

PERFORMING PEER-TO-PEER E-BUSINESS TRANSACTIONS: A REQUIREMENTS ANALYSIS AND PRELIMINARY DESIGN PROPOSAL

Stephanos Androutsellis-Theotokis, Diomidis Spinellis and Vassilios Karakoidas
*Athens University of Economics and Business, Department of Management Science and Technology
 76 Patission Street, GR10434, Athens, Greece.*

ABSTRACT

On-line business transaction processing systems have so far been based on centralized or client-server architectures. It is our firm belief—and it has also been recognized by the research and industrial community—that such systems may also be based on the constantly evolving decentralized peer-to-peer architectures. The first step in this direction, which constitutes the core of our paper, is a detailed requirements definition and analysis. We discuss requirements preceding the actual collaborations, such as support for discovery of services, merchandise or trading parties, authentication and access control, and negotiating collaboration parameters; requirements referring to the actual collaboration and transaction phases, such as support for workflow and collaboration orchestration, logging and non-repudiation; requirements following the collaboration, such as user ranking and reputation management; and generic non-functional requirements including security, availability and anonymity. A preliminary design proposal is presented, based on our proposed set of requirements, and on implementation solutions from the recent literature. We conclude that current peer-to-peer technology has evolved to the extent that it is able to fulfill many of these requirements to a large extent.

KEYWORDS

p2p, peer-to-peer, business transaction, on-line transaction processing, b2b, on-line collaboration.

1. INTRODUCTION

Business-to-Business (B2B), Business-to-Consumer (B2C) and Business-to-Government (B2G) systems have so far been based on a variety of centralized or client-server models, featuring a central authority that mediates and orchestrates the collaboration procedure, including searching for parties or products, providing authentication and access control, dictating the transaction business logic and supporting non-functional security requirements such as integrity, non-repudiation, privacy and anonymity where necessary.

The main business models currently encountered in the internet, typical of the above trend toward centralization and aggregation include:

- The old-economy “one-to-one” EDI model, for direct point-to-point transactions between companies.
- The “one-to-many” model, in which a supplier provides access to a large group of buyers.
- The “many-to-one” model, representing buy-side solutions (many suppliers, one buyer).
- The “many-to-many”, or net markets model, which aggregates vast numbers of buyers and sellers around a central “hub” providing services ranging from basic directory indexing of products and trading parties to complex transaction logic and workflow processing. Typical examples include eBay and Yahoo.

The research and industrial community has recognized [White et. al., 2003; Schoder, 2004; Choi and Whinston, 2001] that business transaction processing systems could also be based on the new, decentralized “peer-to-peer” model. In the recent years, the wave of decentralized peer-to-peer architectures, exemplified by systems such as Napster, Kazaa and Gnutella that were originally only appropriate for best-effort file sharing and basic collaboration, has evolved to provide the infrastructure and non-functional characteristics required for implementing much more demanding and complex tasks [Androutsellis-Theotokis and Spinellis, 2004]. Factors motivating the use of peer-to-peer architectures in supporting business transactions include:

- Absence of cost and risk of ownership and maintenance of a centralized server and all the relevant business and transactional data.
- Improved scalability and ability to deal with transient populations of users.
- Dynamic, ad-hoc communication, network organization, failure recovery and enhanced reliability.
- Democratic participation (especially regarding SMEs) and censorship resistance.
- Direct, un-mediated and potentially synchronous peer-to-peer transactions.
- Improved access to resources.
- The explicit exposure of trust relationships that are typically implicitly handled in centralized systems.

The rest of this paper is organized as follows: In Section 2 we present our main focus and the limited related work; in Section 3 we present a definition of business transaction, which will guide our analysis; In Sections 4&5 we present our requirements analysis and preliminary design proposal; and the paper concludes with Section 6.

2. FOCUS AND RELATED WORK

To the best of our knowledge, work on designing peer-to-peer systems to support business transactions has so far centered around generic and rather abstract discussions of ideas, problems that need to be solved and potential system requirements. [Charalabidis et. al., 2004] present a relevant yet centralized system, supporting various B2B and B2G transactions through application interconnection.

