Digital transformation of the circular economy: digital product passports for transparency, verifiability, accountability

Leandro Navarro ✉ 🏠 iD
Universitat Politècnica de Catalunya, Spain

Javier Cano Esteban ✉ 🏠
Universitat Politècnica de Catalunya,

Marc Font Miralles ✉ 🏠
Universitat Politècnica de Catalunya,

David Franquesa Griso ✉ 🏠
Universitat Politècnica de Catalunya,

Abstract

While overshooting the planetary limits, billions of ICT goods are sold annually in an increasingly digitalised world. ICT is part of the problem but part of the solution to digitally transform our society to meet climate change goals. The world needs sustainable digital devices: efficient, durable, reusable, respectful to people and the planet, and accountable for their impacts. That requires trusted, detailed information to inform the best decisions and actions in a circular economy, and verifiable data and content about actions. We explore the design, implementation, evaluation and operation of a verifiable registry for digital product passports of ICT products, supported by blockchain technology. Our experimental results confirm the design decisions and the feasibility of delivering these services efficiently and at scale. Therefore it confirms digital product passports as viable instruments to move the ICT sector to the track of transparency and environmental accountability, to become an example for other product sectors to meet the climate change goals the world cannot afford to miss.

2012 ACM Subject Classification Networks → Location based services; Applied computing → Enterprise data management; Social and professional topics → Governmental regulations; Hardware → Impact on the environment; Social and professional topics → Sustainability

Keywords and phrases circular economy, digital transformation, digital product passport, verifiability, blockchain, distributed ledger

Category Pending publication

Acknowledgements This work was partially funded by the NGI Trublo project, NGI Atlantic, Spanish Government under contracts PID2019-106774RB-C21, and PCI2019-111851-2 (LeadingEdge Chistera). Many organisations and people have participated in eReuse circuits and have contributed to improving the model and software tools.

1 Introduction

More than 6 billion new ICT goods are sold annually worldwide. There are estimates of 1.5 billion smartphones [24] in 2021, 126 million desktop computers, 659 million laptops, and 513 million Wi-Fi routers produced every year (2021). These numbers will probably grow exponentially over the next five to ten years with new "smart" technologies [18]. That has a significant impact on the environment as raw materials consumption during manufacturing, consumption of electricity during use, as well as e-waste and pollution at the end of life.

In contrast, science says decarbonisation must tackle the environmental crisis and comply with the 1.5°C global warming objective described by the IPCC Special Report (Paris Agreement). In 2020 the ITU-T L.1470 [17] recommendation made public a considerable challenge for the ICT sector: a dramatic reduction of the environmental impact of about 50% is required by 2030 to align with the 1.5°C climate change trajectory.

ICTs have an environmental impact at each stage of their life cycle, e.g., starting from energy and natural resource consumption and ending in e-waste. In contrast, ICTs can enable vast efficiencies in social and economic life through digital solutions that can improve energy efficiency, inventory management, and efficiency by reducing travel and transportation, e.g., telework and videoconferencing, substituting physical products for digital information. This capacity is second-order or enablement effects.

The circular economy (CE), and the term circularity, is about "designing out waste and pollution, keeping products and materials in use, and regenerating natural systems" [6]. In the context of digital devices, circularity aims at achieving the best use of devices with an extended lifespan.

The digitalisation of the circular economy of digital devices enables these second-order efficiencies in the management of ICT devices. In addition, digitalised information about devices brings various benefits, including efficiency, transparency and accountability. This digitalisation will enable the management of ICT devices from informal to formal environmental accountability, as the sustainable product initiative is proposing [11].

People and organisations relate to the digital devices they use. They buy, rent, use, repair, enhance, sell, donate, and recycle them. In addition, there are businesses focused on managing ICT assets by volume, such as manufacturers, IT asset disposition (ITAD), repairers or recyclers.

These actions are supported and can be testified by user-generated digital content (pictures, geolocation, tags, snapshots, documents) as proofs that can confirm these actions. Histories of these multimedia proofs associated with devices can confirm and trace of the relevant events along with the lifespan of an ICT device. These histories and related data determine social, economic and environmental impacts. That becomes a kind of product CV or portfolio, carefully disentangled from the people's personal data. Devices are grouped by brand and model, and each instance is distinguished by serial number.

These unique instance/individual identifiers, derived from unique serials in the devices, the chassis, and several serialised parts, can be registered as unique identifiers in a verifiable registry. Other relevant events along the life of a device can also be recorded in a verifiable registry, such as ownership transfer, repair, data wipe, reconfiguration, refurbishment, and recycling. Even conditional rewards in the form of an economic deposit or a reputation boost in exchange for future actions, an extended responsibility, such as returning a device after use, or funding recycling by the manufacturer, can be recorded in a ledger.

We propose an inventory of details about products in the form of standardised digital product passports [10] and a ledger of device lifecycle events linked with social interactions (the actions mentioned before). The ledger is recorded in a verifiable registry, metaphorically equivalent to a notary public, that offers transparency and accountability about the detailed data. The verifiable registry relies on an append-only distributed ledger (blockchain), with the ability to apply agreed rules about procedures when a condition is met (smart contracts, inexorability). This registry will allow social, economic and environmental impacts to be reliably determined, even generated as impact reports.

This work has designed, developed and evaluated a verifiable registry. We have integrated it into DeviceHub [7], an open-source device inventory system that acts as a client and intermediary on behalf of human device owners that can record supporting digital details and content as well as generate and deliver digital product passports. We have experimented with the testing and evaluated the performance and guarantees of that registry service. We have developed specifications such as naming schemes, API, and metadata to facilitate the management of devices.

The main innovations come from the system design (architecture, implementation, integration of multiple blockchains). We specify tools and data format to represent device characteristics as they change over the product lifespan (predictable format, by standardising data items and order) to ensure verifiability. Regarding data formats, we are producing innovations in

- Defining a naming scheme for relevant entities: chassis details to derive unique identifiers (CHID), detailed characteristics of product hardware identifiers (PHID), according to the W3C DID specification [28].
- The definition of a DID lookup mechanism for the previous identifiers validated by a demo implementation.
- The definition of verifiable details and a verifiable registry API for circular devices.
- The definition of basic Digital Product Passport (DPP) documents for digital devices.
- The definition of smart contracts implements a verifiable registry that implements agreed procedures for the management of devices.
- The collection of diverse content linked (not embedded) to blockchain transactions for increased verifiability and trust.

We expect the digital transformation of ICT device management will bring the necessary accountability and trust to assess these devices' social, environmental and economic impacts while enabling and promoting more circular and sustainable practices. These models, systems, and specifications are a critical need to support regional or global policies to meet climate change goals that many countries, especially in Europe, as part of the different initiatives related to meeting climate change goals, such as the European Green Deal [3].

2 Related work

Life-cycle assessment (LCA) is a solid research field [2] on assessing the life-cycle of electronic devices as a significant element of assessment of environmental efficiency that we can leverage. It accounts for environmental provenance, and impact assessment is difficult as it requires reliable and trusted sources of information, such as recording diverse factors (about materials, energy, parts and devices) [16]. There have been innovative research about the circular economy of material devices in our research group over the last eight years, about the transition to a collaborative and circular consumption of electronics, such as [13], or distributed ledger technologies applied for the traceability of devices and services in crowdsourced networks as [5], both at UPC. This research has produced software tools and services (eReuse.org) that collect (privacy-preserving) data about devices, as well as assist in preparing devices for refurbishment, and automate the recording of accountability data (e.g. reporting of crucial operations in the lifespan of devices).

There are multiple research, experiments and pilot services related to the traceability of goods using distributed ledger technology. The requirements and functionalities of supply chain integration are explored in [4]. The circular economy has its challenges in practice [19]. Incentives for rewarding cooperative behaviours are an opportunity for distributed ledgers [22] and resulting implications for sustainability and social responsibility [26]. However, we have to relate "real world" multimedia information that becomes proofs for recorded transactions in the ledger through the concept of an oracle [1].

