


Florian Kleedorfer, Christian Huemer

Towards a Web Based Transportation Infrastructure



Towards a Web Based Transportation Infrastructure

Florian Kleedorfer¹, Christian Huemer²

1 – Research Studios Austria

2 – Technical University of Vienna

The level of digitalization within transport companies is much higher than the level of digitalization across organization boundaries. This fact suggests that there is room for improvement. However, this situation is not likely to change as long as there is no financial incentive for the whole sector to cooperate in establishing a shared communication infrastructure. In this paper, we present our approach for building such an infrastructure using the method of design science. The goal is an open, Web based, de-centralized network operated by transport organizations themselves. Based on expert interviews, we argue that the current situation causes frictions that our approach may help reduce, thereby providing the incentive to participate. The proposed system is described in terms of its existing technological base, the Web of Needs, and the extensions needed to provide the required functionality, giving an overview of the current state of implementation.

Keywords: linked data; transport; de-centralization; digitalization

1 Introduction

Companies in the transportation domain have developed their IT infrastructure in recent years, giving clients ever easier access to the services they provide. Mobile apps and Websites allow users to create accounts and comfortably book transports, make pick-up/delivery appointments, and pay their invoices. These developments reduce the transaction cost of becoming and onboarding a new client.

However, this style of digitalization is not without drawbacks. For one, there are considerable costs associated with building and maintaining these tools. Because of that, they are primarily designed to meet the requirements of the transporter, not its clients. They tend to lock users in a streamlined one-stop shop for transportation, whereas the opposite would be in the clients' interest: to be able to compare service offerings and prices across vendors, and to switch between vendors easily. Instead, clients are confronted with a plethora of communication channels, apps and on-line accounts each of which lets them communicate with just one vendor.

Recent economic successes of *matchmaking platforms* like Uber or Airbnb (*Uber* 2017; *Airbnb* 2017) have proved the need for a vendor-independent medium where supply and demand can meet. Such a medium for the management of transactions between market participants obviously serves as a unifying force leading to standardized interfaces between the participants, lowering the cost of finding a partner and enabling trust between newfound partners.

While a remedy to some issues of the more traditional style of digitalization, this approach has drawbacks, too. Most prominently, matchmaking platforms tend to create winner-take-all (or winner-take-most) situations: For any given product or service, it makes more sense for a participant to join the platform that already hosts more of the clients the participant interested in. The result is a small number of dominant platforms that have a de-facto monopoly.

The intrinsic logic of the platform business model is to allow transactions only within the platform, and to allow only the transactions that are directly or indirectly economically profitable for the platform. While the platform model thus allows participants to reach a bigger pool of potential partners than the traditional model (each participant operates their own infrastructure), it introduces new limits on that pool out of the platform's self-interest: only members of the platform are meant to interact.

Our goal is the development of a de-centralized, open, standards-based infrastructure for transportation, providing the adequate technology for all market participants without the limitations induced by the necessary self-interest of the centralized platform model. The vision is to interconnect the existing IT systems of transport companies on a protocol level instead of connecting each of them to one central hub. This approach allows for diverse, inter-operable software products that can be tailored to each participant's needs; all market participants can operate services of their own and thereby retain ownership of their data and they are free to engage in interactions with any other participant that implements the protocol.

In this paper, we report findings from expert interviews that support the relevance of the problem we are addressing and give us a direction for further development. We explain how we harness prior work and show which adaptations and extensions have to be made in order to solve the problem, and finally, we describe the current state of design and implementation.

2 Research Method

In our work, we follow the method of design science in information systems research, proposed by Hevner (Hevner et al., 2004). According to this framework, design science is primarily focused on the creation of technological artifacts. The research topics are obtained from the environment of the artifact in the form of relevant business needs. When finished, artifacts have to be applied in the environment in order to assess their effects. This cyclic process is called the *relevance cycle*. Using the available knowledge base, artifacts are iteratively adapted and evaluated using rigorous scientific methods. Theoretical insights are added to the knowledge base through scientific publication, completing the *rigor cycle* (Hevner, 2007).

