The Case for the Web of Needs

Florian Kleedorfer, Christina Maria Busch
Studio Smart Agent Technologies
Research Studios Austria FG
Thurngasse 8/2/16, 1090 Vienna
fkleedorfer@researchstudio.at, cbusch@researchstudio.at

Christian Pichler, Christian Huemer
Institute for Software Technology and Interactive Systems
Vienna University of Technology
Favoritenstrasse 9-11, 1040 Vienna
cpichler@ec.tuwien.ac.at, huemer@big.tuwien.ac.at

***Abstract*—E-marketplaces on the worldwide Web are information and transaction silos, which in the general case don't allow transactions across their boundaries. The consequence is that the Web, often termed the global marketplace, is fragmented along the dimensions of geography, content domain, supply or demand, user base, and many more. This fragmentation makes it inefficient to buy and sell commodities on the Web. We propose a framework that serves as a foundation for a distributed, decentralized e-marketplace on top of the Web, making boundaries between existing systems disappear from the user's perspective. The framework standardizes the creation and description of objects that represent supply and demand. In addition to this, it allows for independent matching services to connect objects suitable for a transaction and it defines protocols for the message exchange between such objects.**

I. Introduction

The World Wide Web was designed as a general information medium. At its inception in 1990, the guiding idea was to build an open network of documents that could be linked to each other. At the beginning, manually curated directory pages were used to discover documents; however, as the Web grew in size, these became impractical to use and even harder to maintain. The situation gave rise to the idea of applying the methods of information retrieval, hitherto mostly applied to well-controlled collections [1], to documents on the Web in the form of search engines such as WWWW (the Worldwide Web worm) [2] in 1994. These services had been the missing ingredient that allowed the Web to scale beyond human manageability while remaining generally usable. Search services made all the difference as they let users structure the Web's content spontaneously according to their needs.

The World Wide Web evolved into more than an information medium: creators of video, music and of other pieces of art transformed it into a cultural medium; social networking services made it a social medium; e-commerce services made it a medium for business. Its development has been pervasive; people increasingly use the Web to satisfy all kinds of needs. Through all these transformations, it has kept its original bipolar form, a symbiosis of producers and consumers of documents connected by search engines. This form is probably not a coincidence, as it functions like an abstraction of a marketplace, where physical objects are on display on different kiosks while clients have to look for things that satisfy their needs, and eventually buy them.

Marketplaces have been known to man for several thousand years. Their historically consistent overall form, a collection of kiosks with goods being offered, and many independent clients buying there, is due to the fundamental asymmetry at the basis of all trade: the asymmetry between its two poles, between supply and demand. Although conceptually, supply and demand can be said to be of the same kind, namely intentions of taking part in a resource transfer, they differ substantially. Supply appears as the thing being offered, it is thus concrete, can be searched and found, measured and compared. Individual demand, in contrast, denotes the absence of something, thus it doesn't appear as a physical thing in a traditional marketplace.

The fact that the World Wide Web resembles a marketplace so strikingly does not necessarily mean that it has to adopt its weaknesses. It is, on the contrary, the central stance of this work that the Web can evolve to transcend its historic roots and alleviate some of the asymmetry of trade. Web documents can describe demand as well as they can describe supply; all that is needed is an commonly understood description language. For the Web to become a need satisfaction medium, two additional ingredients are required. First, in analogy to Web search engines at the beginning of the Web, matching services connect published needs and published offers with each other. Second, the ensuing negotiations can be held using a set of open communication protocols. With these additions, e-commerce is directly integrated in the Web's infrastructure, which we refer to as the Web of needs (WoN). This infrastructure provides the functionality to publish documents that are meant to be connected to one or more appropriate counterparts, to allow users to make these connections based on recommendations by dedicated services and to exchange messages according to predefined protocols.

More generally speaking, such an infrastructure enables a wide range of applications beyond e-commerce. It provides the basic building blocks for any social application. The necessary prerequisite of all human communication is to identify the appropriate communication partners, which – in the off-line world – happens as a result of context, a personal embeddedness within relations of locality, friendship, family, leisure, work, trading, and others. In parallel to the more specialized forms of supply and demand, this context can be captured in machine-readable Web documents. By the same independent matching services introduced for connecting supply and demand, appropriate communication partners can be identified. The proposed communication infrastructure allows them to connect to each other and exchange messages, and thus functions as the basis for any kind of transaction or conversation between humans, irrespective of its purpose or domain.

978-1-4799-5779-8/14 $31.00 © 2014 IEEE DOI 10.1109/CBI.2014.55

In this vision paper, we first motivate the design of our framework using a simple business case in Section II. We discuss design decisions in Section III, where we also give an overview of the architecture. After discussing the state of the art in relevant domains in Section IV, we illustrate possible consequences of the adoption of the proposed architecture at Web scale in Section V.

II. Motivation

The business case used as a running example is trading in a marketplace, understood in a very general sense – without restriction with respect to content domain or location.

