1 Introduction

> Traveler, there is no path. The path is made by walking.
>
> Antonio Machado, *Campos de Castilla*

Transportation is the backbone of our industry. On the last mile to its destination, almost every shipment is moved via trucks. Not only in the final stage but also on the long haul, trucks are crucial to transporting goods. Filling trucks with loads is nontrivial because, usually, freight is either picked up with an empty truck or an empty truck returns from a delivery. In the European Union, 20% of all truck kilometers are empty runs (Eurostat, 2019). In 2016, this led to a cost of 100 billion EUR in unused capacity (Riedl et al., 2018). Not only is the economic cost high but the environmental impact is alarming, as hazardous greenhouse gases are emitted even while transporting no cargo. Transportation was responsible for 27% of global emissions in 2019, which makes it the second largest sector for emissions after energy and heat generation (IEA, 2021). Thus, it is important to connect load with trucks in an efficient way.

One solution to the problem of matching loads with empty trucks is freight exchanges (Goldsby and Eckert, 2003), also known by the more generic term of electronic logistics marketplaces (ELMs). Figure 1.1 shows the main participants in an ELM: shippers, carriers, and technology providers. Shippers can post their loads onto an ELM, and carriers can bid on them. Another mode of ELM use is for carriers to post available loading space, which shippers can then buy. These platforms are offered by technology providers. While this business began when telephones and fax were used for communication, nowadays, all relevant platforms are available via the internet.

Freight exchanges offer many services, but they also have many shortcomings (Wester and Otto, 2021). Issues reported in the literature (Kleedorfer and Huemer, 2017) include the fact that much communication regarding freight exchanges still occurs over the phone. This adds many manual steps to dispatching loads. Due to the heterogeneous IT landscape in the area of logistics, freight exchanges are often poorly integrated

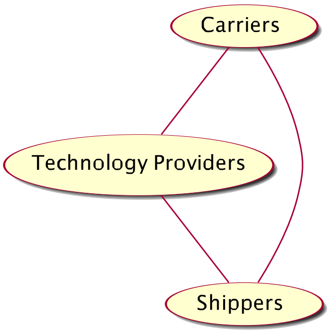

Figure 1.1: Adopted from Wang et al. (2007a).

with carriers' or shippers' transport management systems (Riedl et al., 2018). Therefore, freight exchange processes are inefficient. For instance, payments take up to 60 days after service is conducted (Wester and Otto, 2021). Media disruptions are common, resulting in increased transaction costs. Trust between carriers and shippers is also a serious issue. Cargo may be valuable, and cargo theft by carriers who are not well-validated is a risk (Burges, 2012). Revealing information about customers and trade lanes to a central technology provider is another common concern (Föhring and Zelewski, 2015b), as no participant can be sure about the neutrality of the platform operator. Also, margins in land transport are low (1.8%–5.7% in 2018) (DVZ, 2019; Zacharia et al., 2011), and adding additional intermediaries such as freight exchanges decreases the margins even further (Witkowski, 2018).

Bigger market participants create their own specialized ELMs (e.g., Drive4Schenker (de Sampaio Picão, 2017)), while smaller participants must depend on platform operators for good results. Integration into different ELMs is costly, even if done via manual work, because different markets must be continuously monitored, offers drafted, and offers managed once they are no longer valid. Automation in this area is also costly because every ELM has different application programming interfaces (API) and data formats, if it has them at all.

Startups that want to compete in the ELM market are often prohibited from participating in existing marketplaces, which creates a barrier to market entry (Haas and Seiter, 2020). Even if integrations for supplementary services were possible, the venture would depend on the goodwill of the technology provider. Forward integration from the technology provider would always be possible, but this kind of risk could inhibit the founder's willingness to take risks.

1.1 Motivation and problem statement

Digitalization in the logistics industry is in its beginnings although many areas of the surrounding world are already more digitalized. A BCG study (Riedl et al., 2018) shows problems around "high fragmentation of the market, underutilization of assets, costly manual processes and outdated customer interfaces." This high fragmentation, which should lead to cooperation, according to the study, actually leads to distrust among the actors. Distrust usually results in higher costs because more controls must be established (see Section 2.5). Transparency about "rates, capacity, quality and reliability of the carriers" (Riedl et al., 2018) is missing. The underutilization of assets was discussed in the previous section. Paper-based processes and outdated customer interfaces, used because of legacy systems, are other inhibitors to further automation. There is a call for new marketplaces that would enable new modes of cooperation.