In the following, we describe our research process as an application of this method. The first step is the assessment of problem relevance and the elicitation of business needs, which we do by interviewing domain experts. As a second step we propose a design for technological artifacts to address the business needs we found.

3 Relevance: Qualitative Interviews

The problem we postulate is that there is considerable friction between market participants in the domain of transportation in its current state of digitalization. Our main hypothesis is that a unified open infrastructure that allows for matching suitable business partners with each other can reduce friction.

3.1 Sample Selection and Data Collection

In order to assess the problem relevance and to find out which frictions are deemed important by domain experts, we conducted qualitative expert interviews with individuals who are active in the field of transportation (DiCicco-Bloom and Crabtree, 2006). The interview partners were initially identified by Web search for keywords mainly from the domain of courier and parcel services in Austria. Interviews were conducted with those who showed interest in cooperating. The range then was extended by a snowball technique (Varvasovszky and Brugh, 2000). In the course of this analysis, the focus was widened to include the forwarding sector as initial results pointed toward greater than expected problem relevance in that area. A number of interviews were conducted at an agricultural trade fair, which explains the influence of that sector in the results.

In total, 27 interviews were conducted with a prepared interview plan that evolved a little over time; five interviews could be completed according to the plan, the others had to be shortened due to limited availability of the participants. The interview partners are active in the roles of trader(10), transporter(5), logistics(2), broker(2), or others (8). Data collection was done by note-taking. The notes were subsequently analyzed for commonalities that caught our interest.

3.2 Findings

In the following we highlight our findings, noting the number of participants whose interviews support each one. For each finding, we argue how it influences our assessment of problem relevance. While some findings will directly relate to functionality currently under development, some findings can at this point only serve as motivations for functionality that may be provided by our infrastructure in the future.

1. **Trust.** Especially when freight is high-volume or high-value, established contacts between consignor and transporter are highly valued. Trusting new players is difficult and not deemed necessary once good business relationships have been set up. This finding disconfirms the relevance of our work because players prefer existing business relationships and do not need matching with new partners. On the other hand, it would motivate a feature that enables trust in new partners, and of one that allows representing and leveraging one's business network. Support: 9.
2. **Digitalized communication.** Transportation is digitalized only within companies, and even within companies, the actual execution is often not digitalized (communication with driver or for confirming handovers). If it is digitalized across companies (other than via e-mail or messaging apps), it is set up in such a way that the client company integrates the API or uses the user interface of the transport company. The general view of market participants was that the main means of communication is through phone calls, using E-mail for written confirmation. Companies do have systems using barcodes, matrix codes or electronic identification tokens to trace their consignments, but they do not seem to work across companies. This finding supports problem relevance because it shows there is room for automation of communication across companies. Support: 4.
3. **Documents.** In some constellations, consignors and consignees demand the original paper documents from the consignor be passed on to the consignee, incurring additional complexity. One such document is the CMR¹ consignment note (UNECE, 1956) that needs to be signed by the consignee and find its way back to the consignor - failing to obtain the signed CMR consignment note can incur considerable costs. The consignee, on the other hand, has little to no incentive to send it back, which can cause serious problems.² Together with Finding 2 (digitalized communication), this finding is a strong argument for further technological harmonization of standards for transport documents and their seamless integration in everyday business. Support: 4.

¹CMR is short for the United Nations' *Convention on the Contract for the International Carriage of Goods by Road*

²Overall, delivery paperwork is being replaced by electronic data interchange. Concerning the CMR document it should be noted that at the time of this writing, the digital e-CMR Waybill is available in a number of European countries. (UNECE, 2008). It was not mentioned in any of our interviews, though.