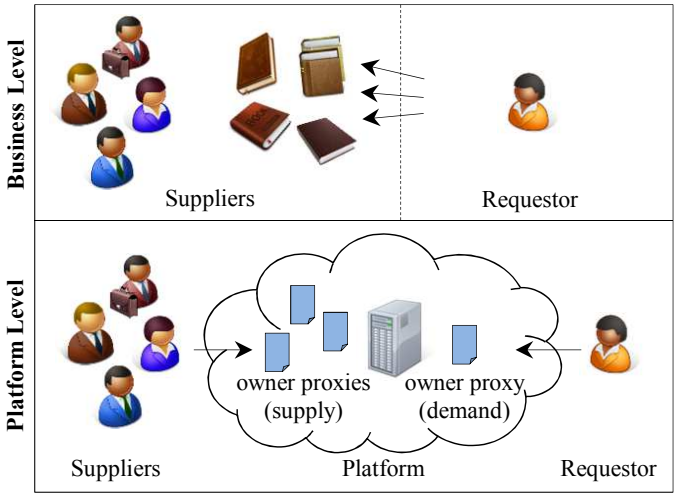

Fig. 1. Overview of the business case solved by the Web of needs.

In such a setting, *suppliers* intend to sell and *requestors* intend to buy commodities. A process of discovery and selection leads to a transaction involving requestors, suppliers, and goods or services. The abstract view is illustrated in the upper section of Figure 1, showing the example business case of a requestor seeking a book and multiple suppliers offering books that may be of interest to the requestor.

In the following we develop an appropriate representation of this business case for a Web based system. Then, we discuss consequences and problems this representation entails as a basis for major design decisions. We concentrate on three main steps a user performs: describing their supply or demand, identifying trading partners, and conducting a transaction.

A. Modeling Supply and Demand

Description. The view initially described uses roles of people or organizations as elements of the representation: the supplier is used to represent the offer, the requestor is used to represent the demand. These are appropriate abstractions for reasoning about personal relationships. For the sole purpose of organizing transactions in a marketplace, however, this description is overly complex: involving the notion of a person evokes many associations that are irrelevant to the purpose, such as gender, age, etc. More importantly, it incurs unnecessary constraints regarding identity: when considering individual transactions, viewing them as occurring between people induces a relationship between the transactions that involve the same people. Such relationships may be interesting and exploitable in many ways but they should only be added when needed. For example, when making recommendations or building a track record of a user's transactions, this information is important. In the general case, though, it is not. In the depicted example, age and gender of the requestor is irrelevant for conducting a transaction, so are other things they are currently looking to buy. By choosing to disregard the person and focusing on the intention, the irrelevant can be ignored.

For identifying the right abstractions, we consider that prior to entering a transaction, each participant must have developed a mental representation of the transaction and the reason for entering it: in our example, the idea of wanting to sell or buy a book. In the case of the supplier, this representation is often called 'offer' when made public. On the requestor's side this representation is sometimes called a request for proposal. In most cases, however, it is never externalized as an individual object. These abstractions, developed by the parties involved in a transaction, are suitable for representing the process of preparing and conducting transactions in a software system: they pertain to exactly one plan or goal and they do not incur unnecessary relationships to other notions or transactions. In our example, the requestor has to develop the idea that they intend to buy a book. This idea may identify a specific book or encompass certain properties it should have (e.g., a cookbook for Italian food). Likewise, the suppliers develop the idea to sell a specific book..

Reified as objects in the system, such representations are used to identify the transaction's endpoints, serving as proxies for their respective creators or *owners*; we will refer to them as *owner proxies* (or proxies for short). These objects contain a description of the reason for entering a transaction, (i.e., supply or demand) and allow their owners to connect to other objects of this kind and communicate. This view is depicted by the lower portion of Figure 1, where we see that the requestor controls a proxy representing their demand for a book and the suppliers control proxies representing their offers.

When describing a commodity, one chooses a level of granularity depending on the context in which the description is needed. When the commodity, the requestor, and the supplier are physically present, and the type and conditions of the transaction are prescribed by the circumstances, descriptions are mostly unnecessary - as is the case in traditional public marketplaces. In electronic marketplaces, the same commodities may be described with the highest level of detail. A generally useful marketplace solution must be flexible enough to support both cases equally well, therefore users should be given the options to describe their intentions in unstructured form using properties such as title, description, and tags. In addition to this there should be support for the input of generally useful properties such as time of availability, price information and location information. For expressing fine-grained properties, it should be possible to use an appropriate data structure. The example use case may occur in an informal setting, as may be encountered when trading used books. In such a setting, using title and description may be enough. The owner proxies representing an industrial publisher's books may contain very detailed domain-specific data such as ISBN, year, and other properties.

Identification of trading partners. Traditionally, the identification of a transaction counterpart is the responsibility of the requestor, who may have to spend quite some effort to do this. In the general case, however, the requestor's need is almost identical with those of many other requestors, all of whom have to go through the same identification process. Therefore, a lot of time could be saved if this task was done by a party dedicated to it. Search services and recommendation systems have been applied to tackle this problem in on-line settings. Both approaches have two drawbacks: First, they are only requestor-oriented and second, they do not offer the option of specifying what is needed. The representation developed above, in contrast, allows for independent services dedicated to matching supply and demand. Such services keep track of the owner proxies and in case new ones are encountered, they are compared to the known ones. The only required functionality such a service must offer is to send a message to the owner proxies that are suitable transaction partners, informing them of one another. In the book-selling example, a party independent from suppliers, requestors, and the platform operators creates the link between their owner proxies.