The Product Circularity Data Sheet (PCDS) [21] and [20] is an initiative that provides a public specification for a primary source of verifiable data about how manufacturers design their products. It can help establish how circular a product is and inform about the circular path it was designed and manufactured. The PCDS offers a standardised format with trustful data without scoring or ranking these aspects. It has three objectives: provide basic data on product circularity, improve the sharing efficiency of circularity data, and encourage the circularity performance of products. It is inspired by the (material) safety data sheet [25]. At each stage of product manufacturing or transformation, a new PCDS is created. Every supplier passes the relative PCDS one step up the supply chain to allow for its integration into the next tier's product. Each manufacturer is responsible for storing the information related to the PCDS statements and for making such information accessible to other stakeholders upon request. PCDS is designed to be integrated through the supply chain. Our work builds on these principles. The digital product passport we propose focuses on closing the gap to achieve a circular economy that resolves the lack of trustworthy, verifiable, useful information to facilitate use, reuse and recycling. DPPs are expected for many product categories. In the EU, this is part of the Sustainable product initiative in the Circular economy action plan [8] of the European Green Deal [3], with electronic products and batteries being the first, as the new EU Battery Regulation [9].

With eReuse.org and the IOTA Foundation, we have explored the concept and implementation of the digital product passport to deliver guarantees of efficiency, transparency, trust and accountability. As reported in [15], we are exploring with Alastria.io and the IOTA Foundation to integrate decentralised identity and verifiable credentials for participants in our verifiable registry. In collaboration with the Obada Foundation [12], we have explored the concept of physical NFT. We are exploring the sustainability requirements for global digital product passports with ITU-T [23].

We have experience with experimentation and using permissioned blockchain infrastructures, such as Ethereum PoA, or the T Network in Alastria.io, in our EU NGI Ledger participation. In addition, we have managed to set up experimental permissioned blockchains with a minimal environmental impact footprint by using low power servers in a permissioned model with highly efficient consensus algorithms based on the proof-of-authority model and with even asynchronous operation (block time 0).

We go beyond state of the art and our research group experience and results by extending our circular ledger to allow the collection of off-chain diverse media linked (not embedded) to blockchain transactions and the digital product passport. That will allow genuine content to be associated with these transactions. The result will increase quality and trust in the traceability and impact reports about devices over a circular lifespan. That enables participants and third parties to verify the claims, increase the reputation of the participants (individuals and organisations) and enable incentives as rewards.

3 Model and design

In a circular economy, the lifecycle of digital devices can go as follows: after using raw and secondary materials to produce parts, devices are assembled at factories and sold by brands. These devices usually have unique identifiers (serial numbers) that may come linked or labelled with details (information sheets) about their composition, characteristics, instructions for maintenance, repair, and even recycling. Devices can be repaired, upgraded, transferred internally to new use, or disposed of to be transferred (sold, donated) to a new owner, dismantled for parts, or recycled to recover secondary raw materials or dumped in a landfill. All these actions are related to events and documents that prove and document these actions. These miscellaneous details (multimedia content) can help the accountability and verifiability of these processes and motivate and reward human participants.

Digital support systems, like the eReuse ecosystem of tools, have the structure in Figure 1. Multiple organisations ($OrgA..OrgZ$) have multiple devices ($ID_1..ID_X$), with all details stored and updated as digital data in their organisational inventory system. Devices have tags with digital identifier codes using unique physical ID tags as data carriers to facilitate identification for tracking and handling these assets during the usage cycle, maintenance, and end-of-use phase before final disposal. Commonly, these physical tags include a written identifier and a machine-readable, optical (e.g., QR code) or electromagnetic (e.g., RFID, NFC) element to facilitate reading. Any maintenance, repair, upgrade, trading, reuse, and final decommissioning is usually associated with updates to an inventory system, while documents may accredit these actions. These documents can be grouped as "portfolios" per person (about all devices she had, actions done with them), per device (about all events in its lifespan), and per organisation (all devices owned and managed).

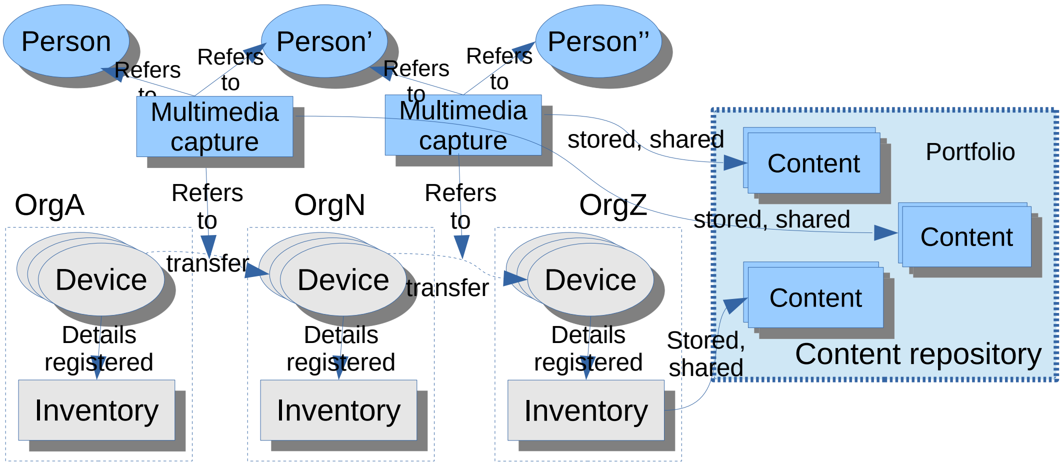

Figure 1 A data and system model for the lifecycle of digital devices, people and content

This model in Figure 1 shows the detailed information collected in the inventory that relates to or includes the DPPs of initial or modified products as mentioned in the next sector and described in Section 3.3.1. However, trust requires the ability to account for and verify that information. A verifiable data registry can provide this. According to W3C [27], it facilitates the creation, verification, updating, and/or deactivation of decentralised identifiers and DID documents, including cryptographically-verifiable data structures such as verifiable credentials.

3.1 A scenario

We identified the following workflow for a Consumer Electronics Digital Product Passport and the corresponding required interactions with a digital ledger infrastructure. That scenario translates into functional requirements:

- The *manufacturer*, reseller or first owner registers the *chassis* (a modular computer) on the digital ledgers using a unique Chassis ID (CHID). It also publishes its first DPP, which refers to the initial detailed hardware configuration of the product that includes a Product Hardware ID (PHID). A *DPP* for that specific configuration is identified by an ID composed of *CHID:PHID* identifiers.
- The *first owner* configures, switches on and registers the complete hardware information (a *hardware profile* including manufacturers, models, serials, characteristics of all identifiable elements). They do this running software, Workbench in our implementation, on the device. The Workbench software captures and sends the information to an inventory service of the owner organisation, DeviceHub (DH) in our implementation. A QR code gets printed and attached to the device as a *data carrier*. The DH backend must store detailed data and records in the ledger a summary (a *fingerprint*) of the captured information associated

with the same device chassis (including the URL and summary of supporting documents). The recorded hardware configuration results in a new DPP must include the new hardware composition, therefore identified as CHID:PHID[1].