4. **Privacy.** Keeping trade secrets is an important issue for players in the transportation domain. Several experts voiced concerns that their contracts and bids may be inspected by other parties in an open system, and therefore took quite a reserved stance toward our work. Multiple experts stated that if requests for quotes and their bids could be revealed to a selected audience only, it would be a system they would welcome. Support: 4.
5. **Appointments.** Depending on the means of transportation and the situation of the consignee, the complexity of making appointments for pick-up and delivery as well as the consequences of errors can vary greatly. The main medium for making appointments and informing about changes is telephone. Large logistics companies have dedicated communication channels (apps or automated SMS). This finding generally supports the relevance of our work if the developed technology supports the integration of processes and resources for coordinating appointments so as to standardize the situation, removing the burden of last-minute improvisation and providing a reliable record of the events that led to a problem. Support: 3.
6. **Calls.** Being able to talk to their counterpart in addition to written communication is described as important by some experts. It allows for quick problem solving, helps building trust quickly (video calls or meetings in person even more so), and it is a good way of maintaining the customer relationship. Ideally, it would be possible to reach the relevant person (e.g. call the truck driver to inform them of an important change). This finding provides support for an integration of voice or video chat in a technology managing transport relationships. One broker noted that a conference call with multiple participants would be ideal. Support: 3.
7. **Certificates.** Some experts stated that friction is caused by having to prove that all parties in the logistics chain hold a certain certificate, such as the GMP+ certificate for sustainable production. This finding supports problem relevance and may warrant the use of blockchains. Support: 3.
8. **Framework agreements.** High-frequency consignors do not set up each transport assignment individually. Rather, they select transportation partners, often on a yearly basis, set up a framework contract, and then execute many transports with easily calculated rates and quick planning. This finding seems to disconfirm the problem relevance as a system that matches transportation partners for each individual transport is not

needed for this situation. However, if the solution allows for finding transportation partners for framework agreements, and setting up individual transports based on a framework agreement, possibly in a private call for proposals among partners with whom a framework contract has been set up, the described situation might even be improved. Support: 3.

9. **Return freight.** For a one-off transport, getting a transportation quote can take a long time (one week is not uncommon for agricultural goods) because of the need to find a freight for the return journey. This finding supports problem relevance as the delay causes friction in the information or negotiation phase that might be reduced by a more efficient process of finding a transport request for the return journey. A solution feature that could be especially useful here would be the option to offer a transport only on the condition that a corresponding return contract is found, and to represent and monitor the condition satisfaction automatically. Support: 3.
10. **Price sensitivity.** Multiple experts stated that transportation is very price-sensitive - price is the main decision factor for clients. Quality of service (keep appointments, response speed and accuracy, no damages) is only a negative decision factor in the case that it does not meet the expected standard. A matchmaking technology like the one we are proposing clearly addresses a market need of clients but it does so at the expense of transporters, who seem to feel threatened by our vision. Support: 3.
11. **Price negotiation.** Some experts stated that price negotiation is a cornerstone of the business. A system that does not allow for negotiating is not acceptable. This finding does not affect our assessment of problem relevance but points toward a negotiation feature. Support: 2.
12. **Quality dimensions.** Some experts stated they would not want to compete world-wide with everyone because of price sensitivity, but they felt they had a higher-quality offering than cheaper alternatives. It would be necessary for open competition to be able to advertise verifiable quality dimensions such that clients understand the higher price. This finding shows that the technology should provide the possibility to differentiate offerings from each other in multiple dimensions and to make verifiable claims about service/offer quality. The deeper problem this finding touches upon is that apart from price, clients often do not know about aspects of quality they require, either because these requirements are implicit or because clients lack experience. Therefore, in addition to

providing a means for verifiable claims, it would seem beneficial if the technology provided a means for conducting a conversation with the client about service qualities with the goal of educating the user about alternatives so as to allow them to make a more informed decision. Support: 2.

13. **Wrong Assumptions.** In some cases, important information is present only in an implicit form. For example, the recipient of a shipment of agricultural products may only be able to unload a dumper truck, not any other kind of truck - but this fact is assumed to be known to the transporter, which may lead to the wrong kind of truck being used, incurring delays and additional cost upon delivery. This finding provides support for machine-interpretable, standards-based, automatic information interchange that may uncover incompatibilities like the one above automatically. Support: 1.
14. **Freight exchanges.** Transporters are reluctant to use freight exchanges to find new assignments because the prices they can ask there are very low, to the point that it is cheaper to make an empty return journey than to obtain a contract there. This finding may be seen to disconfirm relevance of our work, as the technology could have a similar effect for transporters overall. It is possible, however, that this effect is related to systems that focus mostly or solely on price and that in this case, the solution features discussed in connection with Finding 12 (quality dimensions) could ease the situation for transporters. Support: 1.
15. **Contract negotiation.** Contract negotiation may not be necessary, but when it is, it can be a lengthy and error-prone phase. In some cases, one side simply dictates the contract's conditions. This is the case when a transport partner is being searched for a longer term business relationship, and transporters make bids, or when a transport company offers standardized services together with standardized terms. In some situations, contracts have to be set up for individual transport assignments, which is usually done using MS Word documents sent back and forth via e-mail. At least one of our interview partners expressed the hope for a better integrated solution that allows for more clarity during this phase. Support: 1.