Conducting a transaction. When doing business, the context in which transactions are made and the nature of the commodities in question determines the communication that takes place. In some settings, complex protocols involving a number of participants have to be followed in order to seal a deal, while in others it is sufficient to exchange a few short sentences in natural language. Therefore, defining one protocol for all kinds of transactions is unrealistic. We can, however, identify the bare minimum that is required for communication: the participants have to decide whether they want to communicate. The process leading to a conversation can be described as a handshake. One party proposes the conversation, the other party may accept or decline. At any point, any party may decide to abort an ongoing conversation. The message exchange within a connection is governed by a protocol chosen by the partners, allowing for flexibility.

B. Consequences

Reification and Symmetry. The main advantage of this representation is that both endpoints of a prospective transaction are modeled as objects that can be manipulated. Thus, third parties can automatically or manually identify and propose appropriate transactions; users don't need to rely on their search skills or a recommendation system nor do they need to equip a personal software agent with information about their supply or demand and trust it to find the best solution. In the illustrating example, the requestor is not required to search for books. Instead, the wish for a book is represented by an owner proxy that can be found by matching services and connected to by other users.

Compositionality. A more subtle consequence of the objectification of both supply and demand in the form of owner proxies is that compositionality becomes possible on both sides. In the general case, the supply side defines and offers a commodity in the hope of meeting demand. The commodity may be a composition of goods and services, but it is intrinsically an atomic entity: if for some reason only a part of that bundle of goods and services were to be sold, that part would become a commodity in its own right; the same is true for the composition of commodities. However, there are costs involved in the creation of new commodities through composition or decomposition. This effort is only made if there is sufficient demand. All-inclusive holiday offers are a well-known example of such complex commodities. As objectification becomes possible on the demand side, it is no longer futile to express demand for combinations of commodities that don't exist yet because each element can be obtained from a different supplier. Moreover, such complex demand specifications are valuable in themselves as they encapsulate knowledge about how a certain goal can be reached through a combination of goods and services, and they can be re-used and re-mixed. Our book-selling example may not justify such complexity, but consider planning a clubbing event. It requires finding and coordinating artists, catering, a suitable location, marketing, and probably much more, all of which must be available at certain, inter-dependent dates and in inter-dependent quantities. Such a plan could be represented as a composition of rather simple owner proxies that are tied together with relationships expressing these constraints. The first formulation of such a plan might require expert knowledge and take a lot of time, but subsequent instantiations may not require quite as much expertise. The frequent application of the same plan may lead to incremental improvement and the development of different versions for different scenarios.

C. Selected Problems

Privacy and single points of failure. The standard approach for solving a business case as described so far consists of creating a Web site with e-marketplace functionality that offers the required services. If such a marketplace were to be used for any and all transactions, it would, however, represent a single point of failure in many respects (technical, political, security-wise), which is not desirable. Moreover, the party providing the services would automatically become all-knowing with respect to every user's supply and demand. This data is very personal and sensitive. Placing it in the wrong hands may ensue severe privacy breaches.

Trust. Trust is widely recognized as an essential aspect in the adoption of e-marketplaces [3], [4]. Approaches for solving the trust problem vary between different types of platforms (B2B, B2C, C2C) as well as within each category. Some platforms, like Craigslist [5], do not seem to manage it at all. Others like Ebay [6] have specialized systems that leverage users' ratings of their transaction partners. Some peer-to-peer marketplaces such as Airbnb [7] have adopted a very sophisticated approach to instilling their users' trust in each other, using multiple levels of profile information verification, mediating phone calls on behalf of their users, and even verifying personal ID documents. Findings reported by Keetels [8] support many of these measures as being positively related to users' trust in the facilitating platform or, respectively, each other. Such measures cannot be taken in de-centralized systems.

III. Architecture

Having derived the high-level view of the framework from the main business case, we describe the approach for its realization. We begin by explaining design decisions, then we give an overview of the framework's architecture.

A. Design Decisions

In an earlier work, the following non-functional requirements for the Web of needs were defined: access, usability, fairness, simplicity, scalability, timeliness, robustness, security, privacy and/or anonymity. [9] Together with the results from the previous section, these are taken into account in the design we develop in the following.

De-centralization. As stated before, privacy is essential, as is the neutrality of the infrastructure so as to ensure fairness. Both are much less of a problem in de-centralized systems than in centralized ones, as there is no single party with access to all the data and no opportunity for one party to control access. Moreover, the cost of running and growing a de-centralized infrastructure is naturally shared among those operating nodes of the network, so no or only little initial centralization of capital is required to establish such an infrastructure. Consequently, we envision the Web of needs as a network of nodes, which we call *WoN nodes*.

Simple creation of owner proxies. We have argued that it is important that needs, offers, or other intentions be represented as standalone entities, disconnected from the person or organization that created it; this entails that each user controls a considerable number of such entities. Clearly, for the system to remain usable, the process of creating such owner proxies must be simple and straightforward.