The only design and implementation attempt we are aware of [Gao, 2004] is based on the JXTA peer-to-peer infrastructure and suggests practical but simplistic solutions for a very narrow and limiting subset of the range of business transactions that should be supported by a complete solution.

What we believe is currently lacking in this direction is a detailed and specific requirements definition and analysis (see also [Schoder, 2004]), including an analysis of desired features, characteristics and properties, and accompanied by an account of whether and how these features can be supported by existing peer-to-peer technologies, and what additional research work is required, if any. This is the main focus of our short paper.

3. DEFINING BUSINESS TRANSACTIONS

According to the International Organisation for Standardisation (ISO), a business transaction is “a predefined set of activities or processes of organisations to accomplish an explicitly shared business goal” [ISO/IEC 14662:1997]. A business transaction may involve any number of participants, it may be instant or last for years, and it can have various degrees of complexity.

We base our following analysis on the concepts of *Business Transaction* and *Business Collaboration*, as defined by the ebXML Business Process Specification Schema [ebXML, 2001]. Figure 1 illustrates the following definitions:

- **A *Business Transaction*:**
 - Can only involve two parties.
 - Is an atomic unit of work that can result in either a success or a failure.
- **A *Business Collaboration*:**
 - Can involve any number of parties.
 - Is a combination of choreographed Business Transactions, and defines the ordering and transition between them.

In our following analysis we discuss:

- Requirements pertaining to the events prior and after the actual collaboration.
- Requirements concerning both the entire collaboration level, and the individual transaction level.
- Generic, non-functional requirements.

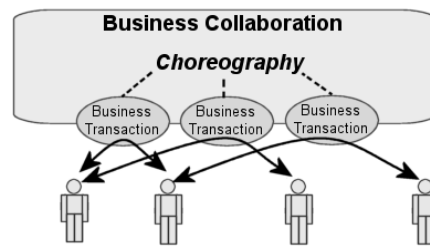


Figure 1: Schematic illustration of a hierarchical *Business Collaboration* between various parties, consisting of choreographed Binary *Business Transactions* between pairs of parties

4. REQUIREMENTS DEFINITION AND ANALYSIS

The following requirements analysis focuses on different aspects of the collaboration and transaction procedure. The proposed peer-to-peer model is considered with reference to the corresponding centralized models. A detailed presentation of the various peer-to-peer based solutions referred to can be found in [Androutsellis-Theotokis and Spinellis, 2004] and the references therein.

4.1 Pre-collaboration requirements

4.1.1 Discovery of partners, merchandise and services

Collaborations are often preceded by a discovery phase, in which the partners, merchandise or services are located. In a centralized approach, this is achieved through the use of centralized catalogues or indexes maintained and updated on the server.

In the proposed peer-to-peer approach, a naming, description and network-wide searching infrastructure would be required. A variety of robust, efficient and adaptable search algorithms exist, capable of coping with highly transient peer populations and constantly changing data. Currently, however, such approaches rely either on basic keyword searching or on the use of globally unique file identifiers. Searching based on a rigorous and flexible (probably XML-based) metadata model would need to be supported for the description and location of services, products, partners, and collaboration details.

4.1.2 Authentication and access control

It is of paramount importance for the authenticity of the identity of all partners engaging in a collaboration to be verified. In a centralized approach, this is certified by the server, through a membership / subscription mechanism or reliance upon trusted third parties (credit card issuers or certification authorities).

In peer-to-peer networks, the research community is still struggling to overcome the problem of peers assuming multiple identities or continuously changing identities. A solution that has received considerable attention and offers promising prospects is the use of “webs-of-trust” in peer-to-peer environments [Khambatti et. al., 2004], to certify the trust-worthiness of peers.

It should be noted that due to the legally and financially binding nature of business transactions, there is the additional requirement that any member of the proposed peer-to-peer system is required to possess and be able to certify a “real-life” identity.

4.2 Collaboration level / Orchestration Requirements

4.2.1 Rendezvous / Collaboration Protocol Agreement

In a centralized approach, the steps and rules of the collaboration are pre-defined and the orchestration of the collaboration is performed by the server.