We judge Findings 2, 3, 5, 7, 9, and 13 as supporting problem relevance. Findings 1, 4, 6, and 8 do not clearly refer to problems calling for a solution. Rather, they describe current strategies evolved to avoid friction as much as possible.

Insofar, we would not say they contradict problem relevance, but we cannot include them in the set of supporting findings, either. Findings 10 and 14 disconfirm problem relevance from the point of view of the transporters because the solution may mean more pressure on prices. Findings 11, 12, and 15 are interpreted as being neutral with respect to problem relevance.

In total, the interviews do seem to indicate the problem we are tackling is relevant. This assessment is backed up by the fact that twelve experts said they would be happy to try out our system and six welcomed our our plans explicitly. We were able to collect information about expectations and requirements that may lead to future developments, most of which need further refinement. In our design process, we focus on fundamental functionalities first, which are generally taken for granted in such interviews and therefore do not come up as important topics. Therefore, the full consequences of the findings on concrete design decisions may materialize at a later point in the process.

In the following, we present our approach to the technical design of the solution. Any influence of our interview findings will be noted explicitly.

4 Knowledge Base

4.1 Web of Needs

The main building block for our solution is the Web of Needs (WoN) technology (Kleedorfer, Busch, Pichler, et al., 2014). Its overall functionality is depicted in Figure 1. Participants publish supply and demand on a de-centralized network (1a, 1b), on servers (*WoN nodes*) of their choice, in a standardized form such that they are machine-readable and can be found by dedicated, independent *matching services* (2) that identify suitable (supply,demand) pairs and inform them by sending a *hint* message (3). The participants controlling the supply and demand objects can choose to establish a communication channel and exchange messages.

WoN is domain-independent, focusing on providing a framework for common interaction patterns while leaving space for creating domain-specific specializations. In order to achieve this goal without introducing a break between the framework and the content, WoN is entirely based on RDF (Manola and Miller, 2004) as a data description language and the Web as its basic framework: the whole data

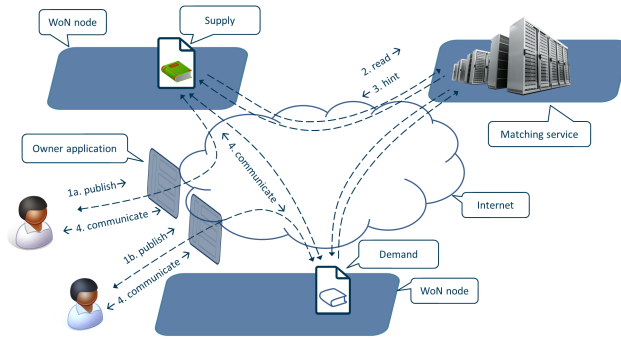


Figure 1: Interaction diagram showing interactions of two participants and a matching service in Web of Needs. Image is based on Kleedorfer, Busch, Huemer, et al., 2016, licensed under CC-BY-SA 4.0.

structure, comprising supply, demand, their connections and the complete data exchange are represented as linked data (i.e., RDF that can be accessed on the Web, Bizer, Heath, and Berners-Lee, 2009). The participant's privacy is protected by representing needs (supply or demand), and not users, in the system; each need can have a dedicated cryptographic key pair which is used to verify any message it sends. The integrity of the communication history is automatically ensured by iteratively signing a hash derived from the past message history (Kleedorfer, Panchenko, et al., 2016). This approach allows for using the communication channel between two participants as a shared RDF database to which they can both add data by sending messages, but not delete or otherwise change past data. By representing the content domain adequately in RDF, this system allows for cooperatively creating and manipulating a shared machine-interpretable model of any business transaction.