Flexible description language. The description of intentions represented by owner proxies can take a variety of forms. The formalism for expressing these descriptions in machine-readable form must be highly flexible, allowing for arbitrary structured data with clear semantics. For these reasons, we chose RDF [10] as data language.

World readable descriptions. The identification of communication partners by independent services can only be done if descriptions of owner proxies available to those services, therefore, they must be world-readable, with all consequences for user privacy. In order to allow this in a standardized way, information describing owner proxies is published as linked data [11]

Anonymity. The fact that owner proxies must contain world-readable information makes it an imperative that the person behind the owner proxy can be concealed - otherwise it would be possible to collect all information published by a person and make an accurate profile of their economic and social behavior. This form of anonymity is weak: the WoN node operator can still link the owner proxy to the owner. Still it protects the users' privacy against all other parties. For an optional strong form of anonymity, an onion routing scheme can be applied in which messages are routed through a chain of WoN nodes, each acting as an owner application toward the next one in the chain and routing the traffic back to the proxy's creator. [12]

Identity. In order to serve as an endpoint for reliable transactions, an owner proxy must be able to prove its identity to others. This is achieved by integrating cryptographic technology by using WebID. [13]

Always-on-line nodes and standardized proxy creation. Retrieval of information about owner proxies must be possible at all times to allow for asynchronous communication and matchmaking. Therefore, desktop computers or mobile apps cannot host their owner proxies themselves. Rather, for each owner proxy, they may choose from a set of available WoN nodes and use a standardized protocol to create an owner proxy there. The WoN nodes are servers dedicated to the task of hosting owner proxies and are always on-line. Thus, users can create different owner proxies on different nodes, achieving a minimum of privacy by distribution.

Access Control. While in general, owner proxy descriptions are publicly readable, there are situations in which only a limited circle of users should be given access to a certain owner proxy, for example, if a private party is being announced. For such cases, optional role-based access control is possible based on WebID. [13]

Standardized matching messages. For matching services to be independent of the other elements of the infrastructure, communication between matching services and owner proxies must be standardized. There is conceptually only one kind of message to be exchanged, namely a *hint* message sent by the matching service to the owner proxies whenever a match is found.

Communication based on messaging. *Owner applications*, programs for users to create owner proxies and to communicate with them, must at times be contacted from WoN nodes (e.g., when a hint message was received from a matching service). However, owner applications cannot be required to be always on-line or even publicly visible, so realizing their functionality using Web services is not an option. Consequently, messaging is chosen as the communication paradigm for our framework. In order to allow for flexibility, the messages are written in RDF.

Standardized establishment of connections. Whenever users decide to connect with each other, their owner proxies follow a simple protocol that leads to a *connection* being established.

Layered communication protocol. When an owner proxy is created, the user decides which *facets* it is going to support. A facet identifies a protocol by which the messages exchanged between two owner proxies are interpreted. When establishing a connection, the initiator names the facet of their owner proxy they want to use in the connection as well as the facet of the remote proxy that should be connected to. For example, in order for a requestor and a supplier to have a conversation, they can use the chat facet, which defines that the creation of a connection is interpreted as the willingness to exchange text messages. One of them connects its chat facet to the chat facet of the counterpart, and as soon as the other accepts the connection, they can exchange chat messages until one of them closes the connection. For informal and infrequent trading, this facet is sufficient. For distributed and automated transactions, protocols like WS-BA [14], realized as facets, are more suitable. The concept of the facet generalizes the framework such that the realization of a global distributed marketplace is only one of many applications that can be built. For example, the proxy may support a facet for 'friending' and 'following', and thus represent its owner in a distributed social network.

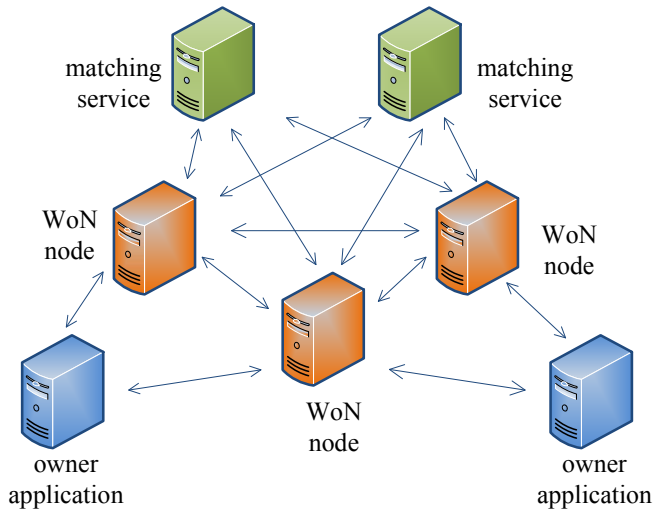

Fig. 2. Deployment diagram indicating types of nodes and communication paths.

Public connection information. When a connection is established between two owner proxies, each one creates an object representing that connection and publishes it as linked data, which means that it is publicly shown which proxies are connected and what the state of their conversation is. The decision of what information is published about a connection (for example, if individual messages are published or not) is the responsibility of the facet and determined by the protocol obeyed by the facet.

