An Infrastructure for Meta-Auctions

C. Bornhövd, M. Cilia*, C. Liebig, A. Buchmann
Dept. of CS, Darmstadt University of Technology
Darmstadt, Germany, D-64283
{*bornhoev, cilia, chris, buchmann*}@*informatik.tu-darmstadt.de*

Abstract

Auctions have been a popular trading paradigm for centuries but have gained new interest through world-wide trading on the Internet. Many B2C sites are embracing reverse auctions as an additional service to registered customers. In all these cases the efficient notification of the participants is essential. In this paper we develop the notion of a meta-auction that allows a potential buyer to roam automatically across auction sites and we identify critical communication and notification requirements of the next generation of Internet-scale trading systems: First, today's information systems are limited in their growth and interaction potential because the typical client-server and n-tier system architectures are solely based on a request/response interaction; second, the user-initiated query metaphor from the database domain is the primary means for information acquisition; and third, many assumptions about the meaning of data and notifications provided and exchanged through the Internet are left implicit.

We argue that Internet-scale business applications require publish/subscribe as an additional interaction paradigm, should leverage proactive information dissemination and caching mechanisms, and that there is a compelling need for metadata-based infrastructures providing common vocabularies for semantically meaningful exchange of data and notifications. We illustrate these points through examples from the auction domain and the development of the meta-auction concept.

1. Introduction

Auctions are a popular trading mechanism when multiple buyers compete for scarce resources. Famous auction houses, such as Sotheby's or Christie's for art or high-priced collectibles come immediately to mind. The advent of auction sites on the Internet, such as eBay, Yahoo or ricardo.de has popularized the auction paradigm and has made it accessible to a broad public that can trade in a consumer to consumer interaction anything from beeney-babies to electronics and from comics to vintage fountain pens. The mechanism has become so popular that many e-businesses are now offering a reverse auction mechanism as an additional service to their registered customers.

Serious art collectors have used similar services for centuries. Agents or gallery owners notify a potential buyer whenever an article that might interest a customer becomes available. In the more mundane world of Internet-auctions collectors would like to enjoy a similar service that would alert them whenever a certain object comes on the market. In addition, a collector might prefer to deal with one common auction portal instead of registering her interests with multiple auction sites. Therefore, we introduce the notion of a meta-auction. A meta-auction allows a potential buyer to roam automatically and seamlessly across auction sites for auctions and items of interest.

To realize the meta-auction several problems must be solved. We argue that today's systems that are based primarily on user-initiated communication are not adequate and will not scale properly. The large number of interconnected users and systems, as well as their wide-area distribution imposes particular restrictions with respect to response times and network bandwidth. Internet-scale information systems therefore must leverage proactive information dissemination and caching techniques. However, typical client/server and n-tier system architectures are merely based on a request/response interaction and do not take into account the asymmetric nature of such systems [1, 16], where the significant data flow is from a (database) backend-tier which stores operational data – such as items to be sold – to the application-tier which then provides access for end-users through a Web gateway.

Furthermore, the query metaphor from the database domain is currently the primary means for information acquisition, which results in the user polling for changes and happenings of interest. We argue, that notifications about events, such as the placement of a highest bid, and their timely delivery to the user represent valuable information. Therefore, publish-subscribe as an additional interaction paradigm is needed to make the efficient dissemination of process related information possible.

* Also ISISTAN, Faculty of Sciences, UNICEN, Tandil, Argentina.

Each site participating in the meta-auction system provides information about items and the auction process but does not share a global data schema nor may we assume a global schema for notifications. Still, all participants come from the same application domain and at least conceptually, share a common – domain-specific but participant-independent – vocabulary. While in most of today's systems the vocabulary is left implicit, we propose an ontology-based infrastructure for explicit metadata-management on top of which the meta-auction service can be realized. The suggested ontology-based infrastructure provides common vocabularies for semantically meaningful exchange of data and notifications, and supports incremental integration of participating information systems as needed.