One way of building new infrastructure for exchanging data in a trustworthy way independent of intermediaries is through blockchain technology. Instead of trusting in partners' promises, blockchain users place their trust in technology. Blockchains offer a democratized infrastructure that provides the basis upon which to build new business models. An immutable distributed ledger captures the state of the system and is transparently shared with all participants.

Smart contracts enable the execution of alignments in a way that does not require the involvement of a third party to enforce the contract in case of later disagreements. Digital currencies enable payment for services in connection with smart contracts so that additional manual processes can be reduced. The use of digital currencies as part of the blockchain also resolves the last barrier to complete automation, enabling not only abstract processes to be modeled but also real value transfers to be made. With a market capitalization of \$3 trillion in 2021 (Ossinger, 2021) blockchain tokens have gained some relevance. As the ledger is immutable, participants can trust that input data will not later change. This enables participants to learn about other users' transactions on the ledger, which can be used to form opinions of the fit of potential business partners.

The concept of the blockchain seems to address the above mentioned problems. However, can it be used to improve the situation of ELMs? What are the actual requirements of ELM users, and how would these requirements change if the ELM were on a blockchain?

An examination of the current literature offers no answers to these questions. Given that blockchains might improve the value of ELMs for users, how would one construct such an ELM? What kind of design theory would enable the construction of different kinds of ELMs? What existing knowledge can be used to design such a theory?

Blockchain technology research is in its early stages (Wang et al., 2021c). Therefore, several unresolved areas influence the success of an ELM. All data are shared with all participants in a public blockchain, which can lead to privacy problems. Companies do not want to share their confidential data publicly. Also, laws such as the General Data Protection Regulation (GDPR) restrict what can be shared and when the sharing or persistence of data must cease. In opposition to this, blockchains should enable transparency. Therefore, mechanisms for securely sharing data are needed on public blockchains. Public blockchains are also not scalable enough to handle the workload of the logistics industry, which involves millions of shipments per day (see Chapter 3). One reason for this is that in storing data, they are always replicated many times to all full nodes. A common practice is to store only representations[1] of data on the blockchain and to store the actual data on a different system. Even this solution might not be adequate to represent the shipments handled in Europe (see Chapter 6). Sidechains or second layer systems are used to address such problem classes. However, the properties of such systems and the protocol that represents the business flow are not known. The use of smart contracts usually means that computations must be conducted several times on different nodes to ensure the integrity of the chain. In addition, a protocol is needed to balance on-chain and off-chain computations.

Energy consumption in public blockchains is high because many nodes must work continuously to ensure that the immutable state remains immutable. Some methods, such as proof of work, are also energy-intensive. Reducing CO_2 emissions on the road but producing them in data centers might not be beneficial.

Consortium blockchains are intended to address the above issues. They are geared more toward enterprise users needs. They allow strict control of data sharing. Additionally, they scale well because trust in the validation of the blockchains' consistency builds trust toward its operators. Energy consumption is much lower compared to traditional public blockchains because the networks are smaller and expensive mechanisms such as proof of work are not needed. While public blockchains act as infrastructure, and their users can rely on the availability of this infrastructure, users of consortium blockchains must operate them on their own. The setup of such blockchains can be highly complex because of the management needed to organize their operation. In addition to running a complex set of blockchain server processes that enterprise operations teams usually do not support, users must coordinate the deployment and maintenance of new smart contracts not only within the data center but also within the consortium. Would the benefits of this model

---

[1] Mainly hash values of the actual data are used.