Expressing relationships between owner proxies. Complex structures of owner proxies, as described in Section II-B can only be created if the description language for owner proxies allows to create relationships between them that are meaningful with respect to the execution of transactions. For example, expressing that all needs described by a number of owner proxies must be satisfied in order for the transaction to finish successfully.

Trust as a Layer. As stated before, measures that instill trust in users can quite easily be implemented in centralized services. The reason for this is that there is a single instance that controls transactions, sanctions user behavior, etc. In a decentralized infrastructure no party can be forced into behaving in a certain way by technical means. It is, however, possible to keep a cryptographically validated track record of transactions that can be shown to and verified by prospective transaction partners. As for many relationships, such a track record is not required, it is realized as an optional layer in the WoN protocols. The Owner proxies controlled by a person may link to one owner proxy managing this person's track record. The transaction partner may choose to sign the entry for the current transaction, leaving feedback and ratings. The reputation information is formally described using the concept of the *reputation object* [15].

Thin clients. In order to allow for controlling a user proxy intermittently from different applications (e.g. a mobile application and a Web site), all information is kept on the WoN node hosting the proxy. As all relevant information is available as linked data, owner applications (i.e., clients) do not need to keep track of all events pertaining to their proxies; whenever needed, local information can be refreshed, provided that the application can authenticate itself as the proxy's legitimate owner. However, the client is only 'thin' with respect to one owner proxy. Managing tens or hundreds of owner proxies is considered normal use; owner applications therefore have to provide the functionality to keep track all those proxies and their login credentials.

B. Architecture Overview

In the bird's eye view, the Web of needs consists of three types of elements, as depicted in Figure 2. The basic functionality – allowing CRUD operations on owner proxies, establishing connections between them and receiving hints from matching services – is provided by *WoN nodes*. *Owner applications* allow users to create and control their proxies. *Matching services* crawl the proxies' data found on WoN nodes and send *hint* messages to owner proxies. Figure 2 shows which components communicate directly.

The goal of our work is to define a framework for an open, de-centralized worldwide marketplace. In order to reach this goal, existing technologies are arranged so as to provide a basis for end-user oriented applications. Figure 3 illustrates the most prominent base technologies in the bottom layer. The basic protocol layer provides secure CRUD operations, communication between owner proxies and discovery of suitable communication partners. The optional layers provide strong anonymity, standardized links to identities, inclusion of trust mechanisms, and different flavors of transactionality. On top of this stack, arbitrary facets can be implemented to map important social interaction patterns, which in turn can be combined in end-user facing applications. For example, the combination of a facet for conducting distributed transactions, a facet allowing to review transaction partners, and a facet for chat may suffice for establishing a distributed e-marketplace.

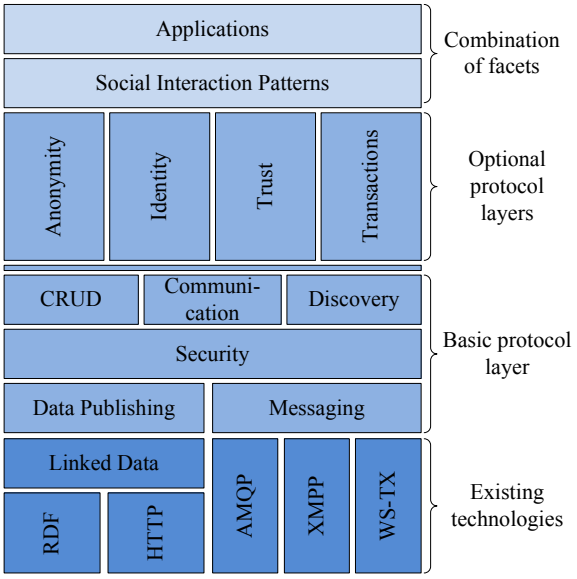

Fig. 3. Diagram illustrating the layers in the protocol and technology stack.

IV. State of the Art

Having motivated and explained our approach we now give an overview of related and relevant systems or concepts and compare them to the Web of needs.

E-Marketplaces. To the best of our knowledge, there is no prior work on global, open, web based market infrastructures to learn from or build upon. The well-established and closely related concept of e-marketplaces, though, has been scientifically examined. The approach pursued in the Web of needs, we argue, differs from traditional e-marketplaces insofar as it does not create separate vertical marketplaces for different niches; rather, it creates one unified marketplace on the Web. This helps to lower transaction cost below the current levels in the long run and provides a marketplace for niches that aren't profitable enough for a dedicated Web site to emerge. We do not have estimates for the size of these 'long-tail markets' where supply and demand must be assumed to exist but lack mediation; if the long tail phenomenon translates from other domains to this one, however, it should be considerable.

Traditionally, e-marketplaces have been used in business-to-business (B2B) or business-to-consumer (B2C) context. In addition to B2C and B2B marketplaces, examples of successful consumer-to-consumer (C2C) e-marketplaces have emerged in recent years. These include classified ads portals and auctioning websites like Craigslist [5] or Ebay [6] as well as specialized platforms such as Airbnb [7] or Couchsurfing [16]. In C2C e-marketplaces, individuals are responsible for both offerings and purchases. All of these types of marketplaces are realizable on the basis of the Web of needs infrastructure, which allows transactions that span domains that are currently organized in vertical marketplaces (such as holiday apartment rentals, taxi services, or restaurant bookings), thereby clearly offering additional useful functionality that cannot be provided by any such platform alone, namely the unified access to all the verticals, and hence the possibility to combine transactions, for example, to buy opera tickets, reserve a table in a restaurant and order a taxi without the burden of switching marketplace.

