Foundation of Web3: Fundamental Building Blocks, Architectural Principles, Design Space

Ahto Buldas[1,4][0000−0002−6413−8258], Dirk Draheim[2]✉[0000−0003−3376−7489], Mike Gault[3][0000−0002−4860−7077], and Märt Saarepera[4][0000−0003−1808−9915]

[1] Centre for Digital Forensics and Cyber Security, Tallinn University of Technology, Akadeemia tee 15a, 12618 Tallinn, Estonia ahto.buldas@taltech.ee

[2] Information Systems Group, Tallinn University of Technology, Akadeemia tee 15a, 12618 Tallinn, Estonia dirk.draheim@taltech.ee

[3] Guardtime, Avenue d'Ouchy 4, 1006 Lausanne, Switzerland mike.gault@guardtime.com

[4] Guardtime, A. H. Tammsaare tee 60, 11316 Tallinn, Estonia {ahto.buldas,mart.saarepera}@guardtime.com

Abstract. The Web3 vision takes blockchain disintermediation to a next level by making it ubiquitous, encompassing not only payments and financial services but also digital identities, data and business models. Recently, Web3 has gained massive attention by major analysts such as Gartner, Forrester, Forbes Technology Council and the Harvard Business Review. Albeit the current enthusiasm about Web3, we are lost in a state of confusion about what Web3 actually is – or could be. The purpose of this paper is to mitigate the gap between the perceived usefulness of Web3 and its potential implementation. We take a descriptive design science approach. We provide informed arguments for a potential foundation of Web3 in terms of fundamental components, architectural principles and a Web3 design space. We demonstrate the usefulness of the provided Web3 foundation by describing Alphabill, a platform that allows for universal asset tokenization and transfer as a global medium of exchange. The findings of this research enable policy makers, decision makers and information systems architects alike to make informed decisions about Web3 and its potential implementation as follows: (i) The Web3 can be characterized as the integration of digital rights exchange into the (application layer) internet protocols. (ii) The Web3 has the potential to revolutionize today's information systems's landscape by turning today's information systems into deeply standardized views on a huge, single underlying information structure. The killer application of a well-founded Web3 is the Web3 itself – being the currently missing backbone (value-added middleware) for all of today's and future enterprise applications and business-to-business communication. (iii) The scenario-based evaluation of the provided Web3 foundation reveals the described Alphabill platform as a Web3 enabling technology.

Keywords: Web3 · Blockchain · decentralized finance · DeFi · Alphabill

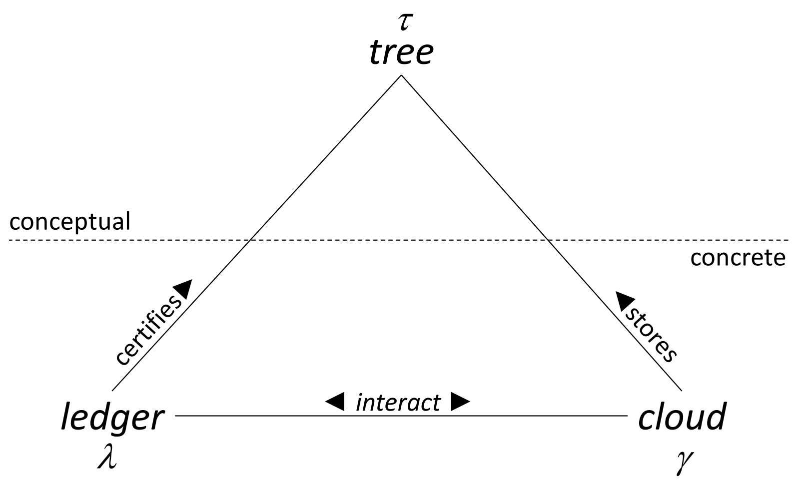

Fig. 1. Web3 fundamental components: tree, ledger and cloud.

1 Introduction

The Web3 vision takes blockchain disintermediation to a next level by making it ubiquitous, encompassing not only payments and financial services but also digital identities, data and business models. Where the the vision of a ubiquitous integration of emerging technology has become widely known as Internet of Things (IoT), the Web3 narrative can be characterized as the *Web of Everything*, and even more, the *Web of Everything and Everybody*, since the idea of being "owned and operated by its users" [1] is the key ingredient of Web3. Although Web3 is still in its infancy, it has gained massive attention by major analysts such as Gartner [2], Forrester [3] and Forbes Technology Council [4] as well as the Harvard Business Review [5,1,6], and the expectations are high towards Web3 being "our chance to make a better internet" [1]. In Table 1, we have summarized a series of Web3 characteristics that we find most significant for the current Web3 narrative[5] by comparing them to corresponding Web 2.0 characteristics.

Albeit the current enthusiasm about Web3, we are lost in a state of confusion. The Forbes technology article titles "Why Web3 Is So Confusing". The Forrester article is titled "Web3 Isn't Going To Fix The Shortcomings Of Today's Web". And when both the Gartner article [2] and the Harvard Business review article [5] are titled "What Is Web3?", they rather aim at giving an overview of the current Web3 narrative than answering the question. And, actually, since the Web3 does not yet exist, the question of 'what *is* Web3', can only be about overviewing its current narrative. Therefore, for an engineering research endeavor, the question to be asked has to be 'what *will be* Web3'?

[5] not to be confused with Web 3.0 [7] (related to the Semantic Web [8]).

In this paper, we discuss a potential foundation of Web3 in terms of fundamental components, architectural principles and a Web3 design space. We postulate that any implementation of Web3 can be explained in terms of three fundamental components, i.e., tree, ledger and cloud (Fig. 1) that adhere to a series of Web3 architectural principles and thus form the basis to elaborate the full Web3 design space. We find that the killer application [9] of a well-founded Web3 is the Web3 itself – being the currently missing backbone (value-added middleware) for all of today's and future enterprise applications and business-to-business communication.

In [10], we have contributed the architecture of the Alphabill platform – a platform for universal asset tokenization, transfer and exchange as a global medium of exchange. In this paper, we discuss Alphabill as an enabler for Web3 – in terms of the suggested Web3 foundation.

We proceed as follows. In Sect. 2, we provide a discussion of related work. In Sect. 3, we provide an analysis of Web3 characteristics as perceived in the recent discourse. In Sect. 4, we provide a description of the suggested Web3 foundation. In Sect. 5, we delve into the Web3 design space. In Sect. 6, we discuss the Alphabill platform. We finish with a conclusion in Sect. 7.

2 Related Work

2.1 Socio-Econmic Aspects of Web3

Sadowski and Kaitlin investigate the movement around Web3 [11]. The study investigates how the Web3 movement formed as a result of discourse about notions of decentralization, political economic interests and venture capital endeavors. Furthermore, Sadowski and Kaitlin provide an analysis of obstacles and barriers to successful Web3 development [11].

Qin et al. analyze the concept of Decentralized Autonomous Organization (DAO) [12,13,14,15] from the "perspective of organization and operation" [16] aiming at "providing a more precise definition of DAOs" [16] encompassing a discussion of many aspects ranging from cyber-physical-social systems, over parallel intelligence, digital twins, metaverse, to Web3.

Murray et al. analyze the promises of Web3 and aim at understanding how can firms prepare for Web3. Murray et al. identify fungible tokens (cryptocurrencies), nonfungible tokens (NFTs), decentralized autonomous organizations (DAOs), and metaverses as central blockchain-enabled Web3 applications. [17] conclude that, in order to be prepared for Web3, companies should "reconsider the role of network effects" [17] in the sense of [18], "tokenize intellectual contributions" [17], "reconsider organization hierarchy (to a degree)" [17], and "double down on [Web3] stakeholders and communities" [17], however, also to "exercise caution when investigating in Web3" [17].

Kelsie Nabben is able to "demonstrate how the origins of Web3 reveal the intentions of its creators as a political tool of prefiguration, yet its practices reveal the inherent tension of expressing these ideals in coherent technical and institutional infrastructure" [19]. On the basis of this, she argues "that one of the fundamental challenges Web3 is negotiating through technical and governance experiments is »how to self-infrastructure?«" [19].

2.2 Central Bank Digital Currencies

Recently, central bank digital currency (CBDC) gains more attention by policy makers all around the globe. The Hamilton Project [20] is a concept study conducted by the Federal Reserve Bank of Boston and Massachusetts Institute of Technology Digital Currency Initiative in order to evaluate the feasibility of two different approaches of CDBC implementations. The project investigates a sharded, UTXO-based blockchain solution (so-called *atomizer architecture*) and a *2PC architecture* (2-phase-commit architecture), which utilizes standard transaction system technologies found in today's banking. RSCoin [21] is a sharded, blockchain-based CBDC. RSCoin's architecture is based on shards and a trusted central component. The shards create consensus on the next valid blocks, which is then certified by the central component. The shards communicate indirectly via the wallets adhering to a two-phase-commit protocol. The Alphabill platform is different from both the Hamilton project and RSCoin, as they are designed as platforms for universal asset tokenization. KSI Cash is different from both the Hamilton project's atomizer architecture and RSCoin, because KSI Cash achieves ultra-scalability in a fundamentally different way. Both RSCoin and the Hamilton atomizer architecture are UTXO-based, whereas KSI Cash utilzies the bill money scheme. RSCoin and the Hamilton project need to deal cross-shard transactions, which is a severe issue. Instead, KSI Cash deals with the comparatively moderate issue of increasing amounts of smaller money denominations [10].

2.3 Federated Blockchain Technology

Polkadot[6] is a "heterogeneous multi-chain" [22]. Polkadot provides a so-called *relay chain*, which may host "parellizable" chains – so-called parachains. With that concept, Polkadot aims at blockchain scalability ("In principle, a problem to be deployed on Polkadot may be substantially parallelised – scaled out – over a large number of parachains." [22]) and blockchain heterogeneity ("In other words, Polkadot may be considered equivalent to a set of independent chains (e.g. the set containing Ethereum, Ethereum Classic, Namecoin and Bitcoin) ..." [22]). Parachains can dispatch transactions in other parachains or the relay chain. The dispatched transactions are fully asynchronous, i.e., "there is no intrinsic ability for them to return any kind of information back to its origin" [22]. The interchain communication is based on queuing: "Interchain transactions are resolved using a simple queuing mechanism based around a Merkle tree to ensure fidelity. It is the task of the relay-chain maintainers to move transactions on the output queue of one parachain into the input queue of the destination parachain." [22]. There are

[6] `https://polkadot.network/`

two key differences between Polkadot and Alphabill. The first difference is in their approach to decomposition [23,24,25], i.e., Polkadot is a *federation* of multiple blockchains, whereas Alphabill is a single partitioned blockchain. Second, they differ in their communication models. Polkadot chooses asynchronous messaging, whereas Alphabill establishes rigorous atomicity control for each cross-partition swap [10].