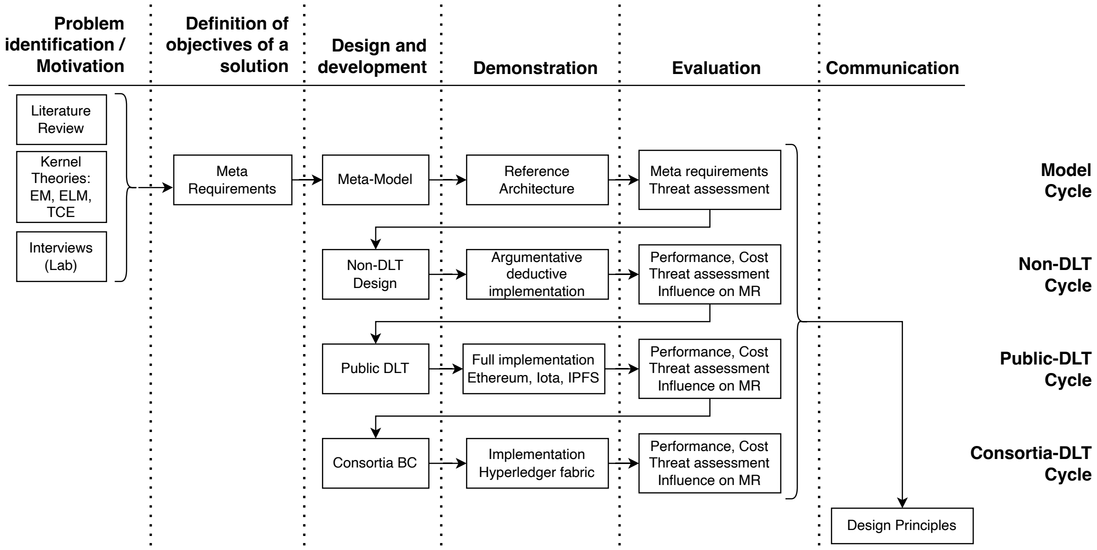

Figure 1.2: DSR cycle according to Peffers et al. (2007).

outweigh its drawbacks, which tend toward centralized management?

1.2 Target definition and research methodology

The goal of this research was to determine whether DLT would improve the usage of ELMs. If so, how could an ELM be constructed to achieve the highest proportion of benefits? The main research question, therefore, is as follows:

Research Question. *How do Distributed Ledger Technologies contribute to the design of Electronic Logistics Marketplaces?*

To answer the research question, a multidisciplinary study must be conducted. First, the logistical problem must be clearly defined and understood, and then the field of computer science used to formulate and evaluate possible solutions in the area of DLT. Information systems research connects both directions (Avison and Elliot, 2006) and is thus where the present research is positioned.

The goal of this research was to build an artificial (Simon, 1996) artifact, a meta-model, that would serve as a blueprint for concrete ELM models. The dissertation follows the design science research (DSR) methodology (Gregor and Hevner, 2013; Hevner et al., 2004; Peffers et al., 2007) because it aimed to create a novel artifact to solve a problem in an innovative way (Hevner et al., 2004), and DSR provides a rigorous process to create this artifact. This process is an iterative approach that structures the search for the research outcome through several phases (see Figure 1.2) to find an answer to the research question. Although the meta-model is presented in a single chapter of this dissertation, it was designed and refined through four DSR cycles. Every instantiation and evaluation of the meta-model refined its design. The target audiences are researchers and practitioners.

The first iteration started with motivating and identifying the problem. To better understand the problem, a literature review of the knowledge base around ELMs was conducted so that existing theories around electronic marketplaces (EM) that enabled the research into ELMs could be used. The review of existing research also uncovered the current inhibitors of the usage of ELMs, and it also showed which enablers improve their usage. These inhibitors and enablers would serve as meta-requirements (Walls et al., 1992) against which the later designs could be tested. Additionally, transaction cost theory served as a kernel theory (Walls et al., 1992), providing justificatory knowledge (Gregor et al., 2020) to explain "why the design works" (Hevner and Chatterjee, 2010).

In addition, a study with experts in logistics and blockchains was conducted to evaluate the current gaps in ELMs, as well as how an ELM business model could leverage the distinctive properties of blockchains to produce an ELM that better fits the needs of its users. Interviews were conducted to gain additional insights into current problems.