- Any time, the current owner or any qualified device operator must be able to find out a *DPP document* from a DPP ID or even list all DPPs (hardware configurations) associated with a chassis ID over a lifespan. The inventory service must provide an ID lookup/search and deliver these documents, with full to no detail according to who the requester is, containing standardised and human and machine-readable data[2].
- Every time the product gets sent for repair, upgrade, refurbishment or remanufacturing, when the *hardware changes*, the product operator shall issue and supply a *new DPP* identifier and content.[3]
- After a few years, the device gets decommissioned, no longer used. The owner generates a *(proof of) data wipe* (using a specific data wipe software). The device is transferred and refurbished by a certified refurbisher. These *proofs*[4] about actions get stored in the ledger for verifiability.
- The new owner can optionally record a *proof-of-use* (hours used). The Workbench software runs and records usage hours and sends it to the inventory service (DH), recording linked to its chassis ID, with a proof sent to the ledger for verifiability.
- Years after, the new owner stops using the device and brings it to a certified recycling centre. The certified recycler shall record a *proof* that the device has started the *recycling* process in the ledger, and the recycler has acquired ownership of it.
- After a few months, a third-party *auditor* or the initial owner wants to see the history and check the events and lifespan of a device by querying its chassis ID. The search for this chassis ID on the *verifiable registry* (ledger) will bring them one or several DPP IDs and other details associated with that product (e.g. document IDs, summaries (hashes) and associated timestamps). The search of these IDs from the inventory service will result in documents[5]. These can be validated by matching identifiers and comparing them with hashes and timestamps retrieved from the ledger. That *double check* allows the matching of document summaries for integrity. That brings reliability and allows to confirm dates and other details, increasing trust and preventing fraud. A case of auditor actor is the product passport *registry* for authorities. Auditors shall even be able to generate a *sustainability impact report* with an estimate of extra usage hours beyond the first use and CO_2 impact, etc., counting to environmental accountability.

This scenario and requirements align with the requirements set in the Sustainable Products Initiative documents adopted by the EC in March 2022 [11].

The diagram in Figure 1, combined with this scenario and the W3C verifiable data registry model, results in a more comprehensive data and system model in Figure 2. The model relates content as documents associated with devices, recorded with metadata details in an inventory service by their owners. This service controls and delivers DPPs associated with decentralised identifiers. The verifiable data registry keeps a ledger record of identifiers for devices, people and documents, and proofs of relevant actions on devices. The entries in the ledger have

[1] Ledger selects human-centric projects.

[2] NGI Atlantic experiment: blockchain asset disposition alliance

[3] Tessera, Consensys.

[4] T-network: 126 regular (data storage), 9 blockmaker (validation), 3 boot (permissioning)

[5] Definitions from the UN ITU working document for the L.GDSPP work item: [23]

timestamps and block numbers as references. Identifiers for all relevant entities are common in all zones (documents, metadata and DPPs, verifiable registry).

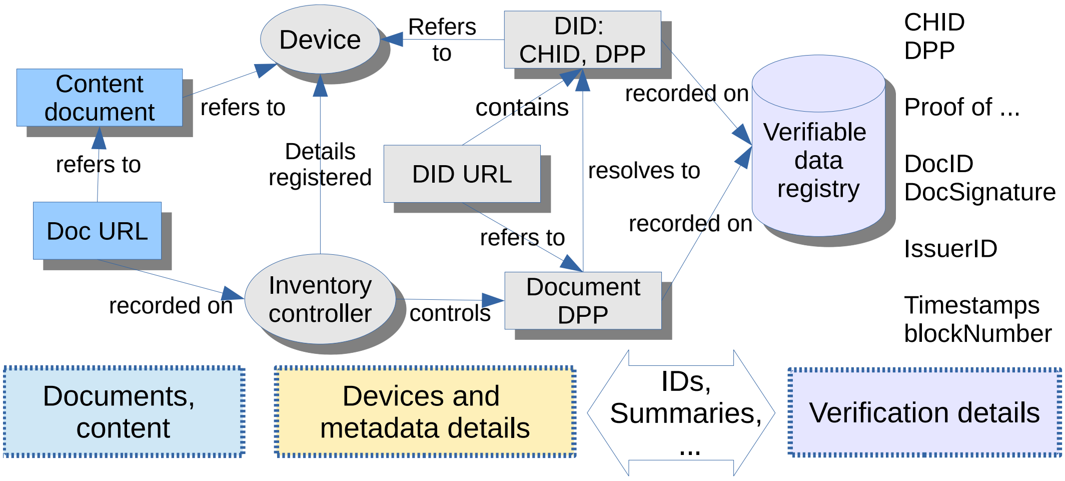

Figure 2 A data and system model for the documents, devices and metadata, and verification details

3.2 Safety properties

Diverse actors may be interested in this information for different purposes. While sometimes what matters is the information stored, other times what matters is certainty about the action done (verifiable proof, attestation), and the link between both is critical. The *global record of devices* (GRD) [13] is a comprehensive record of transactions about relevant events. Some of the main events about a device are manufacturing, purchase, repair, upgrade or modification, decommission, transfer, data wipe, sale, recycling, and loss. Reporting these events in an inventory can be helpful to its owner. Still, in a circular economy, a device may go through multiple actors and organisations along with its lifespan, and these may not trust each other or may even risk colluding. Therefore, beyond the inventory systems for device owners in each organisation, a common verifiable data registry is needed to record transactions and supporting claims that affect any device and its complete lifespan.

The *verifiability* requirement translates into *irreversibility* of recordings (those operations already recorded cannot the undone or modified, sometimes referred to as *immutable* in the sense of *append-only*). It leads to introducing data *replication* with ledger updates coordinated by a consensus majority decision to prevent any attempt of manipulation of ledger books by faulty or even malicious actors.

However, irreversibility and the ability to be accessed by multiple actors raise a requirement to preserve personal privacy and business confidentiality. Nothing in the ledger can be private or confidential, as it could not be removed without destroying the ledger completely. For that reason, most of the details about transactions in a ledger should contain verification information (e.g. proofs, hashes, signatures, timestamps) that enable the holder of any data to prove it was present or produced by the time a transaction was recorded: If a transaction record is a tuple (actor, signature, timestamp), an actor can prove they had that data at the timestamp instant, as the data was hashed or signed by that actor, and that can be repeated now for verification by the data holder.

Furthermore, agreements in a community about rules and procedures, such as eReuse or Obada, can be translated into code (smart contracts) encoded and executed in the distributed ledger virtual machine. That brings *inexorability*, the confidence that rules will eventually be automatically applied as agreed and implemented in computationally binding code, preventing the insecurity of discretional and different rules applied in the future after a recent decision (equivalent to retroactive changes in law, that can generate insecurity).

The original data, the details, can be stored in any private inventory database (for device owners) or a public multimedia object like a document, picture, geolocation, scanned document, etc. In contrast, the ledger only stores the summary information about transactions for verification as a link, hash, or signature. This global ledger log allows searching for events linked to a device and its lifespan across multiple organisations. That allows verifying traceability, impact information and other transactions about devices.

Combining verifiability details stored in a verifiable registry with details from data stored outside the ledger accounting books, including inventory data, links to documents, social networks, allows us to generate verified circularity and social impact reports and metrics about specific devices, as well as generate rewards, as tokens or credit, for all participants to promote positive behaviours, and ensure sufficient sustainability impacts (in economic, social and environmental terms).

3.3 ICT devices as unique products

Our product is a computer, so we do not cover pre-use and post-use, just the *use phase*. That includes procurement, sale, use, reuse, repair, and modification until final disposition (e.g., a recycler). In the use phase, devices can be used and transferred for reuse until they are no longer useful and deactivated for recycling or dumping.

ICT devices are products that usually have a unique identifier per instance. However, they can change detailed hardware composition over their lifespan. That common identity over their lifespan is defined by the tuple "manufacturer/model/serial", which we call chassis identifier or CHID. The different hardware configurations with additional or different components resulting from reconfiguration (e.g. modular devices with add-on cards, changes due to repair or upgrade) can be summarised as a product hardware identifier or PHID. Therefore the chassis (CHID) will be the same for each hardware configuration, and the details (PHID) will change. That means *one product (CHID) can have multiple associated PHID*, one for each configuration (different parts).

That leads to defining a naming scheme for main entities: chassis unique identifiers (CHID), detailed hardware characteristics CHID:PHID. These identifiers follow the W3C DID specification.