Lukaj et al. utilize a federated blockchain [26] to improve IoT (Internet of Things) service provision by decoupling "the communication between IoT devices and the CA [(certification authority)]" [27], i.e., federated blockchains "are adopted to guarantee a distributed storage for information from different certified data sources" [27].

Nguyen et al. have designed a framework [28] for federated-blockchain systems that utilizes sidechain technology for cross-chain transfer of assets; i.e., the framework's cross-chain transfer protocol builds on *simplified payment verification* (SPV) [29]. Furthermore, as part of the framework, Nguyen et al. suggest a proof-of-stake consensus mechanism together with an incentive mechanism using a Stackelberg game model [30].

Ma et al. have designed a decentralized authentication protocol for self-sovereign identity (SSI) that allows for "fine-grained access control of identity information" [31]. In the protocol, "distributed identifiers are used as global identifiers in decentralized domains based on federated chains" [31] along the lines of the policy-controlled signature scheme suggested in [32].

Mohey Eldin et al. suggest to utilize an *institutionally* federated blockchain system for the healthcare domain that they have designed on the basis of a previously conducted survey [33]. Technologically, the suggested system architecture consists of three module blockchains (called "shards" in [34]); i.e., a top-level "authority shard" (using proof of stake (PoS)); a permissioned "cache" shard that "acts as a temporary storage" [34], for which Mohey Eldin et al. recommend proof of authority (PoA); and a "master" shard, which is a "consortium blockchain that contains the mid-level healthcare organizations and is being federated by the high level healthcare organizations" [34] using proof of authority (PoA) to increase scalability.

2.4 Asset Tokenization

Ethereum is widely used to create new tokens. In an ICO (initial coin offering), the Ethereum development standard ERC-20[7] can be utilized to create new tokens. These new tokens represent the coins of a newly offered cryptocurrency. However, typically this new cryptocurrency is eventually build as an entirely new, separate blockchain. This blockchain would implement the innovations outlined in the white paper of the respective ICO. These innovations can be about processes in a business domain or about technological advancements such as performance, security, decentralization, and integration of emerging technologies. The new cryptocurrency now exists in the form of two kinds of tokens, i.e.,

[7] https://ethereum.org/en/developers/docs/standards/tokens/erc-20/

its ERC-20-based tokens in Ethereum and the genuine coins of the separately implemented blockchain. Both coins can be traded, and investors can exchange the ERC-20-based tokens to the genuine token via a blockchain bridge[8]. In such an ICO lifecycle, Ethereum has merely the role of a crowdfunding platform, but the new token is not genuinely "build on top of Ethereum" [35]. In principle, it is possible to build the new cryptocurrency entirely as an Ethereum token, i.e., without separate implementation and without blockchain bridging. In such case, all innovations of the new cryptocurrency would be reflected in Ethereum smart contracts. These innovations can be business domain innovations but no blockchain technology innovations, and the new cryptocurrency is limited by the technological capabilities of the hosting Ethereum platform. In contrast, Alphabill is designed as a platform for universal asset tokenization from scratch [10].

3 An Analysis of Perceived Web3 Characteristics

3.1 A Brief History of the Web

A standard explanation of the history of the Web – as typicaly found in text books, lectures, blogs etc. – is provided in terms of Web 1.0, Web 2.0 and Web 3.0 (see dashed line "Tim Berners-Lee" in Fig. 2). Here, Web 1.0 stands for a descriped initial situation of the Web, where Web users are divided into content creators and content consumers. In Web 1.0, the mass of Web users consume content created by a relatively small number content creators. Then, Web 2.0 stands for moving of ordinary web users from content consumption to content creation, i.e., getting *everybody* involved in Web content creation. Then, the Web 3.0 vision, called *semantic web*[9] [8] was (it actually never took off) about systematic meta data and a Web of linked data. Practically, and not reflected in this standard story, the Web is about e-commerce, which started in 1996 with the availability of Netscape's SSL 3.0. This means that the whole dotcom bubble happened in less than half of a decade – the second half of the "roaring nineties" [36].

Web 2.0 has not developed as envisioned. Bigtech has taken over with huge social media platforms dominating the scene, leading to a massive commercialization of customer data and a further divide between consumers and producers. Tim-Berners Lee, being dissatisfied with how the Web developed [37], initiated the Solid [38] project[10] to advance the Web for a better. For similar reasons, to be analyzed in more depth in Sect. 3.2, Gavin Wood coined the term Web3 (not to be confused with Web 3.0) and the Web3 movement formed from inside the Ethereum community as a dedicated blockchain-based vision of a next generation Web.

[8] https://ethereum.org/en/bridges/
[9] https://www.w3.org/standards/semanticweb/
[10] https://solidproject.org/

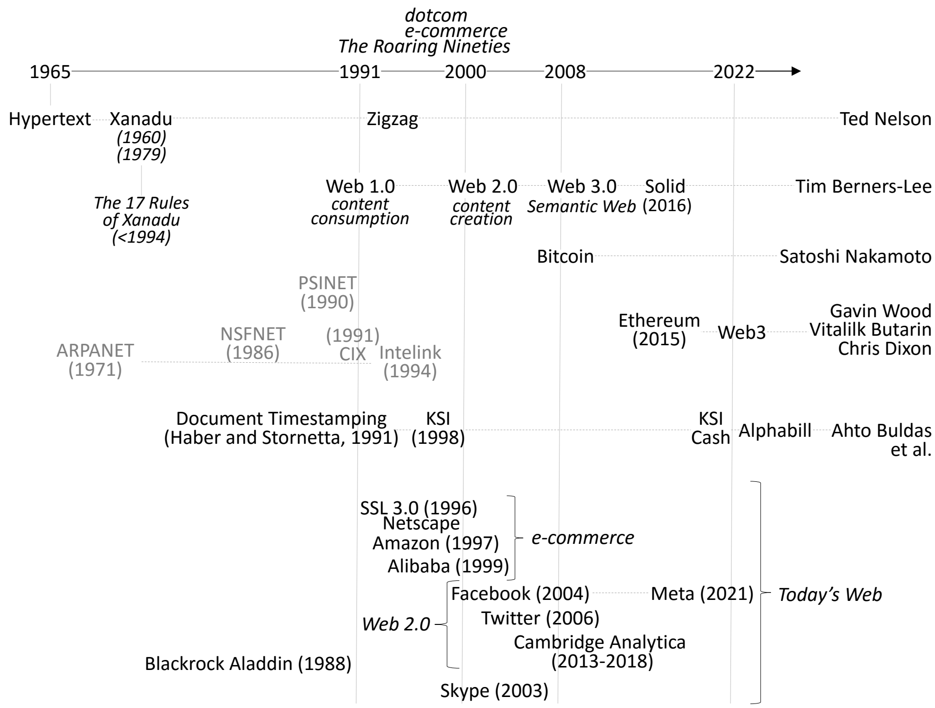

Fig. 2. A brief history of the Web.

3.2 A Comparision of Perceived Today's Web and Web3 Characteristics

Gilad Edelman has summarized the intention of the Web3 movement follows: "The Web3 movement seeks to liberate us from Big Tech and exploitative capitalism – and to do it using only blockchain, game theory, and code." [40] . Throughout Sections 3.2 to 3.2, we aim at further analyzing the characteristics of the Web3 vision in terms of several categories (payments, financial services, identity concepts, data ownership, trust anchors, protocal characteristics, business models, uses cases), as perceived in the current discourse, see Table 1 for a summary.

Payment Characteristics The success of today's Web is in large part due to the widespread adoption of e-commerce. Payments in today's Web are based on online bank transfers. Online bank transfers are realized between accounts hosted by commercial banks, i.e., they are realized in the tier of M1-money. This is also true for "digital payments in existing currencies – through Paypal and other »e-money« providers such as Alipay in China, or M-Pesa in Kenya" [41]. Instead, the Web3 relies on cryptocurrencies. Payments are realized as direct payments

	Today's Web	**Web3**
Payments	Online bank transfers.	Cryptocurrencies; payments without intermediaries
Financial services	Not considered as critical.	Built-in DeFi; considered as integral part (disrupting established banking).
Identity concepts	Public key infrastructure; cloud-based e-identity; ISO 29003.	Self-sovereign identity.
Trust anchors	Authorities; companies.	Peer-to-peer; consensus protocols.
Data ownership	Data owned and utilized by companies.	Data owned and utilized by users.
Protocols	Stateless.	Stateful.
Business models	Big Tech; e-commerce; social platforms; commercialization of customer data.	DeFi business models; DAOs.
Use cases	Content consumption, content production; B2C; rather less: B2B.	All use cases of today's Web, disintermediated; disintermediated B2B is considered essential; NFTs.

Table 1. Most significant Web3 characteristics – compared to today's Web. As perceived in the discourse.

between web users, without intermediaries. Cryptocurrencies are neither owned by a central bank nor collateralized M1-money. Cryptocurrencies do not belong to the tiered monetary system. Central bank digital currency [42], which is M0-money, is usually not considered part of Web3.

Financial Services In today's Web, financial services might be made accessible through web-based e-commerce services, however, they are not perceived as playing a critical role in todays's Web. Instead, for Web3, the notion of decentralized financial (DeFi) [43,44,45,46] is crucial. Financial services are considered integral part of Web3, disrupting both established commercial banking and investment banking.

Identity Concepts Today's Web relies on established eID concepts [47,48,49] such as based on public key infrastructure (PKI) and cloud-based e-identity, including established routines of personal identity proofing such as ISO 29003 [50]. Instead, Web3 is centered around the notion of self-sovereign identity (SSI) [51].

Trust Anchors In today's Web, trust anchors are provided by authorities and companies. In Web3, trust is created peer-to-peer [52] by consensus protocols of underlying blockchain technology.

Data Ownership In today's Web, data is perceived to be owned and utilized by companies as part of their Web-based business models, see Sect. 3.2. In contrast, the Web3 vision is about data owned and utilized by the Web users.