In the proposed peer-to-peer approach, there cannot be any presupposition that the two parties are aware of and will follow the same collaboration rules. In an initial “rendezvous” phase, the peers will negotiate and agree on the collaboration protocol (expressed in an open “language”), including the scope, time bounds, business information semantics, and determination of success or failure.

4.2.2 Collaboration Orchestration

In a centralized approach the collaboration orchestration is the responsibility of the server.

In the proposed peer-to-peer approach, the progress and state of the collaboration must be maintained (“mirrored”) by all peers. A synchronization and constant agreement on what is the current state of the collaboration must be enforced throughout all collaborating entities.

The peers should be able to detect the opening, transfer of control, completion or failure of transactions, and update the collaboration state appropriately, based on the collaboration protocol and rules, including the ability to define, detect and handle timeouts and exceptions. A roll-back functionality should also be provided for failed transactions, including multi-party transactions that fail when a member of the population leaves the peer-to-peer network.

4.3 Transaction level requirements.

In a centralized approach, the server is recognized as a neutral authority or mediator, responsible for providing the necessary accountability and non-repudiation characteristics. It is thus often entrusted the roles of logging, time-stamping and maintaining all transaction information, which may be used to resolve potential disputes or legally enforce collaboration outcomes, apart from guiding the collaboration workflow.

In the proposed peer-to-peer approach, the parties engaging in a transaction will usually have conflicting interests, so it is not feasible to assign this task to only one party. Current peer-to-peer technologies allow the tamper-resistant cryptographic distribution of data across various other parties, in a secure way that allows immediate recovery upon demand.

The inter-user communication and document exchange that takes place directly between the parties engaging in a transaction remains to a large degree unaffected in the peer-to-peer approach.

4.4 Post-collaboration Requirements.

It may be desirable (or required) for some additional procedures or events to be carried out after the completion of a collaboration, including post-transaction user ranking or reputation information management and dispersion. In a centralized approach, the above tasks can be performed by the server, where the results are also stored.

In the proposed peer-to-peer approach, a variety of distributed reputation management mechanisms that have been designed and deployed can be utilized.

4.5 Generic non-functional Requirements.

Non-functional requirements for supporting business transactions include integrity and confidentiality for documents, user data and transaction data (whether stored or in transit between peers), availability of services and information, fault resilience and a guarantee of some degree of quality of service.

Peer-to-peer systems have evolved considerably in these respects and currently offer solutions including various encryption techniques, replication and caching for availability and performance and secure routing algorithms. An additional potential requirement is the provision of user anonymity, which has also been dealt with in peer-to-peer environment.

5. A PRELIMINARY DESIGN PROPOSAL

Based on the requirements analysis presented above, and on the inherent architectural components of any purely distributed peer-to-peer system, we present in Figure 2 an abstract, preliminary design proposal. The

figure illustrates the various components of a node of the proposed peer-to-peer e-business transaction processing system. For most of the components we also suggest some possible design/implementation approaches, as presented in the recent literature by other researchers. All references in the figure can be found in [Androusellis-Theotokis and Spinellis, 2004].