The immutability property of WoN conversations makes it look similar to blockchain systems at first glance (Nakamoto, 2008). The immutability property is not quite as reliable in WoN as it is in a popular blockchain: signatures are only made by participants of a conversation, as well as by the WoN nodes they use, so threat model and trust model differ greatly between blockchain and WoN. It remains to be seen if trust and threat models in WoN are acceptable for practical applications;

it should be noted, however, that WoN provides functionality on a different level than blockchains: expressing supply and demand, matching, and messaging are not concerns of blockchain systems.

The reason for WoN to include a cryptographically assured message history is the observation that in practice, the conversation that precedes a transaction must be interpreted as the contract both parties agree to in the absence of any more formally defined agreement. It is therefore important that they can rely on an unchangeable message history in order to prove what has been agreed to in the event of a dispute. In order to achieve this, it is sufficient to have each others' signatures on each message.

4.2 i-Cargo Ontologies

The EU project i-Cargo (Hofman et al., 2016; A. Garcia, 2015), finished in 2015, is a highly relevant research endeavour we can build upon. One of the goals of i-Cargo was the creation of an open freight management ecosystem spanning multiple organizations and countries. The project produced a number of artifacts that can be re-used for our work, most importantly a collection of ontologies specifically developed for the transport sector (Daniele and Ferreira Pires, 2013). These ontologies are the Logistics Core Ontology *LogiCo* (Daniele, 2013a), the Logistics Services Ontology *LogiServ* (Daniele, 2013b), and the Transport ontology (Daniele, 2013c). They provide a basis for describing the entities relevant for the transport domain using RDF. Using these ontologies, it is possible to describe such entities as consignor and consignee, means of transport, consignments and transported goods, packaging, delivery and pick-up options, transportation requests and transport execution plans.

5 Artifact Design: Open Logistics Networks

The application and extension of WoN to support transport and logistics creates a virtual medium in which any number of real-world market participants can be represented in self-organized open networks, hence the name of our project, *Open Logistics Networks*.

The general idea is to describe the relevant entities according to the *LogiCo*, *LogiServ* and *Transport* ontologies (or possibly according to simplified versions

thereof), to place them inside the needs according to the Web of Needs approach, and to publish those needs on the Web.

Figure 2 shows a transport-related interaction diagram to illustrate this approach. Matching services can use the domain specific descriptions for accurate matching and allow the needs to establish an RDF-based communication channel. Each such connection naturally gives rise to an RDF model that both parties have access to via HTTP. It consists of the set of RDF triples defining the needs and the triples exchanged in the communication channel. If additional resources are referenced in any of these triples, they are added to the shared model as well.

For example, consider an example motivated by Finding 13 (wrong assumptions), with one need describing the delivery event for a consignment of oil seed at an oil mill, seen from the side of the oil mill. It references an RDF description of the technical capabilities for unloading. The connected need describes the delivery event as seen from the transporter, linking to an RDF description of the vehicle, among other details. Both needs have access to the union of all these triples, all of which have traceable provenance information. If information is missing or needs to be changed, it can be requested and provided on the communication channel in the form of RDF triples. Any state change in the shared transaction (e.g., a change to the expected time of arrival) is also represented as new triples that are added to the channel. An actor organizing a bigger part of the transport chain controls not just one, but a number of such needs, which are naturally combined in one RDF model that just covers a larger part of the transport transaction chain than the models accessible to each of the other partners.

In the remainder of this chapter, we describe adaptations required for applying WoN in transportation use cases. The interested reader will notice that none of these topics are uniquely specific to transportation. They must, however, be dealt with before we can tackle domain-specific functionality that addresses problems revealed in our interviews.