Protocol Characteristics In today's Web, protocols can be considered as "stateless", in the sense that protocols connect siloed applications and regulate the "transmission of data, not how data is stored" [53]), as opposed to Web3, which envisions "stateful" protocols in the sense of "collectively maintained universal state for decentralized computing." [53]. Data storage, versioning, access rights, and identity might be subject to such "stateful" protocols. Compare also with the discussion of the *http* protocol conducted in Sect. 4.2.

Business Models Today's Web is perceived as dominated by Silicon Valley tech giants such as Alphabet, Amazon and Meta. Web-based business today is about super-scaling e-commerce. Equally, it is also about social media and social networks via commercialization of customer data. For Web3, genuine DeFi business models are essential, such as decentralized payment services, decentralized fundraising and decentralized contracting [45]. Also, the Web3 movement aims as novel forms of organizations such as the decentralized autonomous organization (DAO) [13].

Use Cases Following the usual narrative, today's Web is about content consumption (Web 1.0) and content production (Web 2.0) includung social media and social networks, see Sect. 3.1. Practically, business-to-customer (B2C) e-commerce is an extremely important use case, see Sect.3.2. This was true for the dotcom-era [36] and is true for the post-dotcom era. An increasingly important use case are are collective intelligence systems [54]. Despite the fact that service-oriented architecture SOA (incl. Web services) [55] is considered as belonging to the Web technology stack, we find that business-to-business (B2B) is rather not considered a typical Web use case today. This might simply be due to the perspective of the average Web user, who is an end consumer. When it comes to Web3, all of today's Web use case, are also Web3 use cases, however, always disintermediated. Furthermore, isintermediated B2B is considered an integral part of Web3. A currently widely discussed Web3 use case are non-fungible tokens (NFTs) [56,57], which "can represent real-world items like artwork and real estate" [58], and "can also ... represent individuals' identities, property rights, and more" [58].

4 A Web3 Foundation

We postulate that any implementation of Web3 can be explained on the basis of three fundamental components [59], i.e., tree, ledger and cloud (abbreviated as $\tau + \lambda + \gamma$), see Fig. 1, that adhere to a series of Web3 architectural principles and thus form the basis to elaborate the full Web3 design space, compare with Table 2. We say that the Web3 fundamental components together with the Web3 architectural principles and the elaborated Web3 design space form the foundation of Web3.

- Web3 Fundamental Components
  - *Tree.* Partitioned information tree. Each partition establishes its own language and rules. Access rights are an integral part of the tree.
  - *Ledger.* Provides certificates for Web3 information. Fully certified complete protocol Web3 log.
  - *Cloud.* Stores the full version history of the web tree.
- Web3 Architectural Principles
  - *Pervasive Digital Rights.* Web3 is about the integration of digital rights exchange into the (application layer) internet protocols.
  - *Data Abstraction Principle.* The Web3 state tree is manipulated and *only* manipulated through Web3 protocols.
  - *Web3 Livestream.* All Web3 protocol activities are recorded in the Web3 ledger.
  - *Maximizing Data Protection.* Data privacy/anonymity are protected maximally against *everybody* (in tension with data transparency required by *regulators*).
  - *Ultra Scalability.* Ledger transaction performance is the *sine-qua-non* precondition for the Web3 to be turned into reality.
- Web3 Design Space
  - *A Better Web*
  - *Amalgamation of Intranet and Internet*
    - *A Massive Enterprise Application Backbone*
    - *A Massive Interorganizational Application Backbone*
    - *A Massive Devops Backbone*
  - *A Web of Everything*
    - *A Web of Manufacturing and Logistics*
    - *SDN, IoT and Blockchain*
  - *A Web of Everybody*
    - *Disintermediation*
    - *Governance*
    - *Collective Intelligence*

Table 2. A Web3 foundation.

4.1 The Web3 Information Tree

The fundamental component τ represents the latest Web3 information as a tree[11]. The tree structure is extended to a graph structure by nodes that contain node addresses and are interpreted as references – as we are used to from hyperlinks [60] and transclusions [60,61]. Conceptually, the Web3 information is actually a graph. We stay in the tradition of modeling web information as a tree. We do not do so for the sake of tradition in its own right. It is the language-oriented stance of the tree approach that is beneficial for us, when we conceptualize an essential ingredient of Web3: each partition of Web3 defines its own domain-specific language. Furthermore, the tree approach allows for a convenient ad-hoc addressing scheme: paths in the tree[12].

– The tree τ is a child-ordered, node-colored tree.

Edges, child-ordering and node colors of τ are used to express information. We call the color of a node its *node information*. Depending on the context, we call the color a node also the *label* of the node (i.e., we use *node color*, *node information* and *node label* as interchangeable). We use the label l of a node v also to identify the sub tree σ that has v as a root, and call l also the *label of sub tree* σ. We understand τ as an abstract syntax tree that adheres to a context-free grammar that we call the *grammar of* τ[13,14]. We call the language that τ belongs to the *Web3 base language*, denoted by W (the grammar of τ is the grammar of W, and $\tau \in W$).

Web3 Partitions The Web3 tree τ is partitioned. A partition is a *sub tree* of τ. The purpose of a partition is to hold the information of a specific asset or a specific domain. The list of example partitions is sheer endless. Basically, all of the cryptocurrency-based platform visions seen during ICOs (initial coin offerings) in the last decade can be realized as partitions in Web3. Examples of partitions could be: a cryptocurrency, a real-estate tokenization platform [64] (ideally connected to the official cadastre; or even being the official cadastre), an e-procurement system for the public sector [65,66], a nation-wide healthcare information system [67], a business-to-business vending platform, a particular relational database of a certain company etc. Each partition owns a partition-specific language (that is used describe the content of the partition) and is governed by partition-specific rules.

Figure 3 illustrates an example Web3 tree τ. The topmost part of τ (in Fig. 3, consisting of labels of the form τx) realizes the addresses of the Web3 partitions.

---

[11] latest= at each point in time. Versioning comes in through the other fundamental components: cloud γ and ledger λ.

[12] compare to X-Paths: `https://www.w3.org/TR/xpath-31/`

[13] Abstract syntax trees are indeed child-ordered, node-colored trees.

[14] As suggested by the *integrated source code paradigm* [62,63], we overcome concrete syntax (we exclude concrete syntax from our considerations) and work with grammar and a direct utilization of its abstract syntax.

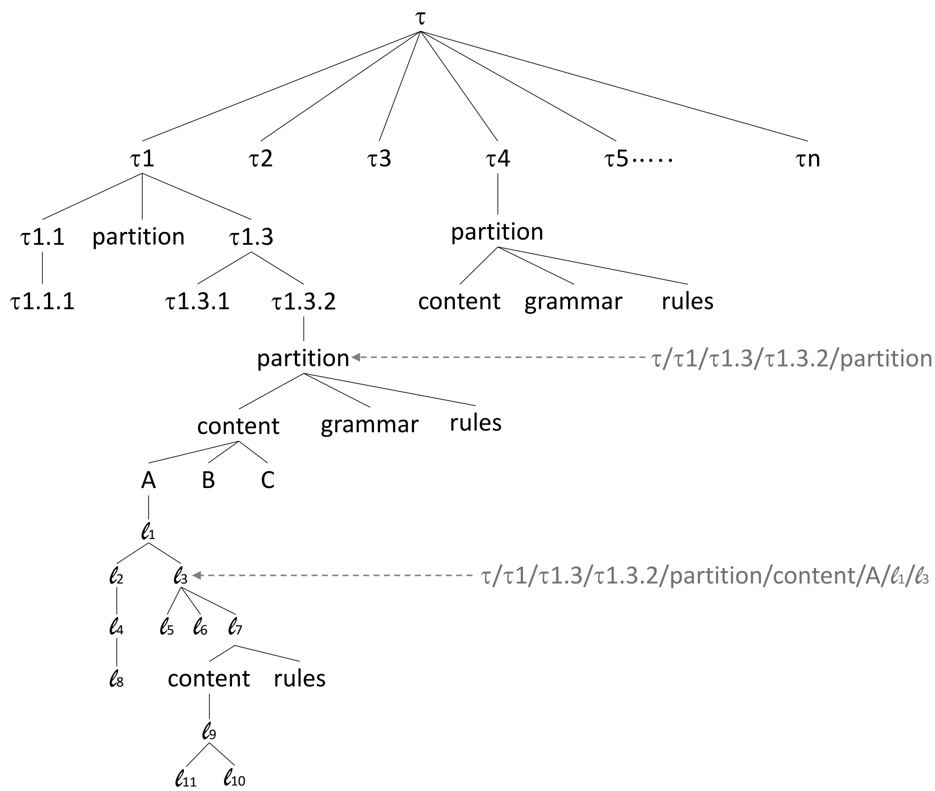

Fig. 3. The Web3 information tree τ.

Figure 3 depicts three partitions having the addresses $\tau1/partition$, $\tau4/partition$, and $\tau1/\tau1.3/\tau1.3.2/partition$. Actually, addresses of partitions are not special, each node in the Web3 tree can be addressed, see the address of the node labeled l_3 in Fig. 3 as an example.

The Web3 base language W contains a language for describing context-free grammars as a sub-language. Each partition has three child nodes: *content*, *grammar*, and *rules* (see Fig. 3). The grammar node contains a the grammar of the partition. All sub trees of the content node has to adhere to this grammar. Furthermore, W contains a programming language as sub language that allows for establishing rules for sub trees of τ. This programming language has the same intentions as Bitcoin's programming language Script and the smart contract [68] languages [69,70] of other blockchain technologies such as Ethereum's Solidity[15]. We call this sub language of W simply the Web3 programming language.

The rules node of a partition contains rules that are written in the Web3 programming language. A rule evaluates to true or false and can have side-effects

[15] `https://docs.soliditylang.org/en/v0.8.16/`

upon its execution. The rules are triggered whenever the content of the partition is about to be changed by a Web3 protocol. Whenever a rule is violated (evaluates to false), the change is rejected. The rules of a partition are used to complete the partition's context-free grammar, i.e., they are used enforce needed context-sensitive properties of the partition. But the Web3 programming language allows for much more. It allows to access the full protocol history that is reflected in the Web3 ledger λ and, therefore, to access the full version history of τ – either directly via λ or indirectly via λ (via hash identifiers provided by λ) by accessing data stored in the cloud γ (γ and λ and options to distribute data over γ and λ will be described in Sect. 4.2).