3.3.1 The digital product passport

A *Digital Product Passport* is digital data that describes a product. In more detail, it is a structured collection of product-related data with a predefined scope and agreed data ownership and access rights conveyed through a unique identifier. A DPP provides each actor access to verifiable digital information to use or operate on that product (all you may need to know about it, including links to documentation). A DPP can be *plural*, referring to either a set of items (a model, a batch) or *singular* (individual) associated with a single product (with a unique identity, a serial number, and perhaps a unique composition or configuration). We focus on singular ICT products. As they can change hardware configuration over a lifespan, a DPP can be referred to by a CHID:PHID identifier. A product, identified by its CHID, can have as many DPP as hardware configurations over its lifespan.[6]

A *product operator* is any actor that can transform and supply modified products. Therefore, it can supply the information a DPP conveys about them as a result of manufacture or other operations such as packaging, configuration, maintenance, repair, upgrade, refurbishment, remanufacturing, and even recycling.

DPP provision is the process and responsibility of collecting, creating, maintaining, validating, storing and delivering data from source(s) to targets, including the service setup and managing the data related to it.

A *DPP supplier* is any product operator that is also responsible for DPP provision (supply) the associated data that is part (included or linked) in a DPP.[7]

The aim is to keep a verifiable registry of the DPPs associated with an electronic product due to changes in specifications done by product operators during their lifespan. Along with the circular life of an electronic product, while some events affect the characteristics of a given electronic product (i.e., reconfiguration, refurbishment, resale, recycling), other events just affect the chain of ownership (chain of custody). The DPP is definitely about the first, not necessarily the second (ownership), privacy-sensitive too.

The detailed data about a product that is metadata and supporting documents such as manuals and certifications can be stored and/or linked to the product in the inventory service and therefore linked in the product's DID document, as Figure 2 shows. A search/lookup on a DID allows to find the corresponding DPP document in the inventory *(off-chain)*. A search on the verifiable data registry *(on-chain)* results on proofs about relevant events (DPP issuance, related CHID, related documentation and summaries as hashes) that may support audit and verification of claims. Therefore, DID can find and crosscheck information on both sides: descriptive details off-chain and verifiability details on-chain, clean from personal privacy and business confidentiality details.

At each stage of product manufacturing or transformation, a new passport has to be created as the *product changes*. Every DPP supplier may refer to the previous DPP for previous details, and the first for details from the manufacturer, and be responsible for storing and supplying the information related to the current passport's statements. Therefore, the *issuance/hosting* of the product passport (a new ID resolving to a URL for details) appears to be the agent's responsibility that modifies the device. (A DH instance does this in our system).

The *chain of custody* does not change the product; it changes the track record of the device across owners or users. The management and visibility of this information are open for discussion but may reside only in an inventory service and may or may not be reflected in a DPP. In other words, *traceability* changes related to the chain of custody may be recorded in a "supplement" to the passport but do not change the main part of the information if the device is not physically modified.

The passport must be *accessible* in machine-readable format (e.g. JSON-LD) and may also be in human-readable format (e.g. HTML or PDF). It can be a URL in an Asset inventory manager instance or a DID (W3C) that, routed through a portal, leads to the right server instance supplying the DID document. Then a separate URL per device is not needed (the example of DOI and the doi.org resolution service). A DID registered in a verifiable registry can be resolved over a content addressable lookup service, as done with content-addressable networks,

[6] Equivalent as passports for people: one single tax ID and fingerprints for life, but different passport IDs every few years as our "hardware configuration" or aspect changes over lifespan.

[7] `https://irtf.org/icnrg` and several RFC in `https://trac.ietf.org/trac/irtf/wiki/icnrg`

such as information-centric networking (ICN)[8], IPFS[9] or Bittorrent magnet links [29].

3.3.2 Identifiers

We want a bidirectional mapping of a device and its ID, a *metadata twin* for individual items, that is bidirectional:

- *Idempotent* (repeatable): Two independent hardware inspections on the same device at different times by different actors must result in the same ID. In other words, a collision of identifiers means we are in front of the same device.
- *Bi-univocal*: the mapping of ID to hardware is unambiguous.

A device is identified by its main serial number, but that does not reflect changes in hardware parts. The main *fully qualified serial* for a device (chassis ID) could be based on the common practice of combining the data strings about manufacturer + part-number + serial. That does not include all components; it just identifies its main serial. We have found so-called “clone devices” without a serial number or duplicated. We have recurred to a UUID or the MAC address of an Ethernet port in the motherboard. However, standardised agreements are needed to ensure that the twinning between a chassis and its ID is unique. These identifiers are name-based UUID (not time-based), such as RFC 4122[10] version 5.

The data strings need to be normalised to have repeatable results. That requires again standardised agreements to canonical names for manufacturers and models. The fully qualified serial with minimal details (3-tuple) determines the chassis ID. In contrast, the maximal details (n-tuple) determine the maximal details of a complete hardware configuration that must include at least the unique IDs of the changed or all serialised replaceable elements. Still, a standardised agreement is needed on the elements to include and their order. The result of these min or max data can be mapped to fixed-size values, which a hash function can achieve. A cryptographic hash function such as sha256 is a one-way function that is practically infeasible to invert. However, the binary result must be transformed into a readable text, such as Base64 (RFC 4648) or Base58 (Bitcoin) or Base85 (RFC 1924), and truncated to provide a shorter, more convenient fingerprint (as PGP does). Again, this algorithm requires standardisation[11].

These CHID and PHID become part of metadata twins of the physical device. A DPP has an ID that includes the *main fully qualified chassis ID + a hash of the IDs of some or all replaceable hardware*. That, expressed as W3C DID could follow a scheme like this[12]:

$CHID = baseX(hash(min components in order))
$PHID = baseX(hash(max/replaceable components in order))
did:eCHID:$CHID
did:eDPP:$CHID:$PHID

These DID identify devices and can be used both on-chain to record actions on a verifiable registry through its API, or for the lookup resolution on the off-chain inventory service over HTTP with redirection using URLs, resulting in a DID document (a JSON data document in our implementation).

[8] This is being pursued by the Obada Foundation [12] with its obit public specification to become a standard. For instance obit = version(0000)+trunc(base58(sha256(manufacturer + part_number + sha256(serial))), 12) + checksum for protection against reading errors.

[9] `https://datatracker.ietf.org/doc/html/rfc4122#page-13`

[10] IPFS.org

[11] If adopted by the obada Foundation, that namespace will be prepended: did:obada:

[12] For instance the Obada network or the Alastria network.

3.4 Relevant actors and roles

We have the following users of the overall system as actors and roles: the owner and user of an ICT device, the operators (e.g. installer, repairer, refurbisher, recycler, ITAD) for acting on devices, the internal (organisational environmental manager) and external auditors for verification.

The use cases make more sense when we consider a verifier role (an auditor) is involved. This is the case of verification of any circularity claim, such as impact assessment or traceability reports. Verification implies checking information, either through verification operations (of the integrity of a supporting document or through a summary or source with a digital signature) or from separate sources (third-party): one actor performs a task (operator), the other actor acts as a witness (with a document).

For example, someone wipes a disk drive (a human actor) and records that event accompanied by a deletion certificate. This certificate can originate from a human actor or a software agent (signed on behalf of the person who maintains or certifies the software).

Many patterns of operator-witness come from a human actor. He decides/acts, accompanied by a document generated by a software tool or a person accrediting the action and which are activated and registered together (in a single invocation to the inventory service).

There are three main generic roles in the domain of application of the circular economy of digital devices:

- *Operator:* operates on devices, as an issuer (who wants/can publish a DPP), register (importer, manufacturer, distributor, even buyer if not done before).
- *Witness:* provides an observation. When we want to record an observation, we provide a document (which provides testimony and verifiability).
  An operator can act as a witness if his action is accompanied by a document generated on behalf of them (by a software tool).
- *Verifier:* an auditor of any claim may need to confirm the claim supported by additional details that confirm the supporting details and facts.