5.1 Service Descriptions

The basic element of the Web of Needs are entities representing concrete interest in a transaction, for example, someone may offer a book for sale, and another user wants to obtain a book. Both users could express their intention in the form of a *need*. Both users would get a hint message from a matching service and engage in




Figure 2: Interaction diagram showing transport specific interactions of consignor, consignee, and transporter in the Web of Needs. In the depicted situation, consignor and consignee have agreed to send the consignment (the book), and publish needs representing the transport and the part they each take in it. The transporter controls needs representing pick-up and delivery. The diagram shows how messages are routed between transporter, consignor and consignee. Image by Kleedorfer, Busch, Huemer, et al., 2016 is licensed under CC-BY-SA 4.0.

a conversation. Both sides use the need object to maintain a reference to that specific business transaction. In the case of general service offerings, however, this pattern is not sufficient. A company offering transport services in a given region could publish an object in WoN, receive hints and establish connections with its clients using that object. This would lead to this object collecting more and more connections over time, possibly impacting data management and application performance. Therefore, we decided to enable a new interaction pattern: a need that creates new needs programmatically for each specific case. In analogy to the concept of a *factory* in software development, an object that creates other objects, we call these special needs *factory needs*.

General service offerings can be represented as factory needs that are informed of any new, potential transaction partner. For example, if the service to be offered is a courier service, the factory need is informed whenever there is a potential client who wants a package delivered. The client, however, should not be informed of just any existing service announcement. They already stated their need and expect to be contacted with concrete offers, not links to services that they might want to explore. Upon being informed of the existence of a potential client, the logic behind the factory need is activated and makes a concrete offering (possibly taking into account situational factors such as current traffic, concurrent requests, or available resources). This concrete offering is represented by the newly created need object, which initiates a connection with the client's need.

Such a scenario requires different matching logic than the symmetric case that had been possible in WoN, a case in which both needs in a match are notified of the match. In the case of factory needs, only the factory should be notified. This requirement caused us to extend the way matching works in WoN such that by default, both parties are notified, but a party can choose to suppress notification for the counterpart as well as the notification for itself. This change was made to the prototypical implementation and tested functionally. Though not a full solution, the fact that one can now define that others should not be notified of one's needs, may contribute to easing the issue raised in Finding 4 (Privacy).

5.2 Information Requirements

As described in Finding 13 (wrong assumptions), the information used for matching is not necessarily identical with the information needed to execute a transaction. In some cases, all the needed information is present when the match occurs.

In other cases, additional questions need to be asked. The fact that the whole data exchanged between participants, including their initial need descriptions is available to both sides as RDF, it is possible to query this data in order to see if the required information has already been stated or not. Based on that property of the system, it is possible for a participant to define information requirements in a declarative way. When a connection is made, the information requirements are automatically checked and missing information is identified. User interfaces can use this information to generate interface components (e.g., forms) to elicit the missing information. User agents may be able to fill in missing information automatically from personal data stores.

In order to realize a system that allows participants to define their information requirements, we need to be able to formulate and check them automatically. We plan to use SHACL (Knublauch and Kontokostas, 2017) for this aspect of WoN. SHACL is a formalism for defining *shapes* that an RDF graph can be validated against, the shapes themselves are written in RDF. This allows a need author to include SHACL in the need description, and any party with access to the conversation content can evaluate the content against the shapes. Transmitting the information requirements (shapes) in declarative syntax instead showing a service-defined user interface (e.g. a HTML form) allows for great flexibility; user agents can choose the user interface technology for eliciting the required information from the user. Moreover, SHACL allows for referencing shapes from any Web location, which allows for re-use of shapes and thus an avenue for standardization of distinct use cases, as well as the evolution of such standards. A complete prototypical implementation is to follow.

5.3 Message Retraction

The functionality of information requirements has to operate on a higher level than the level of simple message exchange - one might call it the level of meaning exchange. Each new message contributes its content (RDF triples, grouped in named RDF graphs) to the meaning of the conversation. Sometimes, participants need to change the meaning, for example, to correct a mistake. Therefore, the communication protocol allows for a *retraction message* that links to a message already contained in the conversation, which is thereby marked as retracted. Each participant can only retract their own messages that were sent earlier. Retract messages themselves cannot be retracted. The effect of retraction is that the payload of each retracted message is disregarded in further operations. This planned

feature would contribute to functionality that expects users to make and correct mistakes, such as the issues raised in Findings 2 (digitalized communication), 15 (contract negotiation), 5 (appointments), and 11 (price negotiation).