Web3 Substructures In general, substructures of the Web3 tree recursively follow the same principles of the Web3 tree as expressed in Sect. 4.1. We can then speak of sub-partitions, with sub-partitions potentially defining their own grammar and rules. One option is to built in such recursiveness into the Web3 tree concept, however, we can also leave it to the specification of the grammar and rules of the partition that hosts a sub-partition. The grammar and rules of a partition can specify that there are additional rules under a certain path of the partition's tree that apply to a certain sub tree of the partition. See sub tree l_7 in Fig. 3 for an example. Here, the rules are meant to apply to all sub trees of the corresponding (sibling) content node. However, there is full flexibility in regard of specifying where rules appear and to which sub tree they would apply. Furthermore, the grammar and rules of a partition can even introduce a new language to be used as more specific Web language in this partition or in parts of the partitions. Such language specifications can also be re-used by other partitions of the Web3 tree as part of their grammars and rules. We denote any of those web languages currently specified in the Web3 tree as w as opposed to the Web base language W.

Pervasive Digital Rights In the Web3, digital rights are a pervasive concept. They are so essential for the Web3 vision that the Web3 can be even characterized in terms of them, i.e., as the integration of digital rights exchange into the (application layer) internet protocols. Digital rights express trusted, certified ownership of digital assets. Digital rights manifest in digital signatures of digital assets stored in the Web3 tree τ. Complex digital rights scenarios can be expressed with the Web3 programming language. The enforcement of digital rights is on a different page. Basic digital rights that are merely about consuming (accessing) digital assets might be enforced (ensured) technologically. However, in general, when digital rights are about re-use of digital assets, they need to be collateralized by appropriate regulations. With the Web3, the notion of digital rights itself seems to become generalized. They are not merely about the utilization of digital assets anymore, instead, they express rights in real-world assets (legal assets or physical assets). Again, such notion of digital rights need to collateralized by regulations. In this strand of Web3, regulations and institutions [71,72] need to co-evolve [73,74] with the emerging Web3. The Web3 need to anticipate (conceptually and technologically) such developments.

Access rights represent a basic form of digital rights. Similar to digital rights, access rights are an essential, integral part of the tree. Access right owners are identified via public cryptographic keys. We consider the access rights as part of the rules. Access rights can be established with the Web3 programming language, allowing for arbitrarily complex, dynamic access right management. Practically, we can assume that the Web3 defines an access rights language (that can itself be considered part of the Web3 programming language).

Eric Schmidt and Jared Cohen have explained the “future of identity” [75] in the “new digital age” [76] as follows: “The shift from having one’s identity shaped off-line and projected online to an identity that is fashioned online and experienced off-line will have implications for citizens, states and companies as they navigate the new digital world.” [75] We postulate that the Web3 principle of pervasive digital rights is of utmost significance for changing the concepts of online identities and identities.

4.2 The Ledger and the Cloud

The Web3 tree τ is a purely conceptual model. It explains the informational structure of the Web3 and, most importantly, introduces the notion of Web3 partition. The Web3 cloud γ and the Web3 ledger λ together provide the concrete realization of the Web3. The cloud γ stores the full version history of the Web3 tree γ. The Web3 ledger λ provides certificates for Web3 information. It is the fully certified complete protocol Web3 log. Occasionally, we therefore call the ledger λ also the Web3 *certification ledger*

The Web3 cloud and ledger can be implemented as overlay network to any internet protocol stack such as, of course, the TCP/IP protocol stack of today’s Internet. The Web3 tree τ manifests merely through application-layer protocols that are kept free from any lower-layer concepts and, therefore, is independent of any changes to lower protocol layers. Today’s dominating Web protocol *http* relies on IP addresses and is therefore intertwingled with the current Internet protocol stack at the Internet layer. DNS (Domain Name Service) is designed as an aftermath to *http*. The functioning of today’s Web tree is anchored in trust into the centralized mechanism of IP address allocation – provided by the Internet organizations ICANN (Internet Corporation for Assigned Names and Numbers) and IANA (Internet Assigned Numbers Authority). Trust in today’s Web tree is rooted in trust into ICANN and IANA. The Web3 tree τ can gain trust from the Web3 ledger λ in its role as the certification backbone of the Web3. Again, the ledger can be kept free from any lower-layer concepts. It is Web3 cloud component γ that needs to be related to a concrete internet protocol stack when it is realized.

Data Abstraction and Livestreaming The concept of the Web3 ledger λ can be explained best through two architectural principles that go hand-in-hand with each other: the Web3 *data abstraction principle* and Web3 *livestreaming*. Data abstraction is a core software engineering principle [77]. In the context of the Web3 it means, that the Web3 tree γ is manipulated and *only* manipulated through a set of well-defined Web3 protocols. This is not so in the Web. From the beginning [78], the *http* protocol had a *post* method, which allows for adding new data to a Web server. Soon after [79], the *http* protocol was enriched by a *put* method that allows for updating a specific web resource. The point is that the *post* and the *put* method are rather seldomly used in practice. Instead, Web resources are manipulated by all kinds of means, i.e., direct writes to the file system, mitigated by a web content management system etc. This means that the complete log of Web protocol activities would not reflect at all the actual Web version history.

In our Web3 foundation, it is an architectural principle that all Web3 protocol activities are recorded in the Web3 ledger λ and we call this principle Web3 *livestreaming*. Of course, this principle makes only sense if the Web3 is always only manipulated through defined Web3 protocols – which is the essence of the Web3 *data abstraction principle*. We postulate, that the Web3 ledger λ is the only authoritative reference for Web3 content. As such is is certified, becoming: the fully certified complete Web3 log.

Following the current state of the art, a natural candidate to implement the Web3 ledger λ is with today's blockchain technology [80,52,81,82]. The reason for this is the efficiency of the blockchain data structure. We can assume that verifying a signature is thousand times more costly than computing a hash [10], which leads to the concept of organizing the data in blocks, computing a Merkle tree per block and signing this tree (via its root hash). This efficiency argument holds independent of the concrete consensus mechanism of a blockchain or the question whether the blockchain is permissionless or permissioned etc.

Maximizing Data Protection There are two fundamental options to distribute Web3 data over the Web3 tree τ and the Web3 ledger λ:

- *Pure certification ledger.* No data is stored in the ledger λ. All Web3 data is stored (only) in cloud γ. A chunk of Web3 data δ that is exchanged via a Web3 protocols is represented in the ledger by a hash value h_δ. The hash value h_δ serves as identifier of δ, i.e., to retrieve δ from γ.
- *Certification/data ledger.* Some of the data that is exchanged via Web3 protocols are stored directly in the ledger λ.

A reason for not storing data directly in the ledger is efficiency. This reason is independent of whether the ledger is public or not. If the ledger is public, a natural pattern (to maximize data protection) is to formulate Web3 (partition) rules only in terms of data stored in the ledger λ – the Web3 rules can now be called *ledger rules*. Then, assuming that the data in the cloud γ is not public (and effectively protected), only such data would be stored in the ledger that is needed in formulating Web3 (partition) rules. Storing clear data in the ledger does not automatically break anonymity.

Ultra Scalability Ledger transaction performance is the *sine-qua-non* precondition for the Web3 to be turned into reality. We discuss ultra scalability as part of the Alphabill scenario discussion in Sect. 6.

5 On the Web3 Design Space

5.1 A Better Web

The Web3 is said to be "our chance to make a better internet" [1]. A "better Web" has been envisioned long before the Web. Already in 1960, Ted Nelson founded project Xanadu[16] [83,84] (See Fig. 2) – the original hypertext [60] project. Today, more than 50 years later, the requirements that have been formulated for Xanadu (Table 3) read like a wish list for the "better internet" including: a document type system, transclusions [60,61], secure user identification, access rights management, data replication etc. Last but not least, a royalty mechanism and payment system for the consumption of digital assets was in the Xanadu list.

1. Every Xanadu server is uniquely and securely identified.
2. Every Xanadu server can be operated independently or in a network.
3. Every user is uniquely and securely identified.
4. Every user can search, retrieve, create and store documents.
5. Every document can consist of any number of parts each of which may be of any data type.
6. Every document can contain links of any type including virtual copies ("transclusions") to any other document in the system accessible to its owner.
7. Links are visible and can be followed from all endpoints.
8. Permission to link to a document is explicitly granted by the act of publication.
9. Every document can contain a royalty mechanism at any desired degree of granularity to ensure payment on any portion accessed, including virtual copies ("transclusions") of all or part of the document.
10. Every document is uniquely and securely identified.
11. Every document can have secure access controls.
12. Every document can be rapidly searched, stored and retrieved without user knowledge of where it is physically stored.
13. Every document is automatically moved to physical storage appropriate to its frequency of access from any given location.
14. Every document is automatically stored redundantly to maintain availability even in case of a disaster.
15. Every Xanadu service provider can charge their users at any rate they choose for the storage, retrieval and publishing of documents.
16. Every transaction is secure and auditable only by the parties to that transaction.
17. The Xanadu client-server communication protocol is an openly published standard. Third-party software development and integration is encouraged.

Table 3. The original 17 rules of Ted Nelson's Xanadu.

5.2 Amalgamation of Intranet and Internet

A Massive Enterprise Application Backbone

[16] `https://xanadu.com.au`

Analysing Today's Enterprise Application Landscape The fact that today's enterprise applications are implemented as web-based applications gives us an idea of another huge opportunity for Web3 that has been overlooked so far: the systematic *amalgamation of intranet and internet* (where we think of the intranet as a potential *enterprise application backbone* [85,86]. Analysing today's enterprise application landscape [85,86] leads to similar requirements in regard to crosscutting concerns as expressed by the rules of Ted Nelson's Xanadu in Table 3.

Today's office work heavily relies on office applications (spreadsheets, text processing, presentation programs) and ERP (enterprise resource planning) systems. Unfortunately, today's digitalized office work is light years away from being optimal as it suffers a huge legacy problem as follows. The problem is that today's office applications and ERP systems are all silos. In office work, huge amount of time is wasted to manually pump around data (copy-and-paste) between office applications, ERP systems and the many further ERP-related enterprise applications: customer relationship management (CRM) systems, document management systems (DMS), business intelligence (BI) applications etc. [87,88]. In addition to the wasted labor efforts, this *copying-and-pasting* introduces errors (directly), leads to unwanted redundancies, data inconsistencies (indirectly over time) and access right violations.