These roles are played by the following specific stakeholders or actors in our circular economy of digital devices:

- *Manufacturer:* an *operator* that assembles a device and can act as or require a witness. It can register the existence of a new device and publish its DPP.
- *Distributor:* an *operator* in charge of commercial distribution and sale can act as or require a witness. It can register the existence of a new device and publish its DPP if not done by the previous one or if the device has been modified (reconfigured).
- *Consumer/user:* procures, uses and disposes of a device, can act as or require a witness of its usage. It can be a witness of proofs and documents, actions that do not change the hardware configuration. As a last resort, it can act as an operator to register the existence of a new device and publish its DPP if not done by previous roles.
- *Refurbisher/Repairer:* an *operator* that renovates a device (remanufacturing, when done by a manufacturer) can act as or require a witness. It can register the existence of a new device if not done previously and publish a new DPP when a device has been modified: anything replaced or added.
- *Dismantler:* an *operator* that collects and disassembles a device and can act as or require a witness. It can deregister the device.
- *Recycler:* an *operator* that processes elements to extract resources, such as raw materials and energy, and manages e-waste, disposing of it according to established procedures, can act as or require a witness. It can deregister the device if not done previously.

- *Auditor:* a *verifier* authorised to review and verify records' accuracy and ensure that organisations comply with laws and regulations. Essentially crosschecking the verifiable registry CHID, DPP, proofs with inventory and supporting documents. They protect organisations from fraud, point out discrepancies in accounting methods (including environmental impact), and occasionally help organisations identify ways to improve their operational efficiency.

Related to the DPP, device-centric, the main actors are:

- *Operators* that "create" the first DPP of a device, ideally the manufacturer, and those that create further DPP associated with a modified device as a result of repair, refurbishment, recycling.
- *Operators* that record any hardware modification "proofs" but not witnesses that record usage counters or data wipe, ownership transfers, in the domain of chain of custody.
- *Verifiers* that look at traceability, impact reports, market surveillance authorities, customs, etc.

3.5 Verifiability requirements

The information in the DLT should allow for the verification of claims. Given a fully qualified serial for a chassis (CHID), the registry provides verifiability records for cross-checking with a search on the inventory service, including the DPP associated with it, such as JSON data. The *verifiability* comes from cross-checking detailed data with verification data under the same identifiers.

As mentioned before, verifying device ownership (chain of custody) is not a direct concern (first-class information) of the DLT. Still, device inventory systems can record proofs of transfer for the verifiability of ownership-related transactions. However, the registry should not administer users (owners, custodians).

3.6 Proofs

Proofs result from the need to record a significant event for the circular life of a device. We consider three types of proofs.

1. *Proof without a document:* This allows to record a statement by an actor about a device at a given timestamp, with no document linked to it.
2. *Proof with a document:* Proofs with a reference to a related document (identifier and summary, no content) such as a purchase, repair or delivery note.
3. *Proof with hardware snapshot document:* Proofs that have a document generated with software (Hardware snapshot). These proofs provide evidence for a fact or the truth of a statement. The Hardware snapshot document is the machine-generated output of executing software (code) on a device. The integrity of the output must be preserved from its execution on the device to the storage of a summary in a verifiable registry. The integrity of the code that results in the proof must be preserved from the open-source repository to the device that runs it. For example, the proof of data erasure is the result of the application of data erasure software. If the erasure was successful, then the proof will include the value of "success", among other details.

4 Technical implementation

We have developed a proof of concept prototype of the above model and design.

We are familiar with smart contracts representing *non-fungible tokens* (NFT ERC721), *rewards* as ERC20 tokens, and smart contracts that record simple proofs as transactions linked to reported actions with associated supporting documents.

We have designed and developed a verifiable registry in a permissioned DLT. This registry service is controlled by a group of trustee organisations, with partial public access to reports[13]. That allows *linking actions with media proofs and DLT entries* to bring trust and reputation to circular behaviours while preserving personal and organisational privacy. That builds and goes beyond our experience with basic permissioned blockchains that record simple device transfer events to generate lists of traceability transactions associated with a device.

The *user-contributed content* (media about the trustworthiness of actions and traceability of devices) requires a *content repository* to store and link content identified by IDs and URLs like third-party web, social media platforms or decentralised content such as in IPFS.

We have experience developing a mobile Android application that scans QR codes as data carriers, can capture images related to devices and people (a multimedia portfolio) to be stored in a content repository, and allows us to invoke a backend to report that data as blockchain transactions (proofs) as *hashes or digital signatures.* These transactions may include linked data between media, devices, and proofs (one-way data like hashes from private keys) that contributed by participants allow them to claim impacts and rewards while preserving personal and organisational privacy. However, given its proven feasibility, we have not extended that mobile Android application to integrate it in our proof-of-concept implementation.

The combined results allow us to validate our hypotheses and adjust the features of the complete system to assess the future potential for adoption, impact, sustainability and scalability while preparing for a pilot with users.

4.1 Technical elements

The minimal *end-user application* takes a picture, scan a QR code, or executes code on the device (WB) to extract details about a specific device. This data is submitted to a server acting as a *content repository* for media and to a *device inventory* service provider for device details. The inventory backend contacts a *verifiable data registry* provider to record relevant operations through an API. The backend API instance acts in the name of the right intermediary (using its *keyring* of private keys stored there) to invoke smart contract operations running on a permissioned ledger that follows agreed procedures to record and act upon these operations. A device inventory service instance can resolve DIDs such as DPP or CHID to linked data about it, or HTTP redirect to another instance.

The minimal *content repository* is the device inventory system for device details and a web server can be included for other multimedia content or an IPFS node for decentralisation and replication.

The experimental *permissioned ledger* testbed builds on the eReuse-Ledger testbed developed in the NGI Ledger (2019-2020)[14] project with an ERC721 NFT contract; conditional rewards with an extended ERC20 contract, and an API for device traceability developed in

[13] For privacy and confidentiality, we may move to disposable identities. For example, in `https://datatracker.ietf.org/doc/html/rfc5876` for SIP telephony, it is roughly defined as a mail address: P-Asserted-Identity and has a mechanism where a proxy can hide the true identity. More in `https://datatracker.ietf.org/doc/html/rfc3323#section-5.3`

[14] E.g. BASE64URL(JWS Signature) `https://datatracker.ietf.org/doc/html/rfc7515` details: `https://en.wikipedia.org/wiki/JSON_Web_Signature`

the NGI Atlantic[15] (2021) project. This testbed provides a verifiable data registry with a backend composed of a DLT (Ethereum PoA using 3-5 geth instances, with Solidity smart contracts). It keeps an internal record, DLT addresses, external references, and device DID. We use in the testbed a Prometheus daemon to collect logs and a Grafana instance for queries and visualisation of log data about usage and the experiments.

The baseline *experimental API and testbed* has incorporated new API calls for device management functions and smart contracts to deal with multimedia content proofs, portfolios, economic incentives, and a reputation system.