The solution space is, on the one hand, enhanced by the features of DLT. These features are introduced and evaluated regarding their influence on the inhibitors and enablers of the ELMs. On the other hand, the solution space is narrowed by the problems associated with DLT.

The meta-requirements, kernel theories, and structured knowledge about the solution space guided the design of the meta-model and provided it with the proper scientific grounding (Gregor and Hevner, 2013). Although blockchains were not mentioned in the meta-model, the model still had to prove that it could fulfill the stated meta-requirements and work in DLT. In the case of an abstract meta-model, the fulfillment of the requirements could be proven only analytically (Hevner et al., 2004). As the security of blockchain solutions and ELMs is important, threat modeling was done to ensure a secure solution.

However, showing only analytical results for a prescriptive model in a dissertation that aims to solve business problems would not prove the validity of the approach. Consequently, the model evaluation was extended to include implementations in blockchains. Only after a rigorous evaluation in that setting could the model gain some generality for the class of ELMs. Three additional DSR cycles that do not contain the *Motivation* and *Definition of objectives* phases were conducted. As identified in the previous section, there are two dominant classes of blockchains. Hence, two different implementations were conducted. Additionally, whether the same benefits could be achieved if the meta-model were used in a non-DLT model implementation was verified.

First, the non-DLT design was tested. As methods of software engineering are well studied, no interesting knowledge would be gained by conducting an actual implementation in this case. Therefore, the implementation was done only "on paper" in an argumentative way, as this was sufficient for the evaluation of the design and its comparison with the other solutions.

Second, the validity of the model was demonstrated with public blockchains. A prototype was developed that uses the model to overcome the weaknesses of public blockchains while still conforming to the requirements. This prototype was intended to be complete in the sense that a user could potentially use it as an ELM. The evaluation of this prototype aimed to show that it is possible to scale an ELM while keeping the data secure. Implementing and testing the whole process also helped to identify shortcomings in the meta-model.

Third, the validity of the model was shown on consortium blockchains. The expectation was that the implementation would be much easier because many problems, such as scalability and privacy, seem to be solved per definition. The problems in this domain, such as centrality and complexity, must be examined carefully, as should the data secrecy topic, as the initial hypotheses might not hold true.

All three instantiations were evaluated for performance, cost, and threats and their influence on the meta-requirements.

Finally, the learnings from the DSR process were summarized into design principles to guide future software architects in designing ELMs. As the model was rigorously tested and validated, it can be used as a prescriptive blueprint for practitioners. The body of knowledge should also be extended by practical applications of blockchain technology in the area of ELMs. Further, this research should contribute to the research on scalable and secure data exchange in blockchains.

1.3 Dissertation structure

The structure of this dissertation is as follows (see also Figure 1.3):

Chapter 2 defines the ELM background and presents the findings regarding the enablers and inhibitors of ELM usage.

Chapter 3 presents the DLT background and discusses the influence of DLT on the meta-requirements identified in Chapter 2.

Chapter 4 shows the results of the workgroup activities in the lab and of the interviews conducted.

Chapter 5 defines the meta-model, introduces a reference architecture, and evaluates the meta-model.

Chapter 6 establishes the three instantiations (non-DLT, public DLT, and consortium DLT) and evaluates them.

Chapter 7 introduces the design principles that were identified during the DSR cycles.

Chapter 8 compiles and discusses the findings of this work, shows its limits, and presents possibilities for research in this area.