Equally worse is what we call the *digital punk* of today's office work. In the *steampunk* genre, anachronistic technologies are used to realize futuristic machines and crafts; in *digital punk*, we use futuristic technologies to realize anachronistic work: electronic documents stemming from office applications are copied and sent around (via e-*mail*) as if they were paper (sometimes together with a confidentially note: fyeo, do not distribute etc.), instead of simply granting access to their content – again leading to redundancies, inconsistencies and access right violations. This means that in today's offices, latest technology is used to simulate 19th century work routines (*mail* has just become e-*mail*!).

All of this leads to a situation that is fragile, error prone, blocks business process re-engineering [89] efforts and prohibits systematic enterprise data analytics.

Usually, enterprise content management (ECM) systems [90] are considered to help with the described problems. However, actually they only help a little as they do not fundamentally improve the situation. With their features (versioning, access right management, and meta data management) and because of their *end-user friendliness* (as opposed to system administrator tools), ECM systems help organizing the chaos. However, they only help to deal with the symptoms but do not tackle the underlying problem, i.e., they do not crack the silos.

Enterprise wikis [91] could help at least with today's offices' *digital punk*, but only to a certain degree. Wikis [92] are great, whenever collaborative editing of semi-structured data is sufficient to support a task. They fail, whenever end-user features of today's office applications are deserved. Even for those tasks for which they would be the optimal choice, their potential is not well understood by today's CIOs [87].

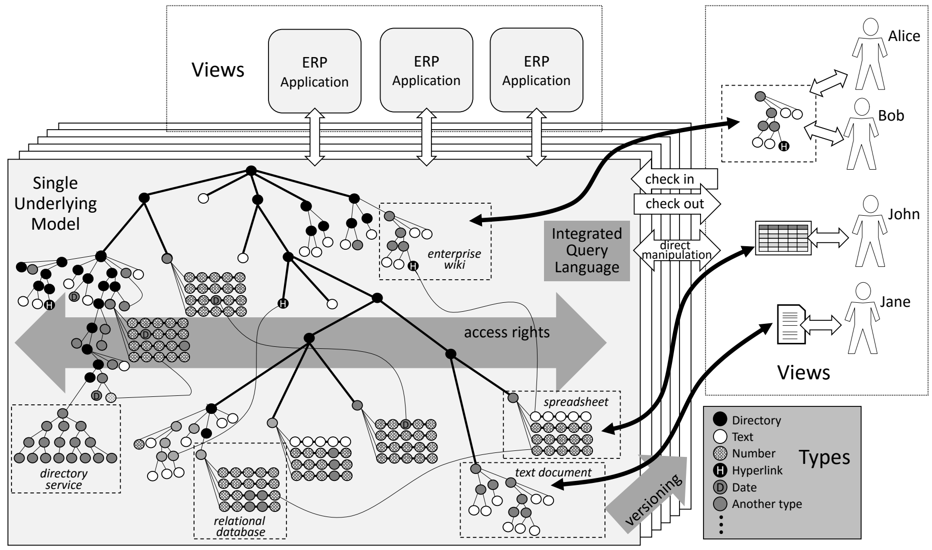

Fig. 4. Web3 as a massive enterprise application backbone – conceptual viewpoint.

Again, also the existing enterprise application integration (EAI) approaches are considered as solutions to help with the current situation. The subsequent discussion applies to all of the existing EAI approaches, including ad-hoc glue code [93], SOA platforms [94,95,96,97,98], and BPM suites [99] [100] [101] [102] [103] [104] [105] [106] [107] [108]. EAI tools are important, but again are no game changers. EAI needs to be used to integrate enterprise applications immediately, whenever manually *copying-and-pasting* data is only a theoretical option, i.e., whenever the data that needs to be exchanged needs to be processed too fast, is too large or too complex. That said, EAI is not very much more than automatic data exchange between enterprise application. The several approaches differ only in the level of support features. Again, EAI does treat only the symptoms, not the underlying problem.

De- and Re-Construction Today's Enterprise Applications Considering all of these problems, we can come to the conclusion that we are trapped in a huge path dependency [109]. Let us try to break free from it, and let us try to fundamentally rethink today's enterprise application landscape. What we suggest is a so-called *massive backbone application*, see Fig. 4, that

(i) hosts all enterprise data of an organization, i.e., all data of all enterprise applications of an organization, as a *single underlying model* (SUM) [110,111,112] and, furthermore,
(ii) provides features for essential cross-cutting concerns needed for all enterprise data and all enterprise applications.

With respect to (i), all enterprise data of an organization is stored as semi-structured data as part of the Web3 tree τ, compare to Figs. 1 and 3, i.e., the tree in Fig. 4 is a sub tree (typically a partition) of the Web tree τ in Fig. 3. The whole tree in Fig. 4 has become the *single underlying model* (SUM) [112,111,113] of all enterprise applications of the respective organization. Each of the several data models of today's information system landscape is enforced by a Web language w. In Fig. 4, we have, for example, depicted a spreadsheet, a text document, a relational database, a directory service and an enterprise wiki. End-users access the SUM via clients of their choice (direct manipulation, form-based [114], versioning systems etc.). Enterprise application access the SUM and this way again create views onto the SUM.

The SUM in Fig. 4 represents a conceptual viewpoint as follows. Theoretically, today's information system products could all be re-implemented from scratch on top of the SUM as a database. However, in the sense of a migration path, a given sub tree of the SUM can also realized as view itself, i.e., as a view onto an existing information system product – turning the enterprise application backbone in Fig. 4 more into a middleware than into a database.

With respect to (ii), we identify the following as essential crosscutting concerns:

- access rights mechanisms [115],
- version control [116,117,118,119],
- an integrated querying language [120].

Through support for those crosscutting concerns the enterprise application backbone in Fig. 4 becomes a value-added middleware for enterprise applications.

A Massive Interorganizational Application Backbone The concept of enterprise application backbone described in Sect. 5.2 is not limited by the boarders of the single organization. It naturally extends to business-to-business communication (B2B) [96,104,96,98,113,85] and [121,122,123], turning B2B communication into cross-organizational enterprise application integration; turning Web3 into a *world computer*. See Fig. 5, where the SUMs of four organizations (each of them as an instance of Fig. 4) all become sub trees of a larger integrated SUM.

Imagine the opportunities of such level of Web3 integration allowing for cross-organizational transclusions [60,61], taking *(web) weaving* [124,125,126] to a next level (just as one example: an organization could grant access to a single column in one of its Excel tables for re-publishing in a different context by another organization, as a paid service – a combination of rules (6), (7) and (9) in Table 3).

A Massive Devops Backbone In the same vein, the Web3 can serve as a *massive Devops backbone*. The SUM (Single Underlying Model) described in the orthographic modeling approach [110] can be integrated as partition into Web3 – enabling both CASE 2.0 as described in [111,112] as well as the model-driven organization as described in [113].

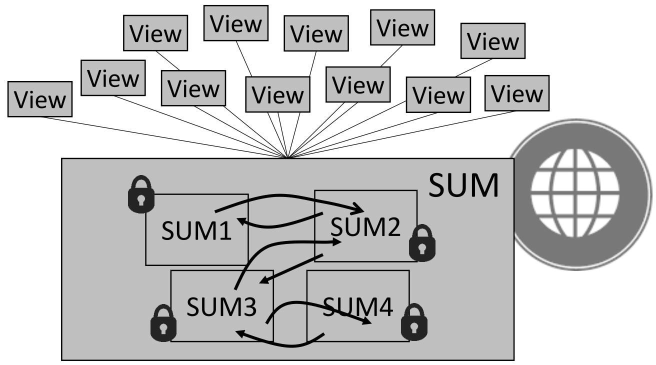

Fig. 5. Web3 as a massive interorganizational (B2B) application backbone – conceptual viewpoint.

5.3 A Web of Everything

As the *Web of everything* our vision of Web3 encompasses the Internet of Things (IoT). And as such, it becomes a *Web of manufacturing and logistics* [127,128,129,130]. The integration of SDN (software-defined networking), IoT and blockchain technology [131,132,133] will stay with us as an important strand of research contributing to Web3.

5.4 A Web of Everybody

As a *Web of everybody*, massive *disintermediation* is the standard narrative of the Web3. Disintermediation leads to re-shaped institutions [71,72,134] as well as entirely new institutions. As the societies' institutional architecture [73], governance needs to be re-thought and re-designed.

Collective intelligence (CI) [135,54,136] systems form an extremely important class of web-based applications with Wikipedia and Reddit being just two examples [54]. CI systems are natural candidates for Web3 partitions. CI systems will stay with us in the future and their importance will even steadily increase. For example, enterprises have started to understand the potential of CI for their endeavors [86] – take Blackrock's Aladdin[17] system and Genpact's Cora system[18] as (particularly important) examples.

6 The Alphabill Scenario

Recently [10], we have described the Alphabill platform and its architecture, see Fig. 6. Alphabill is a platform for universal asset tokenization, transfer and

[17] `https://www.blackrock.com/aladdin`
[18] `https://www.genpact.com/cora`

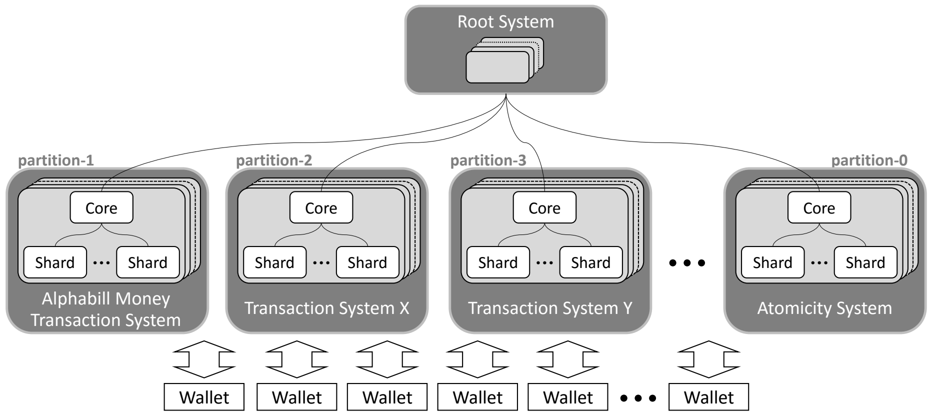

Fig. 6. The Alphabill platform.

exchange as a global medium of exchange. Users of the Alphabill platform can launch arbitrarily many partitions on the platform. Alphabill is a partitioned, replicated, sharded blockchain. Each partition implements an individual token and corresponding transaction system. Alphabill partitions correspond to the notion of Web3 partitions in our Web3 foundation. The Alphabill platform provides the necessary protocols, languages, libraries and toolkits to implement partitions in such a way that they show robustness and unlimited scalability. Robustness is achieved through replication, i.e., highly redundant partitions. The ultra scalability is enabled by a novel electronic money scheme, the bill scheme [24,25]. Each Alphabill partition is sharded. Through its decomposability, the bill money scheme eliminates coordination efforts between shards. Coordination between partitions is achieved efficiently through a dedicated atomicity partition and a novel, three-phase commit protocol.