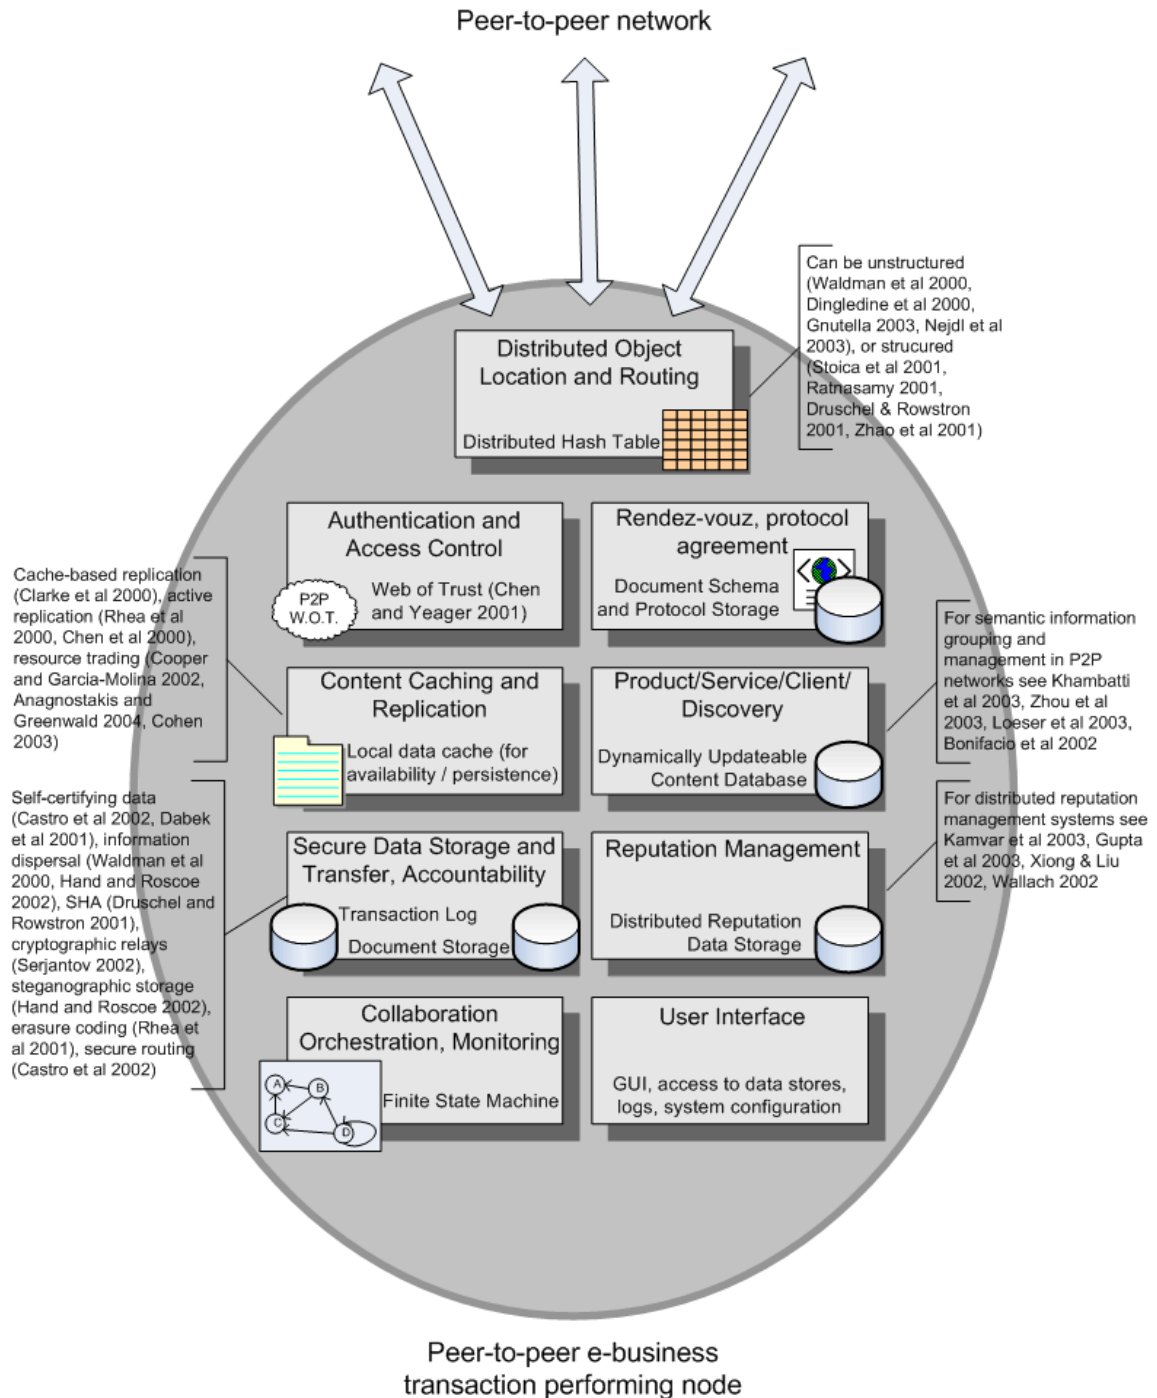


Figure 2: An abstract design of a peer-to-peer e-business transaction processing system. The components of a node of the peer-to-peer network are shown