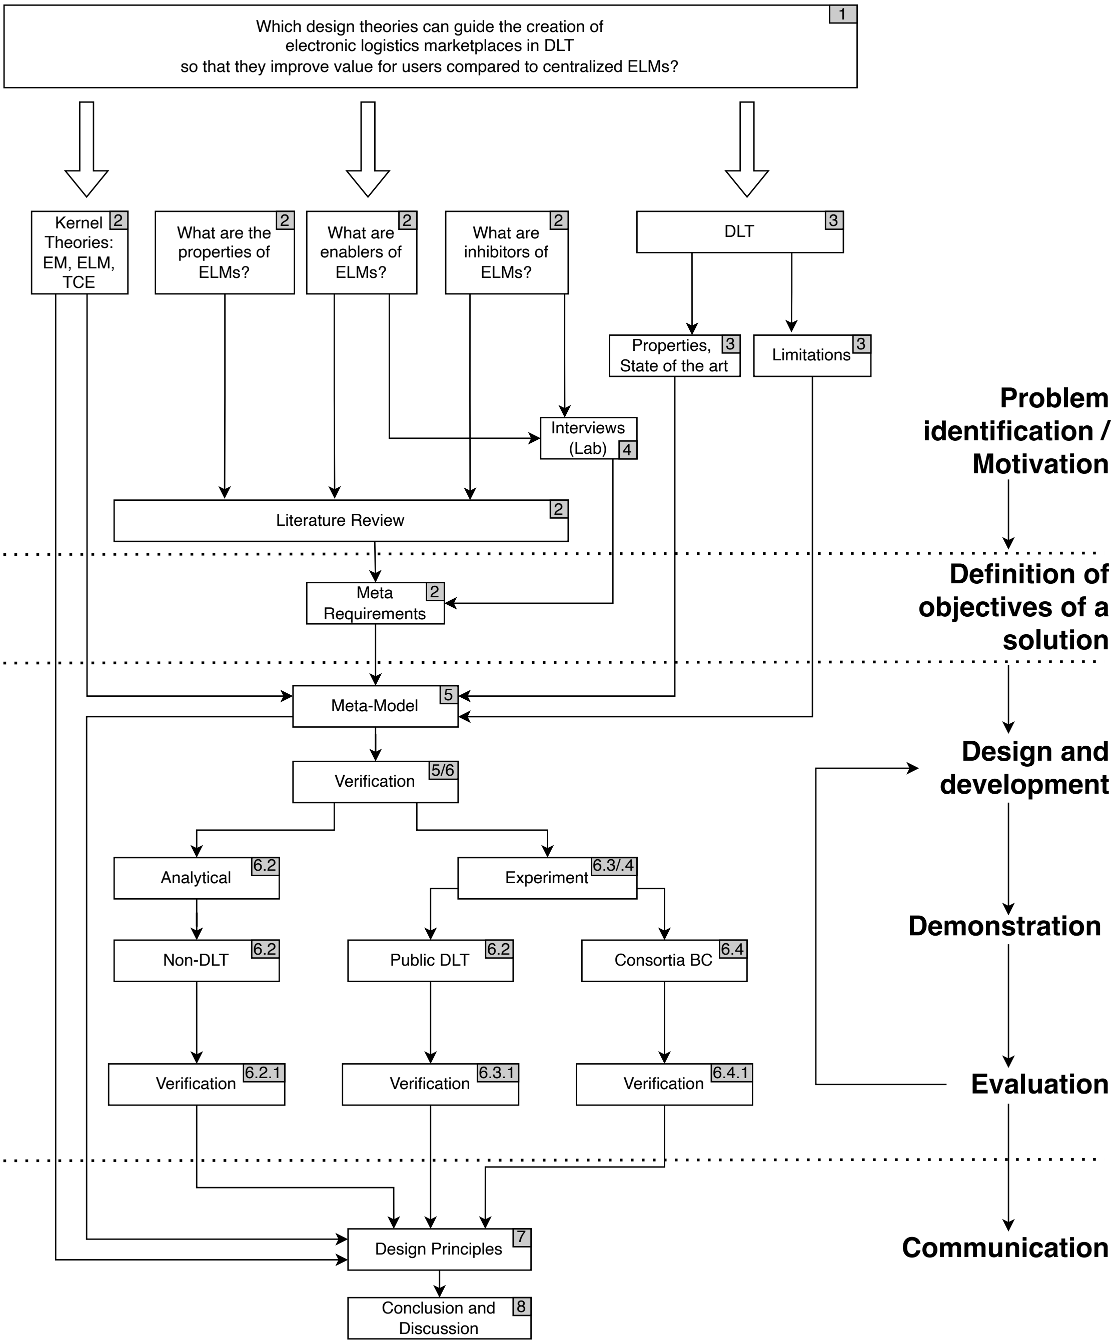

Figure 1.3: Methodological research approach showing the dissertation structure and highlighting its main parts, with the chapter/section numbers in the corners.