As a proof-of-concept, we have successfully delivered the bill-based blockchain technology KSI Cash [10,137,138,25,139]. The performance of the technology has been tested exhaustively, together with the European Central Bank, in order to assess the technological feasibility of a digital euro [10]. The tests achieved: (i) 15 thousand transactions per second, under simulation of realistic usage, with 100 million wallets, and (ii) up to 2 million payment orders per second, i.e., an equivalent of more than 300,000 transactions per second, in a laboratory setting with the central components of KSI Cash.

7 Conclusion

Too often, we think and talk about Web3 in terms of individual Web3 solutions (individual Web3 products, individual Web3 assets, individual Web3 business models etc.) – although Web3 is clearly a vision of a digital ecosystem, and, actually, a vision of the most encompassing digital ecosystem. In our opinion, it is unlikely that Web3 emerges – out of nothing – as a series of Web3 solutions (independent of how much venture capital might be pumped into such individual efforts). What we need in first place, is to shape and to provide excellent (ultra-useful, ultra-easy, ultra-robust, ultra-scalable[19]) infrastructure and tools to enable Web3 solutions. And, we are convinced, now is the time to do so. From an engineering perspective, Web3 is the integration of digital rights exchange into the (application layer) internet protocols. From a design perspective, we need to care more for the completeness of vision of Web3. With the Web3 foundation suggested in this paper, we hope to help with the completeness of vision of Web3.

References

1. Jin, L., Parrott, K.: Web3 is our chance to make a better Internet. Harvard Business Review **10 May** (2022) `https://hbr.org/2022/05/web3-is-our-chance-to-make-a-better-internet`.
2. Wiles, J.: What is Web3? Gartner (15 February 2022) `https://www.gartner.com/en/articles/what-is-web3`.
3. Bennett, M.: Web3 isn't going to fix the shortcomings of today's Web. Forrester (10 May 2022) `https://www.forrester.com/blogs/web3-isnt-going-to-fix-the-shortcomings-of-todays-web/`.
4. Platz, B.: Why Web3 is so confusing. Forbes Technology Council (1 June 2022) `https://www.forbes.com/sites/forbestechcouncil/2022/06/01/why-web3-is-so-confusing/`.
5. Stackpole, T.: What is Web3? Harvard Business Review **10 May** (2022) `https://hbr.org/2022/05/what-is-web3`.
6. Esber, J., Kominers, S.D.: Why build in Web3. Harvard Business Review **16 May** (2022) `https://hbr.org/2022/05/why-build-in-web3`.
7. Shannon, V.: A 'more revolutionary' Web. The New York Times **23 May** (2006)
8. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web – a new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. Scientific American **17 May** (2001)
9. Koziol, M.: 5 questions for Molly White: The "Web3 is going just great" creator on why it isn't. IEEE Spectrum **60**(4) (2023) 21
10. Buldas, A., Draheim, D., Gault, M., Laanoja, R., Nagumo, T., Saarepera, M., Shah, S.A., Simm, J., Steiner, J., Tammet, T., Truu, A.: An ultra-scalable blockchain platform for universal asset tokenization: Design and implementation. IEEE Access **10** (2022) 77284–77322 `https://doi.org/10.1109/ACCESS.2022.3192837`.
11. Sadowski, J., Beegle, K.: Expansive and extractive networks of Web3. Big Data and Society **10**(1) (2023)
12. Ruane, J., McAffee, A.: What a DAO can – and can't – do. Harvard Business Review **10 May** (2022) `https://hbr.org/2022/05/what-a-dao-can-and-cant-do`.
13. Wang, S., Ding, W., Li, J., Yuan, Y., Ouyang, L., Wang, F.Y.: Decentralized autonomous organizations: Concept, model, and applications. IEEE Trans. Computat. Social Syst. **6**(5) (2019) 870–878

[19] Ultra-scalability is the sine-qua-none pre-condition for any Web3 vision to take of [10].

14. DuPont, Q.: Experiments in algorithmic governance: A history and ethnography of “The DAO," a failed decentralized autonomous organization. In Campbell-Verduyn, M., ed.: Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global Governance. Routledge (2017) 1–18
15. Rikken, O., Janssen, M., Kwee, Z.: Governance challenges of blockchain and decentralized autonomous organizations. Information Polity **24**(4) (2019) 397–417
16. Qin, R., Ding, W., Li, J., Guan, S., Wang, G., Ren, Y., Qu, Z.: Web3-based Decentralized Autonomous Organizations and operations: Architectures, models, and mechanisms. IEEE Transactions on Systems, Man, and Cybernetics: Systems **53**(4) (2023) 2073–2082
17. Murray, A., Kim, D., Combs, J.: The promise of a decentralized internet: What is Web3 and how can firms prepare? Business Horizons **66**(2) (2023) 191—-202
18. Afuah, A.: Are network effects really all about size? the role of structure and conduct. Strategic Management **34**(3) (2012) 257–273
19. Nabben, K.: Web3 as ‘self-infrastructuring’: The challenge is how. Big Data and Society **10**(1) (2023)
20. Federal Reserve Bank of Boston and Massachusetts Institute of Technology Digital Currency Initiative: Project Hamilton Phase 1 – A High Performance Payment Processing System Designed for Central Bank Digital Currencies. Federal Reserve Bank of Boston (3 February 2022) [last accessed: 26 March 2022] https://www.bostonfed.org/-/media/Documents/Project-Hamilton/Project-Hamilton-Phase-1-Whitepaper.pdf.
21. Danezis, G., Meiklejohn, S.: Centrally banked cryptocurrencies. In: Proc. of NDSS’16 – the 23rd Annual Symposium on Network and Distributed System Security, The Internet Society (2016) 1–14
22. Wood, G.: Polkadot: Vision for a heterogeneous multi-chain framework, Draft 1 (2016) [last accessed: 30 April 2023] https://polkadot.network/PolkaDotPaper.pdf.
23. Buldas, A., Draheim, D., Saarepera, M.: Secure and efficient implementation of electronic money. In: Proc. of FDSE’2022 - the 9th Intl. Conf. on Future Data and Security Engineering. Volume 1688 of CCIS., Springer (2022) 34–51
24. Buldas, A., Saarepera, M., Steiner, J., Draheim, D.: A unifying theory of electronic money and payment systems. TechRxiv (2021) https://doi.org/10.36227/techrxiv.14994558.v1.
25. Buldas, A., Saarepera, M., Steiner, J., Ilves, L., Olt, R., Meidla, T.: Formal Model of Money Schemes and their Implications for Central Bank Digital Currency. Eesti Pank, Guardtime (2021) [last accessed: 30 April 2023] https://haldus.eestipank.ee/sites/default/files/2021-12/EP-A_Formal_Model_of_Money_2021_eng.pdf.
26. Lukaj, V., Martella, F., Fazio, M., Ruggeri, A., Celesti, A., Villari, M.: A blockchain based federated ecosystem for tracking and validating the authenticity of goods. In: 2022 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCom/CyberSciTech), IEEE (2022) 1–7
27. Lukaj, V., Martella, F., Fazio, M., Celesti, A., Villari, M.: Establishment of a trusted environment for iot service provisioning based on x3dh-based brokering and federated blockchain. Internet of Things **21** (2023)

28. Nguyen, C., Hoang, D., Nguyen, D., Xiao, Y., Pham, H., Dutkiewicz, E., Tuong, N.: Fedchain: Secure proof-of-stake-based framework for federated-blockchain systems. IEEE Transactions on Services Computing (2023) 1–14
29. Singh, A., Click, K., Parizi, R.M., Zhang, Q., Dehghantanha, A., Choo, K.K.R.: Sidechain technologies in blockchain networks: An examination and state-of-the-art review. Journal of Network and Computer Applications **149** (2020) 1–16
30. Han, Z., Niyato, D., Saad, W., Basar, T., Hjorungnes, A.: Game Theory in Wireless and Communication Networks: Theory, Models, and Applications. Cambridge University Press (2012)
31. Ma, B., Zheng, X., Zhao, C., Wang, Y., Wang, D., Meng, B.: A secure and decentralized ssi authentication protocol with privacy protection and finegrained access control based on federated blockchain. PLoS ONE **17**(9 September) (2022)
32. Zheng, X., Zheng, F., Liu, X., Wang, D., Wang, J., Meng, B.: A secure and policy-controlled signature scheme with strong expressiveness and privacy-preserving policy. IEEE Access **9** (2021) 14945–14957
33. Eldin, A.M., Hossny, E., Wassif, K., Omara, F.A.: Survey of blockchain methodologies in the healthcare industry. In: Proc. of ICCI'2022 – the 5th Intl. Conf. on Computing and Informatics, IEEE (2022) 209–215
34. Mohey Eldin, A., Hossny, E., Wassif, K., Omara, F.: Federated blockchain system (FBS) for the healthcare industry. Scientific Reports **13**(1) (2023)
35. Buterin, V.: Vitalik Buterin explains Ethereum (2015) [last accessed: 30 April 2023] `https://www.youtube.com/watch?v=TDGq4aeevgY`.
36. Stiglitz, J.E.: The Roaring Nineties: A New History of the World's Most Prosperous Decade. W. W. Norton (2004)
37. Lohr, S.: He created the Web. now he's out to remake the digital world. The New York Times **10 January** (2021)
38. Mansour, E., Sambra, A.V., Hawke, S., Zereba, M., Capadisli, S., Ghanem, A., Aboulnaga, A., Berners-Lee, T.: A demonstration of the solid platform for social web applications. In: WWW 2016 Companion - Proceedings of the 25th International Conference on World Wide Web, ACM (2016) 223–226
39. Haber, S., Stornetta, W.: How to time-stamp a digital document. Cryptology **3** (1991) 99—111
40. Edelman, G.: Paradise at the cryto arcade. Wired Magazine **June** (2022) 47–59
41. Lagarde, C.: Central banking and Fintech – A brave new world? Bank of England Conference, London, 29 September (2017)
42. Fernández-Villaverde, J., Sanches, D., Schilling, L., Uhlig, H.: Central bank digital currency: Central banking for all? Review of Economic Dynamics **41** (2021) 225–242
43. Grassi, L., Lanfranchi, D., Faes, A., Renga, F.M.: Do we still need financial intermediation? The case of decentralized finance – DeFi. Qualitative Research in Accounting & Management (February 2022) 1–22
44. Zetzsche, D.A., Arner, D.W., Buckley, R.P.: Decentralized finance. Journal of Financial Regulation **6**(2) (2020) 172–203
45. Chen, Y., Bellavitis, C.: Blockchain disruption and decentralized finance: The rise of decentralized business models. Journal of Business Venturing Insights **13**(e00151) (2020) 1–8
46. Schär, F.: Decentralized finance: On blockchain- and smart contract-based financial markets. Federal Reserve Bank of St. Louis Review **103**(2) (2021) 153–74
47. Lips, S., Barhosa, N., Draheim, D.: eIDAS implementation challenges: the case of Estonia and the Netherlands. In: Proc. of EGOSE'2020: the 7th Intl. Conf. on