Very closely related to the Web of needs is the effort of publishing semantically rich offer and demand descriptions on existing e-marketplaces using vocabularies like GoodRelations [17] or Schema.org [18]. The main difference is that these vocabularies are tools for describing entities, but they do not define service interfaces; entities suitable for interaction can be found automatically, but there is no standardized way to establish a connection with them.

Intention Economy. The term *intention economy* denotes an environment in which customers use software systems to manage their relationships with vendors [19], so-called *vendor relationship management* (VRM) tools. Among other functionalities, such a tool supports expressing demand in the form of a *personal request for proposal* (pRFP) [20]; the act of publishing such a pRFP is referred to as intentcasting. The notion of the pRFP is quite similar to our concept of a owner proxy representing a demand, and the technical artifact most similar to the owner proxy is the pico (persistent computing object) in the Kynetix Rules Engine [21], [22]. A difference is the VRM community's focus on commercial activities, in contrast to which the Web of needs is intended for more general use. It remains to be seen how the Web of needs infrastructure fits in with the tools that are being developed in the VRM community.

Discussion Groups. Countless marketplaces are organized on the Internet in discussion groups. The technical bases range from Web forum software over mailing lists to groups in social networking Web sites. The main advantages of these groups are simplicity, openness, and the like-mindedness of participants. They are normally self-regulated with respect to a stated or implicit code of conduct and allow postings consisting of text and images (as opposed to structured data). They work best below a certain frequency of messages, because users have to read all messages so as to decide whether they want to trade. If postings are added with too high a frequency, it becomes hard to follow the updates, which reduces the usefulness of the group for individual users, although in principle the higher number of available options for trading should increase usefulness. At some point such groups tend to split up in multiple smaller ones. Consequently, such approaches lead to a large number of groups, causing users the problem of identification of the right group for a given case. Moreover, there often are a number of redundant groups for the same type of commodity, location etc., in different channels such as Facebook [23], mailing lists, or dedicated Web sites, making the decision for one of them even more difficult. The Web of needs infrastructure alleviates all of these problems and may prove more useful to users in the long run.

Social Networks. The basic functionalities social networking platforms offer are creating and maintaining relationships between user identities and using them for communication purposes. With respect to user-to-account cardinality as described by Dalton [24], such platforms mostly try to achieve a one-to-one cardinality, i.e., one user has one account. Of the big players, only Twitter [25] allows one-to-many and many-to-one cardinalities. Users with profiles on more than one social networking site implicitly have a one-to-many relationship with these profiles, where the platforms set the context for these accounts, resulting in differences in the profile characteristics. For example, Linkedin [26] focuses on business relationships whereas Facebook is more about family and friendship; users shape their profiles accordingly. In contrast, based on the facet generalization introduced in Section III-A, social networking functionalities can be built on the Web of needs allowing any user-to-profile cardinality. In such applications, an owner proxy can represent a user; if credentials are shared, or access is delegated otherwise, it can represent many users, allowing one-to-one and many-to-one cardinalities. Users can create separate social networks as needed, for different contexts such as family, friends, political contexts etc. In each such network, the same person has a different identity. These identities can be aggregated by the user as desired.

In the context of social networking, a highly relevant related work is the project on distributed semantic social networking (DSSN) [27] that takes a more lightweight approach for data access and communication than was chosen for the Web of needs. While for our architecture, social networking is one of many possible applications, DSSN is specifically designed for it but can be combined with other semantic Web based systems to offer richer functionality.

V. VISION

In the following, we discuss possible consequences that large scale adoption of the Web of Needs might have.

In the Web of needs, each Internet user can create and control any number of proxies for any purpose. The proxies are found by services whose purpose it is to connect them. Thus, complex communication structures are created spontaneously to help organize the solutions to complex societal or individual problems. These structures can be understood as recipes for social solutions. They can be observed, revised and re-used so as to help the overall optimization of societal processes.

Users manage all kinds of needs using owner applications, ranging from short-lived tasks such as ordering a taxi cab and tracking its location while waiting, to long-term plans like organizing a vacation, from recurring tasks like grocery shopping to one-time needs such as moving furniture, from satisfying personal needs to industrial procurement. On the other hand, users offer their products and services using these programs. Moreover, using the technology not only for purely economic purposes, people publish their discontent about the state of public affairs, enter group discussions with like-minded people and make common decisions, or collaborate spontaneously to fix neighborhood problems.

All published owner proxies are anonymous and readable only to those users chosen by the owner, in many cases, though, they are publicly readable. Part of this information is the current state of connections with other owner proxies. Such connections represent transactions or dependencies between the owner proxies' objects.