4.2 Data protection and privacy

The aim is privacy and confidentiality "by design":

- *On-chain data*: corresponds to global registered identifiers (product chassis CHID, product parts PHID, digital product passports DPP, clientID, documents) computed as hashes, similarly as document signatures (summaries as hashes). Data is meaningless without access to detailed data after an identifier lookup across device inventory systems and public or private access to details about devices, passports, and supporting documents. Similarly, ClientIDs refer to device inventory servers in addition to internal references, randomly generated by device inventory servers. That protects the identity of individual participants: the mappings are stored privately, and the same person or entity can have multiple identifiers, which means the system provides zero knowledge of sensitive information, either related to personal privacy or business confidentiality.
- The *private keys* for identities involved in DLT transactions (wallet) cannot be simply authorised one by one interactively by a person holding a secret private key. Instead, our verifiable registry API provider acts as a wholesaler (in the name of a retailer). For that, an API provider instance stores the private key(s) (storage in a keychain must be encrypted). Therefore, different API instances can act with different identities or private keys. As with banking operations, personal wallets could be contacted to synchronise validation to a person, preventing delegation of the private key to the API provider, but that would make every transaction interactive, with a human in the middle of the chain.
- *Off-chain data* corresponds to devices and other details stored by device inventory systems. No personal or business-sensitive information is publicly accessible.
- Data derived from data stored in inventory servers is *anonymised and aggregated* before being part of any *open public datasets*: such as summary information about the characteristics of devices registered in the system, such as brand, model, basic characteristics, and durability. [14]
- *Audit* information results from the composition of on- and off-chain information (DLT information combined with device information, linked by common identifiers and data about devices, proofs and documents)
- As a result, the system is *GDPR agnostic* (compliant) as no personal data is stored in the system. Privacy is preserved by the use of one-way functions like hashing, signatures, and the correspondence of references to identifiers in the DLT space to data in DeviceHub space. Therefore details are not reversible or derivable from DLT data.

[15] In the future, with DPP already provided by manufacturers, a lookup will be required to discover any pre-defined DPP for that device. That DPP can include useful details provided by the manufacturer, such as access to information the manufacturer needs to provide (documents, compliance, etc.) Extracting the data details programmatically from the device adds to quality, reliability, and veracity.

4.3 Applicable standards

Apart from the underlying standards of our computing and communications infrastructure, we rely on the following domain-specific standards:

- Smart contracts, written in the de-facto standard Solidity language [16]
- The API follows a de-facto Representational state transfer (REST) software architectural style for Web APIs, where resource's URI elicit a response with a payload formatted in HTML or JSON.
- For reputation accounting, we rely on an ERC20 token contract.
- The architecture of our system builds on the W3C Decentralized Identifiers architecture that includes a naming scheme for resources and the provision of a verifiable registry.
- The generation of bi-univocal device identifiers requires agreements on choices and data normalisation that a community like Obada can provide and aims to standardise.

Our work contributes to the definition of Digital Product Passports for computer devices, as part of the work in ITU, under work item L.GDSPP "Requirements for a global digital sustainable product passport to achieve a circular economy". [23]

4.4 Prototype information

The demonstration consists of a set of software tools: (workbench) that captures data about devices, sent to a device inventory web application (DeviceHub). The base software was developed before this project. In this project, we have extended DeviceHub to integrate calls to new designs and developments, including a verifiable registry with a REST API and a DLT backend running smart contracts also implemented in this project.

The main functionality of our MVP allows us to register devices; register and verify proofs; generate and verify DPPs; generate and verify verifiable reports; All these combine details and verifiability information.[17]

Link to software repositories:

Prior software development (before this research):

- `https://github.com/eReuse/workbench`
- `https://github.com/eReuse/devicehub-teal`

Software development in this research:

- DLT service: `https://gitlab.com/dsg-upc/autoblockchain`
- Smart contracts and API: `https://gitlab.com/dsg-upc/trublo_contracts_api`

5 Experimental evaluation and validation

The experimental validation of the system involves the use of device DID in an inventory system (DeviceHub.org in our case), integrating calls to the verifiable registry API, as well as a proof of concept implementation of a DID resolution (did.ereuse.org/DID) via HTTP redirection to find out the inventory service instance that holds the corresponding DID document. DeviceHub was extended to show CHID or DPP ID details in either JSON or HTTP formats according to

[16] `https://docs.soliditylang.org/`

[17] The proofs contain identifiers, links and summaries (hashes), not detailed information.

HTTP content selection parameters and access credentials (all, less or no detail for the owner, registered or anonymous visitors).

Once the device inventory application has been integrated with one among several verifiable registry service providers. We have performed experiments to validate the functional results (tests to verify meeting functional requirements), performance results, and sensitivity analysis to assess scalability.

5.1 Test plan

We describe the strategy, objectives, infrastructure, and scripts that have been used to test our implementation of the DLT and smart contracts in this project and extract some conclusions through the test results.

Objectives and scope. When testing our implementation, we define three main objectives:

1. Ensure that all our DLT and smart contract functionalities work as expected (e.g., storing a new device, issuing a new DPP).
2. Ensure that our implemented API communicates properly with the DLT and its smart contracts, calling the expected smart contract methods and returning the expected HTTP response status codes.
3. Know and compare the scalability and performance of our smart contracts in different environments (e.g. besu and geth clients, HTTP or WS communication) to find the optimal setup and see if it's suitable for our use case.

Infrastructure. The tests have run onto different implementations of a private Ethereum based DLT, composed of five Linux machines (Ubuntu), each running a Docker container of an Ethereum client (node). In addition, an Alastria node has been used. Every machine has the minimum requirements to run an Ethereum client as stated by the Ethereum foundation: CPU with +2 cores; +4 GB RAM, +8Mbps traffic capacity.

When executing the script in our DLT, all the returned API HTTP response codes were the ones expected. Therefore, we can confidently say that our API communicates successfully with the DLT, and the smart contracts methods execute the functionalities as expected, verifying and accomplishing our main objectives for this test. A file with the script output is available at our repository[18].

5.2 Performance and scalability test

We wanted to test different aspects of different network configurations for our experiments to find out the optimal setup and suitable for our use case.

We compare two Ethereum node implementations: Go Ethereum and Hyperledger Besu. We have used two methods of closing blockchain blocks: periodically timed and on-demand. We have used two protocols when sending transactions to a node: HTTP and WebSocket (WS). Finally, we have evaluated two different evaluation strategies for sending bursts of transactions to the network.

We have used the following two private network configurations:

- Network 1:
  - Running the Go Ethereum client (geth).

[18] `https://gitlab.com/dsg-upc/trublo_contracts_api`

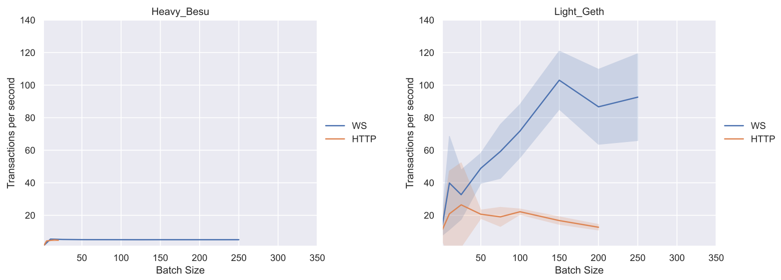

Figure 3 TPS vs batch size for heavy and light transactions in Besu with WS or HTTP.

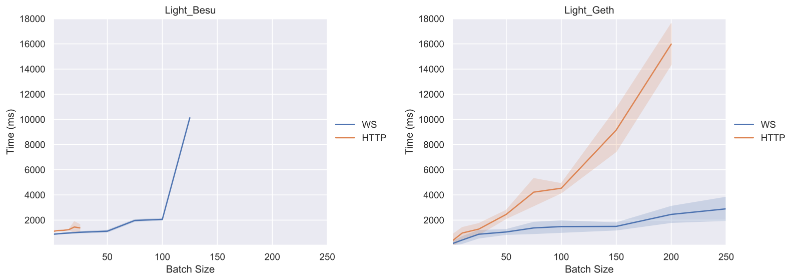

Figure 4 Completion time vs batch size for light transactions in Besu or Geth.

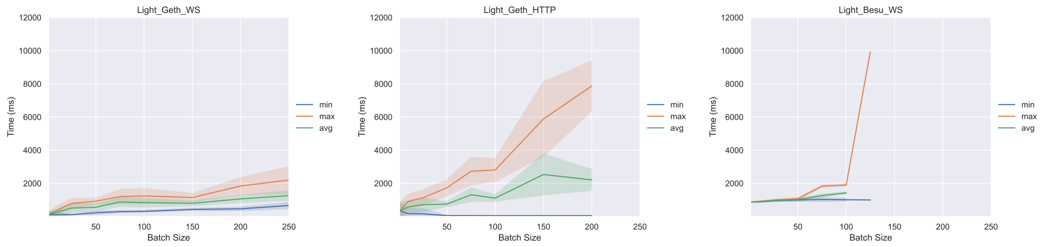

Figure 5 Completion time vs batch size for light transactions in Geth HTTP, Geth WS and Besu WS.