5.4 Distributed Transactions and Long-Running Work

Especially in the context of multi-modal transport, it is required to coordinate more than two parties in order successfully to plan and execute the transport. In addition to multiple parties being involved, such an endeavour can span a time of weeks or even months from initial planning to the final delivery. Both aspects make the concept of long-running work (Bocchi, Laneve, and Zavattaro, 2003) suitable for representing such processes. We decided to use the protocol logic defined by WS-BusinessActivity (WS-BA) to model interconnected and possibly nested activities of multiple business participants. As part of such a plan, a need assumes the role of the WS-BA coordinator vis-a-vis the needs that represent possible solutions for the problem. Vis-a-vis a higher-level aggregation node in this plan, that need assumes the role of a WS-BA participant.

In order to realize such a need in WoN, in either role specified by WS-BA, it is possible to use server-side logic that is invoked for an incoming message after the mandatory logic realizing the basic protocol has been executed. The need thus implements the state machine prescribed by the WS-BusinessActivity protocol specification.

This system of states and state-changing messages is shared between participants and can at all times be used to determine and influence the state of the distributed work. The option to add sub-goals to an overall plan may be one of the building blocks for an integrated appointment making sub-protocol, addressing the situation mentioned in Finding 5 (appointments), and it could be used to realize linked conditional needs that find a one way freight and the return freight simultaneously as required by Finding 9 (return freight). There is an experimental implementation of this feature that will be overhauled.

5.5 Expressing Agreement

In order for transactions or long-running work to function legally, it is required that the two parties involved set up a contract. Theoretically, a spoken agreement

is enough to fulfil this requirement; in practice, though, contracts are written and their contents are an important phase during negotiation. Finding 15 (contract negotiation) hints toward the fact that a unified method of contract negotiation may be beneficial for all participants.

The communication channel established by our technology between the participants retains all messages exchanged in a manner that makes the message history unambiguous and unchangeable. It is thus a suitable medium for contract negotiation as it can always be established who said what in which order.

To allow for negotiation, we introduce the possibility to send a special message that marks earlier messages as *proposed*. The other participant has the option to send an *accept* message, referencing the proposal, thereby creating an agreement. The agreement is identifiable by the URI of the accept message. For calculating the agreement's content, only the message history before the accept message is considered. Consequently, retraction has no effect on agreements. They can be canceled by either side by sending a message that *proposes to cancel* an agreement and the counterpart accepting that proposal. At time of this writing, the protocol has been designed, but not yet implemented.

6 Conclusion and Future Work

In this work we describe the current state of our iterative design process with the overall goal of a de-centralized, open, Web based transportation infrastructure. We use expert interviews to identify business needs and assess problem relevance. Addressing these needs, we propose a design for the most fundamental functionalities needed to achieve our overall goal. Two aspects relevant for practical application have been left aside so far, namely payment information integration and reputation management. We identified both aspects as relevant at the outset of the project; our interviews arguably support these features as well: Findings 1 (trust) and 12 (quality dimensions) quite directly point toward reputation management, and a number of business interactions (e.g. high-frequency transport described in Finding 8 (framework agreements)) may be greatly improved if payment was integrated in the communication instead of being handled in an independent channel.

In line with the iterative approach, we are striving for applying the technology for simple transport problems at first, while gathering requirements for more complex

situations and gradually implementing solutions for them. Practical evaluation of the designed artifacts in real-world applications and follow-up interviews and workshops are planned for later development phases.

Acknowledgements

This work would not have been possible if it had not been for invaluable ideas and criticism contributed by the research team at Studio SAT, which currently includes Heiko Friedrich, Fabian Suda, Kevin Singer, and Soheil Human. We are indebted to Laura Daniele for sharing her knowledge of the i-Cargo ontologies and to the domain experts who provided valuable insights during our interviews. Special thanks go to Dominik Kovacs of DoKo Agromarketing for introducing us to the community of agricultural commerce. We are also grateful for all contributions to the discussion on the topic of negotiation in the semantic Web mailing list.³

Financial Disclosure

This work is funded by the Austrian Research Promotion Agency (FFG) under the project *Open Logistics Networks (OLN)* of the Program *COIN - Cooperation and Innovation*.