These objects, published on the Web, represent the Internet users' intentions, the means of their satisfaction and their dependencies between each other. In total, this structure can be interpreted as the dream of humanity, describing what should be (demand), and its potential, by showing what could be (supply). The collective process of reconciling the two sides is humanity's attempt to make its own dream come true. The success of this endeavor is recorded and visible for all on the Web of needs. Large scale analysis of the data in this system provides insight into market opportunities. By applying methods of operations research, optimization is done on a global scale, calculating and updating plans for (approximately) optimal need satisfaction with minimal effort or cost, which is proposed in the form of hint messages (the messages sent by matching services). Local analysis of one's own owner proxies provides insight into why aspects of one's own wishes are not fulfilled by society and helps organizing with others to overcome the problem.

As people habitually record their wishes and plans in the form of owner proxies, computer programs can help to come to terms with them - show which ones are easily fulfilled, which ones are conflicting, and help to review and prioritize them. Wishes, interests or problems that others share are automatically used to form interest groups that make it easy to get support, meet like-minded people, or act collectively.

For many common use cases, technical systems and social systems are interleaved. For instance, sensors that detect free parking spaces in cities publish them automatically as owner proxies on the Web of needs. Car navigation systems are connected to the same network and connect to such proxies as soon as the goal of the journey is calculated, receiving updates on availability. They establish connections with proxies of the other drivers who want to use the same parking space to get their location updates and so as to be able to calculate if a different location should be selected. Sometimes, buildings cannot be used commercially for some time and the owners do not want them to be unused as this is bad for the building's substance. In such a case they can publish owner proxies describing the availability of space, which is matched with respective demand for office or living space by charities, startups, artists, or other groups or individuals. Transportation and logistics is highly integrated with personal and commercial activities, anyone can offer such services on the WoN, from an individual making a quick buck to a multinational logistics company. On the demand side, such transport requests are seamlessly integrated in the structure of owner proxies involved in, e.g., selling an old couch.

The framework presented here can only rise to its potential if it is widely accepted and implemented. This is the most important challenge facing our work, which is why we explain how we think uptake could develop:

In the long run we expect the Web of needs to undergo a development that is characterized by phases in which new classes of participants start using it, thereby changing its characteristics. With each such class, the technology becomes more and more attractive for new parties until they decide to adopt it. This process may evolve with different speeds at different locations or in different industrial sectors.

C2C Phase. Initially, there is a cold start problem: where there are no requestors, suppliers have no reason to be and vice versa. The infrastructure is suited for non-commercial and symmetric cases, for example for bringing people together in social contexts like sports (finding appropriate partners) or citizen empowerment (finding common problems and organizing groups in order to fix them). The most promising application domain, however, is that of gift marketplaces, in which participants tend to switch between the roles of supplier and requestor, and which, as stated in Section IV, are not ideally managed with current technology, often organized as discussion groups or vertical marketplace Web sites.

Specialist B2C Phase. The first phase creates a marketplace mainly populated by consumers. Small at first, and with local foci, this user base becomes interesting for commercial suppliers of specialized and only locally available offerings as it opens up a cheap new distribution channel for them. Such suppliers could let their decision to move to this marketplace depend on published consumer demand in the relevant niche, thereby lowering the risk of bad investment inherent to such adoption. Examples for such services could be local retail, transportation and logistics (supporting the gift/second hand transactions), or arbitrary services, like running errands.

Mediator phase. When commercial usage of the Web of needs becomes normal, it starts to exert pressure on existing e-marketplaces, to a point where adoption of the technology becomes an advantage for them as they can access a user base untapped by the competition, and they may find that they already offer services, like logistics and bulk processing, that are not yet available on the Web of needs, putting them in a

favorable position as mediators. The effect of this development is a spurt in commercial offerings for industrial products in the system. With the advent of software tools that make it easy to work with the technology for customers and producers alike, adoption becomes cheaper and rises in niche markets, until it reaches the long tail of currently impossible transactions.

B2B phase. When industrial products become available en gros in the Web of needs, it starts to be interesting for obtaining production supplies. Industry finds that the automatic matching services and the ability to tap into a worldwide marketplace make their procurement more and more efficient, to a point where it can compete with the supply chains of big industry.

Industrial market phase. At some point, industry may realize that it can greatly reduce the cost involved in selling their products as they can offer them directly on the Web of needs. Participants offering integrated logistics handle transportation.

VI. Conclusion

In this paper, we have focused on motivating why the Web of needs as a generic cooperation framework can serve as a basis of a worldwide on-line marketplace. We explain design decisions, and communicate the vision of the opportunities such a framework would create. The latter can only be done by sketching scenarios that have to be combined and completed by the reader's imagination.

Our main stance is that a worldwide marketplace must be as de-centralized and open as the worldwide Web. In general, transactions should be recorded publicly. The central element of such a marketplace is the owner proxy, an entity anonymously controlled by a user. It contains a description of the task it has been created for, in our running example is the purchase or sale of a book, and encapsulates the least amount of data and functionality that is required to perform the desired task. Owner proxies are made aware of each other by independent matching services that compare their descriptions and inform them about possible transaction counterparts.

We used a traditional commercial business case for illustrating the proposed framework. The design of the framework was influenced by empirical analysis of sharing communities which is beyond the scope of this paper. This fact, however, explains the current emphasis of natural language messaging and an understanding of transactions that is close to that of informal conversations. Further work is required for enabling the framework for traditional business transactions. We intend to evaluate the framework in the domain of sharing communities before adapting it to a commercial domain.