Electronic Governance and Open Society – Challenges in Eurasia. Volume 947 of Communications in Computer and Information Science., Springer (2019)
48. Tsap, V., Lips, S., Draheim, D.: eID public acceptance in Estonia: Towards understanding the citizen. In: Proc. of dg.o 2020 – the 21st Annual Intl. Conf. on Digital Government Research, ACM (2020)
49. Lips, S., Tsap, V., Bharosa, N., Krimmer, R., Tammet, T., Draheim, D.: Management of national eid infrastructure as a state-critical asset and public-private partnership: Learning from the case of estonia. Information Systems Frontiers (2023)
50. International Organization for Standardization: ISO 29003:2018 – Information technology — Security techniques — Identity proofing. ISO (2018)
51. Mühle, A., Grüner, A., Gayvoronskaya, T., Meinel, C.: A survey on essential components of a self-sovereign identity. Computer Science Review **30** (2018) 80–86
52. Andy Oram (ed.): Peer to Peer: Harnessing the Power of Disruptive Technologies. O'Reilly (2001)
53. Voshmgir, S.: Token Economy – How the Web3 reinvents the Internet, 2nd. ed. BlochainHub Berlin (2020)
54. Suran, S., Pattanaik, V., Draheim, D.: Frameworks for collective intelligence: A systematic literature review. ACM Computing Surveys **52**(1) (2020) 1–36
55. Draheim, D.: The service-oriented metaphor deciphered. Journal of Computing Science and Engineering **4**(4) (2010) 253–275
56. Niforos, M.: The promising future of NFTs remains in a state of flux. Financial Times (20 June 2022)
57. Dowling, M.: Fertile LAND: Pricing non-fungible tokens. Finance Research Letters **44**(102096) (2022) 2–5
58. Sharma, R.: What is a non-fungible token (NFT)? Investopedia **22 June** (2022)
59. Buldas, A., Draheim, D., Saarepera, M.: Towards a foundation of Web3. In: Proc. of FDSE'2022 - the 9th Intl. Conf. on Future Data and Security Engineering. Volume 1688 of CCIS., Springer (2022) 3—-18
60. Nelson, T.H.: A file structure for the complex, the changing and the indeterminate. In: Proc. of ACM'65 – the 20th ACM National Conference, ACM (1965) 84–100
61. Nelson, T.H.: The heart of connection: Hypermedia unified by transclusion. Communications of the ACM **38**(8) (1995) 31–33
62. Draheim, D., Weber, G.: The integrated source code paradigm. In: Form-Oriented Analysis. Springer (2005) 229–247
63. Draheim, D., Weber, G.: Form-Oriented Analysis. Springer (2005)
64. Lazuashvili, N., Norta, A., Pappel, I., Draheim, D.: Integration of blockchain technology into a land registration system for immutable traceability: A case study of Georgia. In: Poceedings of the Central and Eastern Forum at BPM'2109 – the 17th Intl. Conf. on Business Process Management. Volume 361 of Lecture Notes in Business Information Processing., Springer (2019) 219–233
65. Akaba, T.I., Norta, A., Udokwu, C., Draheim, D.: A framework for the adoption of blockchain-based e-procurement systems in the public sector. In: Proc. of I3E'2020 – the 19th IFIP Conference on e-Business, e-Services and e-Society. Volume 12066 of LNCS., Springer (2020) 3–14
66. Abodei, E., Norta, A., Azogu, D., Udokwu, C., Draheim, D.: Blockchain technology for enabling transparent and traceable government collaboration in public project processes of developing economies. In: Proc. of I3E'2019 – the 18th

IFIP Conference on e-Business, e-Services and e-Society. Volume 11701 of LNCS., Springer (2019) 464–475
67. Azogu, I., Norta, A., Draheim, D.: A framework for the adoption of blockchain technology in healthcare information management systems a case study of nigeria. In: Proc. of ICEGOV'2019: the 12th Intl. Conf. on Theory and Practice of Electronic Governance., ACM (2019) 310–316
68. Szabo, N.: Smart Contracts: Building Blocks for Digital Markets. Nick Szabo (1996)
69. Dwivedi, V.K., Pattanaik, V., Deval, V., Dixit, A., Norta, A., Draheim, D.: Legally enforceable smart-contract languages: A systematic literature review. ACM Computing Surveys **54**(5) (2021) 1–34
70. Dixit, A., Deval, V., Dwivedi, V., Norta, A., Draheim, D.: Towards user-centred and legally relevant smart-contract development: A systematic literature review. Journal of Industrial Information Integration **26**(100314) (March 2022) 1–18
71. Williamson, O.E.: Transaction cost economics: the governance of contractual relations. The Journal of Law & Economics **22**(2) (1979) 233–261
72. Williamson, O.E.: Transaction cost economics: How it works; where it is headed. De Economist **146** (1998) 23–58
73. Draheim, D., Krimmer, R., Tammet, T.: On state-level architecture of digital government ecosystems: From ICT-driven to data-centric. Special Issue in Memory of Roland Wagner. Transactions on Large-Scale Data- and Knowledge-Centered Systems **48** (2021) 165–195
74. Buldas, A., Draheim, D., Nagumo, T., Vedeshin, A.: Blockchain technology: Intrinsic technological and socio-economic barriers. In: Proc. of FDSE'2020 – the 7th Intl. Conf. on Future Data and Security Engineering. Volume 12466 of LNCS., Springer (2020) 3–27
75. Schmidt, E., Cohen, J.: The Future of Identity, Citizenship and Reporting. In (Schmidt, E., Cohen, J.): The New Digital Age – Transforming Nations, Businesses, and Our Lives. Vintage Books (2013) 32–81.
76. Schmidt, E., Cohen, J.: The New Digital Age – Transforming Nations, Businesses, and Our Lives. Vintage Books (2013)
77. Liskov, B., Zilles, S.: Programming with abstract data types. SIGPLAN Notices **9**(4) (1974) 50–59
78. Berners-Lee, T., Fielding, R., Frystyk: Hypertext Transfer Protocol – HTTP/1.0. RFC 1945. Network Working Group (1996)
79. Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T.: Hypertext Transfer Protocol – HTTP/1.1. RFC 2616. Network Working Group (1999)
80. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008) [last accessed: 30 April 2023] `https://bitcoin.org/bitcoin.pdf`.
81. Narayanan, A., Clark, J.: Bitcoin's academic pedigree. Communications of the ACM **60**(12) (2017) 36–45
82. Narayanan, A., Clark, J.: Bitcoin's academic pedigree. ACM Queue Magazine **15**(4) (2017) 1–30
83. Nelson, T.H.: Literary Machines: The Report On, and Of, Project Xanadu Concerning Word Processing, Electronic Publishing, Hypertext, Thinkertoys, Tomorrow's Intellectual Revolution, and Certain Other Topics Including Knowledge, Education and Freedom. Mindful Press (1993)
84. Knowlton, K.: Ted Nelson's Xanadu. In Dechow, D.R., Struppa, D.C., eds.: Intertwingled – The Work and Influence of Ted Nelson. History of Computing. Springer (2015) 25–28

85. Draheim, D.: On the radical de- and re-construction of today's enterprise applications CENTERIS'2019 Keynote. In: Proc. of CENTERIS'2019 – the 10th Conference on Enterprise Information Systems. Volume 164 of Procedia Computer Science., Elsevier (2019) 120–122
86. Draheim, D.: Collective intelligence systems from an organizational perspective – iiWAS'2019 Keynote. In: Proc. of iiWAS'2019 – the 21st Intl. Conf. on Information Integration and Web-based Applications & Services, ACM (2019) 3–4
87. Draheim, D.: Smart business process management. In Fisher, L., ed.: 2011 BPM and Workflow Handbook, Digital Edition. Future Strategies, Workflow Management Coalition (2012) 207–223
88. Draheim, D.: Towards total budgeting and the interactive budget warehouse. In: Innovation and Future of Enterprise Information Systems. Volume 4 of LNISO., Springer (2013) 271–286
89. Hammer, M., Champy, J.: Reengineering the Corporation. HarperCollins (1993)
90. Karen Hobert, Michael Woodbridge, M.B.: 2018 magic quadrant for content services platforms. Technical Report G00343925, Gartner (2018)
91. Gotta, M., Mann, J., Drakos, N.: Magic quadrant for social software in the workplace. Technical Report G00270286, Gartner (2015)
92. Cunningham, W.: Design principles of wiki: how can so little do so much? In: Poceedings of Wikis'2006 – the 2nd International Symposium on Wikis, ACM (2006) 13–14
93. Ousterhout, J.: Scripting: Higher level programming for the 21st century. IEEE Computer Magazine **31**(3) (1998) 23–30
94. McCoy, D., Nati, Y.: Service-oriented architecture: Mainstream straight ahaed. Technical Report G00114361, Gartner (2003)
95. Natis, Y.: Service-oriented architecture scenario. Technical Report G00114358, Gartner (2003)
96. Draheim, D.: The service-oriented metaphor deciphered. Journal of Computing Science and Engineering **4**(4) (2010) 253–275
97. Atkinson, C., Bostan, P., Draheim, D.: A unified conceptual framework for service-oriented computing - aligning models of architecture and utilization. Transactions on Large-Scale Data- and Knowledge-Centered Systems **7** (2013) 128–169
98. Atkinson, C., Bostan, P., Draheim, D.: Foundational MDA patterns for service-oriented computing. The Journal Of Object Technology **14**(1) (2015) 1–30
99. Leymann, F., Altenhuber, W.: Managing business processes as an information resource. IBM Systems Journal **33**(2) (1994) 326–348
100. van der Aalst, W.M.P., van Hee, K.M.: Framework for business process redesign. In: Proceedings WETICE'95 – the 4th IEEE Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises. (1995) 36–45
101. Dumas, M., van der Aalst, W.M.P., ter Hofstede, A.H.M., eds.: Process-Aware Information Systems: Bridging People and Software through Process Technology. John Wiley & Sons (2005)
102. Auer, D., Draheim, D., Geist, V.: Extending BPMN with submit/response-style user interaction modeling. In: Proc. of BPMN Workshop, CEC'09 – the 11th IEEE Conference on Commerce and Enterprise Computing, IEEE (2009) 368–374
103. Draheim, D.: Business Process Technology – A Unified View on Business Processes, Workflows and Enterprise Applications. Springer (2010)
104. Atkinson, C., Draheim, D., Geist, V.: Typed business process specification. In: Proc. of EDOC'2010 – the 14th IEEE International Enterprise Computing Conference, IEEE (2010) 69–78