Our proposed system is based on processes that depend on documents and services. Some of the documents will be of a legally binding nature, and must therefore be “issued by”, or registered with elected organisations (e.g. the Department of Finance etc.). This illustrates the potential need for some of the peers in the network to server more specific functions.

Figure 2 also emphasizes the desired modularity of the peer node design. Each box reflects a set of high level services that will be offered by the peer-to-peer network. Some of these services will cooperate with each other, providing a nesting functionality when necessary e.g. the *product/service/client discovery* may need as a prerequisite the *reputation management* service, in order to filter out the result set; and most of the services will obviously rely on the location and routing algorithm. A hierarchical ontological description of the services and their dependencies can be created and updated as new services are provided to the system.

Finally, in terms of an implementation platform, we are currently considering the JXTA peer-to-peer framework [JXTA].

6. CONCLUSIONS

It is our firm belief that the peer-to-peer model lends itself particularly well to the design of decentralized business transaction processing systems. We have attempted to provide an insight into what is currently lacking in this direction, namely a rigorous definition of the potential requirements of such a system, and a view of the degree to which these requirements are satisfied by current “state-of-the-art” peer-to-peer technologies. Our conclusion is that peer-to-peer systems have evolved and matured considerably and are capable of providing practical and satisfactory solutions to many of the problems that arise in the conversion from a centralized to a de-centralized system design.

ACKNOWLEDGMENT

This work has been partially funded by the General Secretariat for Research and Technology of the Greek Ministry of Development, through the “PRAXIS” project of the Information Society Programme.

REFERENCES

- Androutsellis-Theotokis, S. and Spinellis, D., 2004. A Survey of Peer-to-Peer Content Distribution Technologies. *ACM Computing Surveys (to appear)*, 2004.
- Charalabidis, Y. and Karakoidas, V. and Androutsellis-Theotokis, S. and Spinellis, D., 2004. Enabling B2B Transactions over the Internet through Application Interconnection: The PRAXIS Project. In *Proceedings of the e-Challenges 2004 Conference*, Vienna, Austria, October 2004.
- Choi, SY. and Whinston, AB., 2001. Is P2P Right for B2B. In *Cisco IQ Magazine*, March/April 2001. Online at http://business.cisco.com/prod/tree.taf%3Fasset_id=49569&public_view=true&kbns=1.html
- ebXML, 2001. UN/CEFACT and OASIS. ebXML Business Process Specification Schema v1.01. Online at <http://www.ebxml.org/specs/ebBPSS.pdf>
- Gao, R., 2004. Project Venezia-Gondola, A framework for P-Commerce. In *P2Pjournal*, June 10th 2004. Online at <http://venezia-gondola.jxta.org/>
- ISO/IEC 14662: 1997. Information Technologies - Open-EDI reference model.
- JXTA. The project JXTA web site. Online at <http://www.jxta.org>.
- Khambatti, M. and Dasgupta, P. and Ryu, K., 2004. "A Role-Based Trust Model for P2P Communities and Dynamic Coalitions". in *Proceedings of the IEEE International Information Assurance Workshop*, Charlotte, NC, April 2004.
- Schoder, D., 2004. Suitability of P2P for Business Transactions. In *Proceedings of the Peer-to-Peer Systems and Applications Dagstuhl Seminar*, March 2004. Online at <http://net.informatik.uni-tuebingen.de/PS/events/>
- White, A. and Peterson, K. and Lheureux, B., 2003. New P2P Solutions Will Redefine the B2B Supply Chain. Gartner Research Note. Online at http://www4.gartner.com/DisplayDocument?doc_cd=112909