com/codeascraft/cloud-jewels-estimating-kwh-in-the-cloud/>
- [38] Open source community. 2025. *ZK Jargon Decoder*. <https://zkjargon.github.io/definitions/succinct.html>
- [39] Amazon stakeholders. 2018. *Request from Amazon Stakeholders to Board Directors on Climate Change*. [https://drive.google.com/file/d/1\\_WprgbeUpic7\\_t8ovirFsmzxf\\_J0w6h1/view](https://drive.google.com/file/d/1_WprgbeUpic7_t8ovirFsmzxf_J0w6h1/view)
- [40] 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/>
- [41] European Union. 2022. *Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 amending Regulation (EU) No 537/2014, Directive 2004/109/EC, Directive 2006/43/EC and Directive 2013/34/EU, as regards corporate sustainability reporting (Text with EEA relevance)*. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022L2464>
- [42] 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.