105. Draheim, D., Geist, V., Natschläger, C.: Integrated framework for seamless modeling of business and technical aspects in process-oriented enterprise applications. International Journal of Software Engineering and Knowledge Engineering **22**(5) (2012) 645–674
106. Auer, D., Draheim, D., Geist, V., Kopetzky, T., Küng, J., Natschläger, C.: Towards a framework and platform for mobile, distributed workflow enactment services – on a possible future of ERP infrastructure. In: Innovation and Future of Enterprise Information Systems, Springer (2013) 201–215
107. Draheim, D., Kopetzky, T., Küng, J.: How to make mobile BPM robust and intelligent. In: 2013 BPM and Workflow Handbook, Digital Edition. Future Strategies, Workflow Management Coalition (2013) 107–116
108. Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A.: Fundamentals of Business Process Management. Springer (2018)
109. Castaldi, C., Dosi, G., Paraskevopoulou, E.: Path dependence in technologies and organizations. In Augier, M., Teece, D., eds.: The Palgrave Encyclopedia of Strategic Management. Palgrave Macmillan (2018) 1–4
110. Atkinson, C., Stoll, D.: Orthographic modeling environment. In: Proc. of FASE'2008 – the 11th Intl. Conf. on Fundamental Approaches to Software Engineering. Volume 4961 of LNCS., Springer (2008) 93–96
111. Atkinson, C., Draheim, D.: Cloud aided-software engineering – evolving viable software systems through a web of views. In: Software Engineering Frameworks for the Cloud Computing Paradigm. Springer (2013) 255–281
112. Draheim, D.: Case 2.0 – On key success factors for cloud-aided software engineering. In: Proc. of MDHPCL'12 - the 1st International Workshop on Model-Driven Engineering for High Performance and Cloud Computing, ACM (2012) 1–6
113. Tunjic, C., Atkinson, C., Draheim, D.: Supporting the model-driven organization vision through deep, orthographic modeling. EMISAJ (Enterprise Modeling and Information Systems Architectures Journal) **13** (2018) 1–39
114. Draheim, D., Weber, G.: Modelling form-based interfaces with bipartite state machines. Interacting with Computers **17**(2) (2005) 207–228
115. Portable Applications Standards Committee of the IEEE Computer Society, The Open Group: Draft standard for information technology – Portable Operating System Interface (POSIX). Technical Report IEEE P1003.1, Draft 3, IEEE (2007)
116. Shigo, O., Wada, Y., Terashima, Y., Iwamoto, K., Nishimura, T.: Configuration control for evolutional software products. In: Proc. of ICSE'82 – the 6th Intl. Conf. on Software Engineering, IEEE Computer Society Press (1982) 68–75
117. Tichy, W.F.: Design, implementation, and evaluation of a revision control system. In: Proc. of ICSE'82 – the 6th Intl. Conf. on Software Engineering, IEEE Computer Society Press (1982) 58–67
118. Lampson, B.W., Schmidt, E.E.: Organizing software in a distributed environment. SIGPLAN Notices **18**(6) (1983) 1–13
119. Weinreich, R., Ziebermayr, T., Draheim, D.: A versioning model for enterprise services. In: In Proc. of AINAW'07 – the 21st Intl. Conf. on Advanced Information Networking and Applications Workshops, IEEE (2007) 570–575
120. Draheim, D., Weber, G.: Data Modeling. In: Form-Oriented Analysis – A New Methodology to Model Form-Based Applications. Springer (2005) 109–146
121. Taveter, K., Wagner, G.: A multi-perspective methodology for modelling inter-enterprise business processes. In Arisawa, H., Kambayashi, Y., Kumar, V., Mayr, H.C., Hunt, I., eds.: Proceedings of ER'2001 – the 20th International Conference on Conceptual Modeling. Volume 2465 of Lecture Notes in Computer Science., Berlin, Heidelberg, Springer (2002) 403–416

122. Sulis, E., Taveter, K.: Agent-Based Business Process Simulation – A Primer with Applications and Examples. Springer, Berlin, Heidelberg (2022)
123. Norta, A., Ma, L., Duan, Y., Rull, A., Kõlvart, M., Taveter, K.: eContractual choreography-language properties towards cross-organizational business collaboration. Journal of Internet Services and Applications **6**(8) (2015) 1–24
124. Draheim, D., Felderer, M., Pekar, V.: Weaving social software features into enterprise resource planning systems. In: Novel Methods and Technologies for Enterprise Information Systems. Volume 8 of LNISO., Springer (2014) 223–237
125. Pattanaik, V., Norta, A., Felderer, M., Draheim, D.: Systematic support for the full knowledge management lifecycle by advanced semantic annotation across information system boundaries. In: Proc. of the CAiSE Forum 2018 – Information Systems in the Big Data Era. Number 317 in LNBIP, Springer (2018) 66–73
126. Pattanaik, V., Sharvadze, I., Draheim, D.: A peer-to-peer data sharing framework for web browsers – analysis and evaluation. Springer Nature Computer Science **1**(4) (2020) 1–10
127. Siddiqui, S., Shah, S.A., Ahmad, I., Aneiba, A., Draheim, D., Dustdar, S.: Toward software-defined networking-based IoT frameworks: A systematic literature review, taxonomy, open challenges and prospects. IEEE Access **10** (2022) 70850–70901
128. Hameed, S., Shah, S.A., Saeed, Q.S., Siddiqui, S., Ali, I., Vedeshin, A., Draheim, D.: A scalable key and trust management solution for IoT sensors using SDN and blockchain technology. IEEE Sensors J. **21**(6) (2021) 8716–8733
129. Shah, S.A., Seker, D.Z., Hameed, S., Draheim, D.: The rising role of big data analytics and IoT in disaster management: Recent advances, taxonomy and prospects. IEEE Access **7** (2019) 54595–54614
130. Shah, S.A., Seker, D.Z., Rathore, M., Hameed, S., Yahia, S.B., Draheim, D.: Towards disaster resilient smart cities: Can internet of things and big data analytics be the game changers? IEEE Access **7** (2019) 91885–91903
131. Vedeshin, A., Liiv, I., Ben Yahia, S., Draheim, D.: Smart cyber-physical system for pattern recognition of illegal 3D designs in 3D printing. In: SADASC'2020 – the 3rd Intl. Conf. on Smart Applications and Data Analysis for Smart Cyber-Physical Systems. CCIS 1207, Springer (2020) 74–85
132. Vedeshin, A., Dogru, J.M., Liiv, I., Ben Yahia, S., Draheim, D.: A secure data infrastructure for personal manufacturing based on a novel key-less, byte-less encryption method. IEEE Access **8** (2020) 40039–40056
133. Vedeshin, A., Dogru, J.M.U., Liiv, I., Draheim, D., Ben Yahia, S.: A digital ecosystem for personal manufacturing: An architecture for a cloud-based distributed manufacturing operating system. In: Proc. of MEDES'2019 – the 11th Intl. Conf. on Management of Digital EcoSystems, ACM (2019) 224–228
134. Koppenjan, J., Groenewegen, J.: Institutional design for complex technological systems. International Journal of Technology, Policy and Management **5**(3) (2005) 240–257
135. Malone, T.W., Bernstein, M.S.: Handbook of Collective Intelligence. MIT Press (2015)
136. Suran, S., Pattanaik, V., Kurvers, R.H., Hallin, C.A., de Liddo, A., Krimmer, R., Draheim., D.: Building global societies on collective intelligence: Challenges and opportunities. Digital Government: Research and Practice **3**(4) (2022)
137. European Central Bank, Banca d'Italia, Banco de España, Bank of Greece, Central Bank of Ireland, De Nederlandsche Bank, Eesti Pank, Deutsche Bundesbank, Latvijas Banka: Work Stream 3: A New Solution – Blockchain & eID. July 2021 (2021) [last accessed: 30 April 2023]

https://www.ecb.europa.eu/paym/digital_euro/investigation/profuse/shared/files/deexp/ecb.deexp211011_3.en.pdf.
138. Olt, R., Meidla, T., Ilves, L., Steiner, J.: Summary report: Results of the Eesti Pank - Guardtime CBDC Research. Eesti Pank, Guardtime (December 2021) [last accessed: 30 April 2023] https://haldus.eestipank.ee/sites/default/files/2021-12/EP-Guardtime_CBDC_Research_2021_eng.pdf.
139. Eesti Pank: Eesti Pank ran an experiment to investigate the technological possibilities of a central bank digital currency based on blockchain, Eesti Pank, 13 December (2021) [last accessed: 30 April 2023] https://www.eestipank.ee/en/press/eesti-pank-ran-experiment-investigate-technological-possibilities-central-bank-digital-currency-13122021.