References

- Airbnb* (2017). URL: <https://www.airbnb.com/>.
- Bizer, C., T. Heath, and T. Berners-Lee (2009). "Linked data-the story so far". In: *International Journal on Semantic Web and Information Systems (IJSWIS)* 5.3, pp. 1–22.
- Bocchi, L., C. Laneve, and G. Zavattaro (2003). "A calculus for long-running transactions". In: *Formal Methods for Open Object-Based Distributed Systems*. Springer, pp. 124–138.
- Daniele, L. (2013a). *Logistics Core Ontology*. TNO. URL: <http://ontology.tno.nl/logico/>.
- Daniele, L. (2013b). *Logistics Services Ontology*. TNO. URL: <http://ontology.tno.nl/logiserv/>.
- Daniele, L. (2013c). *Transport ontology*. TNO. URL: <http://ontology.tno.nl/transport/>.
- Daniele, L. and L. Ferreira Pires (2013). "An ontological approach to logistics". In: *Enterprise Interoperability, Research and Applications in the Service-oriented Ecosystem, IWEI'13 Proceedings*. Ed. by M. Zelm, M. van Sinderen, and G. Doumeings. ISTE Ltd, John Wiley & Sons, Inc., pp. 199–213.

³See <https://lists.w3.org/Archives/Public/semantic-web/2017Jul/0004.html> [2017/07/07]

- DiCicco-Bloom, B. and B. F. Crabtree (2006). "The qualitative research interview". In: *Medical education* 40.4, pp. 314–321.
- Garcia, A. (2015). *iCargo - Intelligent Cargo in Efficient and Sustainable Global Logistics Operations*. ATOS Spain SA. URL: http://cordis.europa.eu/project/rcn/100869_en.html.
- Hevner, A. R. (2007). "A three cycle view of design science research". In: *Scandinavian journal of information systems* 19.2, p. 4.
- Hevner, A. R., S. T. March, J. Park, and S. Ram (2004). "Design science in information systems research". In: *MIS quarterly* 28.1, pp. 75–105.
- Hofman, W., M. Punter, H. Bastiaansen, E. Cornelisse, S. Dalmolen, Z. Palaskas, B. Karakostas, J. Gato, J. Garcia, et al. (2016). "An interorganizational IT infrastructure for self-organization in logistics: situation awareness and real-time chain composition". In: *International Journal of Advanced Logistics* 5.2, pp. 101–115.
- Kleedorfer, F., C. M. Busch, C. Huemer, and C. Pichler (2016). "A Linked Data Based Messaging Architecture for the Web of Needs". In: *Enterprise Modelling and Information Systems Architectures* 11.3, pp. 1–18.
- Kleedorfer, F., C. M. Busch, C. Pichler, and C. Huemer (2014). "The Case for the Web of Needs". In: *Business Informatics (CBI), 2014 IEEE 16th Conference on*. Vol. 1. IEEE, pp. 94–101.
- Kleedorfer, F., Y. Panchenko, C. M. Busch, and C. Huemer (2016). "Verifiability and Traceability in a Linked Data Based Messaging System". In: *Proceedings of the 12th International Conference on Semantic Systems*. SEMANTICS 2016. Leipzig, Germany: ACM, pp. 97–100.
- Knublauch, H. and D. Kontokostas (2017). *Shapes Constraint Language (SHACL) - Proposed Recommendation*. [Last accessed on 2017/07/06]. World Wide Web Consortium (W3C).
- Manola, F. and E. Miller (2004). *RDF Primer*. [Last accessed on 2016/06/06].
- Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*.
- Uber (2017). URL: <https://www.uber.com/>.
- UNECE (1956). *Convention on the Contract for the International Carriage of Goods by Road (CMR)*. Geneva: United Nations.
- UNECE (2008). *Additional Protocol to the CMR concerning the electronic consignment note (e-CMR)*. Geneva: United Nations.
- Varvasovszky, Z. and R. Brugha (2000). "A stakeholder analysis". In: *Health policy and planning* 15.3, pp. 338–345.