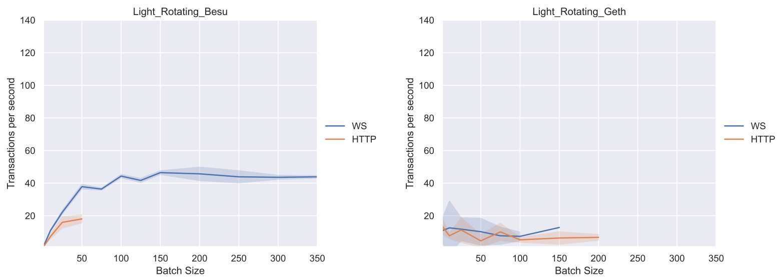

Figure 6 TPS vs batch size for light rotating transactions in Besu and Geth with WS and HTTP.

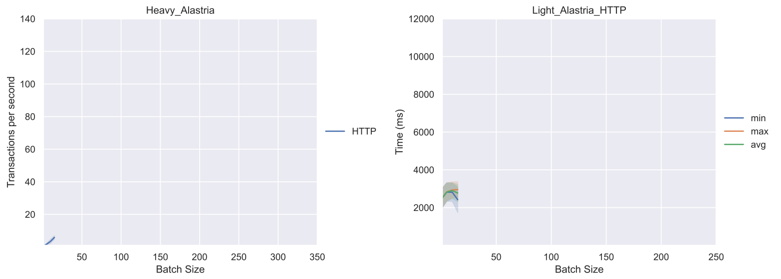

Figure 7 TPS and completion time vs batch size for light transactions sent to Alastria.

  - Five nodes with three validators and two non-validators running a Proof of Authority (PoA) consensus protocol.
  - Blocks closed under demand. That means that blocks can only be closed when transactions are available.

- Network 2:

  - Running the Hyperledger Besu client (besu).
  - Five nodes with three validators and two non-validators running a Proof of Authority (PoA) consensus protocol.
  - Blocks closed periodically every second.

Given the limit on the number of available nodes we can set up in our test environment, we have also tried to perform every test on the Alastria network T. This network runs the Consensys Quorum client and generates blocks at the same rate as our besu network (1 second).

As for the transactions sent, we have tried two different types. First, a “light” transaction consists of creating a proof on one device. This type of transaction requires around 200 thousand gas units to be processed. Second, a “heavy” transaction, corresponding to the registration of a new device, which publishes a new smart contract to the chain. This type of transaction requires around 2.3 million units of gas, which is about 11.5 times more than the light one does. Most of the results will be presented with the light transactions test, as the only difference in the behaviour of the heavy ones is almost always predictable. The only difference is that a block can include fewer of these transactions.

5.2.1 Results

First, comparing the HTTP and WS protocols in Figure 3, we found that the WS interface allows for higher bursts of transactions than the HTTP one. In extreme cases, the HTTP interface stops accepting new requests after about 25 transactions, while the WS one can reach over 200.

On top of that, it is much faster when the right conditions are met, as seen in the second plot, reaching around 5 times more transactions per second processed.

Second, we compare in Figure 4 both clients, geth and besu, by the quality of their RPC interfaces. We found out that the geth interfaces seem to accept a higher number of transactions under the same kind of test.

In particular, the HTTP interface seems to give much better results.

Third, we can compare the methods of closing blocks by seeing how high the latency of each transaction is. The latency is measured as the difference of time between the moment when the transaction is sent by the client and the moment when confirmation of its completion is received. The besu network should have higher latencies as blocks are not closed under demand.

We can clearly see in Figure 5 those minimum latencies are close to the block time in the besu network as predicted. However, across these three plots and earlier ones, a higher variability between samples can be seen in the results of the geth network, which closes blocks under demand.

Fourth, we compare two different strategies for sending transactions to the system in Figure 6. Until now, every plot has been generated by sending the batch of transactions to a single node. In the following plots, we divided every batch. We sent an equal number of transactions to every node in the network, hoping that more machines would increase the number of transactions admitted simultaneously as more individual interfaces would be available.

This worked for the besu network, reaching 350 admitted transactions in a burst, but not for the geth network. This is because closing blocks under demand requires a higher degree of synchronisation across nodes, so the more nodes that receive transactions, the more messages nodes generate to communicate with each other during the on-demand closing of blocks.

Finally, for Alastria, we wanted to run the same kind of tests as in our private networks but found some inconveniences in our way. First, they only expose an HTTP interface by default, so that we couldn't run any WS test. Also, this HTTP interface is exposed through a reverse proxy, giving the following results in Figure 7.

Only around 15 transactions were admitted, and latencies seem worse than on the private networks. We also could not divide the batches to send to every node, as we only have access to a couple of them.

6 Discussion

By analysing the former results, we can develop several takeaways. First, WS seems to provide better results when managing many transactions. We see a higher rate of transactions per second processed and a higher amount of admitted transactions simultaneously. However, the WS protocol needs to keep a TCP connection open, not desirable or valuable when sending a single or few transactions. The HTTP protocol seems more lightweight in this case, reflected by the lower minimum latencies measured in the plots. As both interfaces can coexist and be offered simultaneously by the same node, this shouldn't pose a problem, and the client should decide which interface is the most suitable in each case.

Second, the Go Ethereum client seems to be better regarding the stability of its interfaces compared to the Hyperledger Besu one. It admits more concurrent transactions, and we have not seen crashes when performing the tests, while we have crashed the besu WS interface several times. On top of that, as far as we know, besu doesn't offer the capability to close blocks under demand when using a PoA consensus mechanism.

Third, we believe that closing blocks periodically instead of under demand should work better in networks of moderate activity despite noticing higher latencies. We have noticed that transactions don't propagate to every node in the network and can cause deadlocks under the PoA consensus mechanism. In Ethereum, transactions from a single user have an incrementing nonce. Thus, transaction "n" has to be processed before transaction "n+1", and when a node does not have at its disposal that "n" transaction, it cannot process any of the other ones that it may have at its disposal. In addition to that, it is his turn (in the PoA consensus algorithm) to close the block; it may not have available transactions to process, making the network unable to process transactions under demand and staying in a deadlock until a processable transaction is received. That makes us also believe that the PoA algorithm should not be used, as it is too simplistic and would not work well on a large production network. Furthermore, as seen in the fourth test, sending transactions to different nodes decreases the network's performance under this mechanism, while it does not affect periodically timed one.

As for the period the closing of blocks should be set to, we believe that it should be carefully studied by use case. Producing empty blocks frequently uses up disk space in a not helpful way and results in unneeded synchronisation traffic that is not needed. It should be set depending on the expected activity of the network. However, we have found that setting it too low may make the network struggle to keep up. We have experienced this problem in a 1 second network, where having more than 3 nodes increased the block time by at least 30% of the set time. Further increasing the number of nodes may make this worse. Nonetheless, the Alastria T network keeps up with a 1 second block time while being of decent size (around 200 nodes).

Following up on Alastria, our results are not suitable for any conclusive statement about this network. We can simply say that the number of concurrent transactions that can be sent to a single node is very low under the default configuration. This may improve by not using the reverse proxy and activating the WS interface, as the Consensys Quorum client is based on the Go Ethereum one.

These results are consistent with previous performance analysis done in the NGI Atlantic Open call 2 experiments in 2021 for a deposit and credit system under a similar experimental setup.