At the time of this writing, the basic infrastructure is being developed. Prototypes are working in a distributed fashion, and a primitive owner application and a library for automated publishing and interaction via owner proxies are operational. [28] Work on security-related features, scalability of matching and communication, and on user experience is ongoing.

VII. Acknowledgments

This work was supported by the Austrian Federal Ministry of Science and Research and the Austrian Research Promotion Agency (FFG) in the BRIDGE project *WIN – Web of Needs Infrastructure* and the COIN project *USS WON – Usability, Scalability and Security on the Web of Needs*.

References

[1] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer networks and ISDN systems*, vol. 30, no. 1, pp. 107–117, 1998.

[2] O. A. McBryan, "Genvl and wwww: Tools for taming the web," in *Proceedings of the first international world wide web conference*, vol. 341. Geneva, 1994.

[3] P. O'Reilly and P. Finnegan, "Performance in electronic marketplaces: Theory in practice," *Electronic Markets*, vol. 15, no. 1, pp. 23–37, 2005, (cited in Shish et al., 2013).

[4] P. Patnasingam, D. Gefen, and P. A. Pavlou, "The role of facilitating conditions and institutional trust in electronic marketplaces," *Journal of Electronic Commerce in Organizations (JECO)*, vol. 3, no. 3, pp. 69–82, 2005.

[5] craigslist. Website. [Online]. Available: http://www.craigslist.org/

[6] Ebay. Website. [Online]. Available: http://www.ebay.com/

[7] airbnb. Website. [Online]. Available: https://www.airbnb.com/

[8] L. Keetels, "Collaborative consumption - the influence of trust on sustainable peer-to-peer product-service systems." Master's thesis, Utrecht University, 2012.

[9] F. Kleedorfer and C. M. Busch, "Beyond data: Building a web of needs," in *Proceedings of the WWW2013 Workshop on Linked Data on the Web (LDOW 2013)*, Rio de Janeiro, Brazil, May 2013.

[10] F. Manola and E. Miller, *RDF Primer*, 2004. [Online]. Available: http://www.w3.org/TR/rdf-primer/

[11] S. Speicher, J. Arwe, and A. Malhotra. (2013) Linked data platform 1.0. Linked Data Platform Working Group. [Online]. Available: https://dvcs.w3.org/hg/ldpwg/raw-file/default/ldp.html

[12] D. Goldschlag, M. Reed, and P. Syverson, "Onion routing," *Communications of the ACM*, vol. 42, no. 2, pp. 39–41, 1999.

[13] A. Sambra, H. Story, and T. Berners-Lee. (2013, December) Webid 1.0 - web identity and discovery. W3C. [Online]. Available: https://dvcs.w3.org/hg/WebID/raw-file/tip/spec/identity-respec.html

[14] *OASIS Web Services Business Activity Version 1.2*, OASIS Std., Feb 2009. [Online]. Available: http://docs.oasis-open.org/ws-tx/wsba/2006/06

[15] R. Alnemr, A. Paschke, and C. Meinel, "Enabling reputation interoperability through semantic technologies," in *Proceedings of the 6th International Conference on Semantic Systems*, ser. I-SEMANTICS '10. New York, NY, USA: ACM, 2010, pp. 13:1–13:9. [Online]. Available: http://doi.acm.org/10.1145/1839707.1839723

[16] Couchsurfing. Website. [Online]. Available: https://www.couchsurfing.org/

[17] M. Hepp, "Goodrelations: An ontology for describing products and services offers on the web," in *Proc. of the 16th Int. Conf. on Knowledge Engineering and Knowledge Management (EKAW2008)*, ser. Springer LNCS, vol. 5268, Acitrezza, Italy, September 2008, pp. 332–347.

[18] *Documentation*, Schema.org. [Online]. Available: http://schema.org/docs/documents.html

[19] D. Searls, *The intention economy: when customers take charge*. Harvard Business Review Press, 2012.

[20] Personal RFP. Wiki page. Project VRM. [Online]. Available: https://cyber.law.harvard.edu/projectvrm/Personal_RFP

[21] Kinetic rules engine. Website. [Online]. Available: https://github.com/kre/Kinetic-Rules-Engine

[22] P. J. Windley, *The Live Web – Building Event-Based Connections in the Cloud*. Course Technology, 2011.

[23] Facebook. Website. [Online]. Available: http://www.facebook.com/

[24] B. Dalton, "Pseudonymity in social machines," in *WWW (Companion Volume)*, 2013, pp. 897–900.

[25] Twitter. Website. [Online]. Available: http://www.twitter.com/

[26] Linkedin. Website. [Online]. Available: http://www.linkedin.com/

[27] S. Tramp, P. Frischmuth, T. Ermilov, S. Shekarpour, and S. Auer, "An architecture of a distributed semantic social network," *Semantic Web*, vol. 5, no. 1, pp. 77–95, 2014.

[28] Web of needs on github. Website. [Online]. Available: https://github.com/researchstudio-sat/webofneeds/