7 Conclusions

We have explored the digital transformation of the circular economy of digital devices to bring more efficiency in the management of product-related information and improve transparency and accountability. We have defined device identifiers referencing detailed off-chain document and metadata information and on-chain verifiability information. We distinguish between chassis identifiers and detailed configurations. That leads to the definition and management of digital product passports for individual digital devices. We have tools where devices themselves produce a machine self-generated metadata twin sent to an inventory service that uses the API of a verifiable data registry. This API has access to several DLT backend complemented with ethereum smart contracts that record and keep this data. Our DLTs are an ethereum PoA network with geth, a similar one with besu and a node part of the Alastria T network. We have designed and implemented a prototype system and evaluated it for correctness and performance.

Future work includes the following directions in no particular order: Complete and evaluate IOTA and COSMOS support for our verifiable registry. Integrate partially developed code for an economic model to support transaction fees, reputation credit, and conditional rewards as a deposit to promote the reporting of usage and return of devices for accountable reuse and recycling. Test automation techniques and testing integration into a CI pipeline to improve software quality and early error detection. Pilot with diverse real-world stakeholders to test the system in a realistic environment. Participate in dissemination of the ideas and results globally with the support of ITU, contribute to EU level concertation of DPP specifications for digital devices, expand the work on DPP in the ITU and related standardisation bodies. Finally, maintain and extend the open-source code and publish related datasets.

References

1 John Adler, Ryan Berryhill, Andreas Veneris, Zissis Poulos, Neil Veira, and Anastasia Kastania. Astraea: A decentralized blockchain oracle. In *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, pages 1145–1152, 2018. doi:10.1109/Cybermatics_2018.2018.00207.

2 Anders S.G. Andrae. Life-cycle assessment of consumer electronics: A review of methodological approaches. *IEEE Consumer Electronics Magazine*, 5(1):51–60, 2016. doi:10.1109/MCE.2015.2484639.

3 European Commission. A European Green Deal. Technical report, European Commission, 2019. URL: https://ec.europa.eu/info/strategy/priorities-2019-2024/european-green-deal_en.

4 T. Dahlberg, J. Hallikas, and K. Korpela. Digital supply chain transformation toward blockchain integration. *Proceedings Of The 50Th Hawaii International Conference On System Sciences (2017) - Proceedings Of The Annual Hawaii International Conference On System Sciences*, 2017.

5 Emmanouil Dimogerontakis, Leandro Navarro, Mennan Selimi, Sergio Mosquera, and Felix Freitag. Contract networking for crowdsourced connectivity. In *2020 IEEE International*

Conference on Decentralized Applications and Infrastructures (DAPPS), pages 126–132, 2020. doi:10.1109/DAPPS49028.2020.00016.

6 Ellen McArthur Foundation. What is the circular economy? Technical report, Ellen McArthur Foundation, 2021. URL: https://www.ellenmacarthurfoundation.org/circular-economy/what-is-the-circular-economy.

7 eReuse initiative. eReuse software tools: workbench, deviceHub, etc. Technical report, Ereuse project, 2022. URL: https://www.ereuse.org/software/.

8 European Commission. Circular Economy Action Plan. Technical report, European Commission, 2020. URL: https://ec.europa.eu/environment/strategy/circular-economy-action-plan_en.

9 European Commission. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on batteries and waste batteries, repealing Directive 2006/66/EC and amending Regulation (EU) 2019/1020, of 10.12.2020. Technical report, European Commission, 2020. URL: https://op.europa.eu/en/publication-detail/-/publication/4b5d88a6-3ad8-11eb-b27b-01aa75ed71a1.

10 European Commission. Green digital passport: Product traceability to create circular material loops. Technical report, European Commission, 2021. URL: https://circulareconomy.europa.eu/platform/en/news-and-events/all-events/green-digital-passport-product-traceability-create-circular-material-loops.

11 European Commission. Sustainable Products Initiative products regulation proposal. Technical report, European Commission, 2022. URL: https://ec.europa.eu/environment/publications/proposal-ecodesign-sustainable-products-regulation_en.

12 Obada Foundation. The open blockchain for asset disposition architecture, 2021. Foundation. URL: https://www.obada.io.

13 David Franquesa, Leandro Navarro, David López, Xavier Bustamante, and Santiago Lamora. Breaking barriers on reuse of digital devices ensuring final recycling. In *29th International Conference on Environmental Informatics, EnviroInfo 2015, Copenhagen, Denmark*, pages 281–288, Amsterdam, Netherlands, 2015. Atlantis Press.

14 David Franquesa, Mireia Roura, and Leandro Navarro. ereuse datasets, 2013-10-08–2019-06-03, 2020. URL: https://dsg.ac.upc.edu/ereuse-dataset.

15 Susanne Guth-Orlowski. The digital product passport and its technical implementation, 2021. URL: https://medium.com/@susi.guth/the-digital-product-passport-and-its-technical-implementation-efdd09a4ed75.

16 ITU. Recommendation ITU-L.1410, Methodology for environmental life cycle assessments of information and communication technology goods, networks and services. Draft recommendation, International Telecommunication Union, 2015. URL: https://www.itu.int/rec/T-REC-L.1410.

17 ITU. Recommendation ITU-L.1470, GHG emissions trajectories for the ICT sector compatible with the UNFCCC Paris Agreement. Draft recommendation, International Telecommunication Union, 2020. URL: https://www.itu.int/rec/T-REC-L.1470.

18 ITU-T. The potential impact of selling services instead of equipment on waste creation and the environment - effects on global information and communication technology, 2021. Recommendation ITU-T L.1024. URL: https://www.itu.int/rec/T-REC-L.1024.

19 Mahtab Kouhizadeh, Qingyun Zhu, and Joseph Sarkis. Blockchain and the circular economy: potential tensions and critical reflections from practice. *Production Planning & Control*, 31(11-12):950–966, 2020. arXiv:https://doi.org/10.1080/09537287.2019.1695925, doi:10.1080/09537287.2019.1695925.

20 Luxembourg government. Product circularity data sheet, 2020. Lu. URL: https://pcds.lu.

21 Luxembourg government. Product circularity data sheet: questions and answers, 2020. Lu. URL: https://pcds.lu/wp-content/uploads/2020/11/20201218_WebinarPCDS_QA.pdf.

22 Rumy Narayan and Annika Tidström. Tokenizing coopetition in a blockchain for a transition to circular economy. *Journal of Cleaner Production*, 263:121437, 2020.

URL: https://www.sciencedirect.com/science/article/pii/S0959652620314840, doi:https://doi.org/10.1016/j.jclepro.2020.121437.

23 Leandro Navarro. Requirements for a global digital sustainable product passport to achieve a circular economy. Draft recommendation, ITU-T Q7/5, 2021. URL: https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=17130.

24 Statista. Statista, smartphones industry: statistics and facts, 2021. URL: https://www.statista.com/statistics/263437/global-smartphone-sales-to-end-users-since-2007/.

25 United Nations. Safety data sheets, 2007. URL: https://unece.org/DAM/trans/danger/publi/ghs/ghs_rev02/English/08e_annex4.pdf.

26 Arvind Upadhyay, Sumona Mukhuty, Vikas Kumar, and Yigit Kazancoglu. Blockchain technology and the circular economy: Implications for sustainability and social responsibility. *Journal of Cleaner Production*, 293:126130, 2021. URL: https://www.sciencedirect.com/science/article/pii/S0959652621003504, doi:https://doi.org/10.1016/j.jclepro.2021.126130.

27 Web Consortium. Use cases and requirements for decentralized identifiers, 2021. W3C Working Group Note. URL: https://www.w3.org/TR/did-use-cases/.

28 Web Consortium. Safety data sheets, 2022. URL: https://www.w3.org/2019/did-wg/.

29 Zhang Xinxing, Tian Zhihong, and Zhang Luchen. A measurement study on mainline dht and magnet link. In *2016 IEEE First International Conference on Data Science in Cyberspace (DSC)*, pages 11–19, 2016. doi:10.1109/DSC.2016.106.