In Section 2 we describe the mechanics of a typical auction, and introduce the meta-auction service. Section 3 describes the architecture of a meta-auction site. In Section 4 we address the semantic issues and present an ontology-based framework for creating and managing the needed common vocabulary. Section 5 puts it all together. Section 6 addresses related work while Section 7 summarizes and provides an outlook.

2. Meta-Auctions

Person-to-Person online-auctions, as provided by services like eBay, ricardo.de, and Yahoo look all very similar. Typical entities involved are Seller, Item, Bidder and Bid. *Items* represent those objects that are on sale, each of them is described using a unique id, a title, a description and a picture. Items are classified using a hierarchical *category* organization, like those depicted in Figure 1.

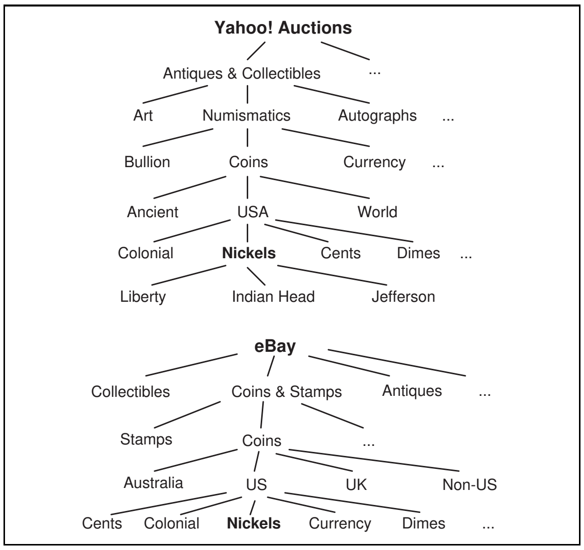

Figure 1. Item categories.

A typical use case for an auction site consists of the following steps: The user searches or browses the categories looking for an item of interest. Once the item is found she can track the auction process by continuously polling the auction site's server. Notice that reaching the deadline of the auction could take some days. In an ascending-price or English auction bidders increase the price during the course of the auction until its deadline. When the auction deadline is reached, the bidder whose bid is the highest "wins". If the minimum price has not been reached, the item might not be sold and the bidder must be notified.

Now consider the case of a collector. With the current auction sites, she has to manually search for the item of interest, possibly visiting more than one auction site. If successful, she might end up being engaged in different auctions at multiple auction sites. There are two obvious shortcomings to this approach: First, the user must poll for new information and might miss the window of opportunity, and second, the user must handle different auction sites with different category setups and different handlings. This motivates the need for the meta-auction broker.

The meta-auction broker provides a unified view of different auction sites and services for category browsing, item search, auction participation and auction tracking, as depicted in Figure 2.

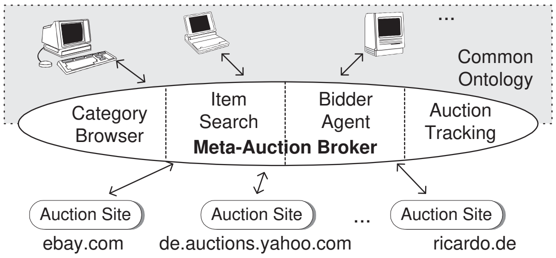

Figure 2. Meta-Auction Broker.

The realization of such a broker service is a major challenge, due to the large number of users, the expected wide-area distribution of the cooperating information systems and the dynamic nature of interconnected systems. Web caching techniques alone are not adequate, as most of the operational data is kept in backend-data stores. Web server content is generated by the application gateways and is thus frequently changing. Therefore, site replication using proactive information dissemination is a better approach.

While user interactions such as placing an offer, placing a bid, or browsing categories are well supported by the request/reply nature of the Web. The above scenario reveals a fundamental weakness of today's WWW infrastructure, namely, the missing support for a publish/subscribe paradigm and timely notifications. In fact, events that arise in the context of an auction process should be treated as first class information and propagated as notifications to the users who subscribed to the event. Propagation of events leads to a useful and efficient non-polling realization of an auction tracking service.

Finally, we identify the need to cope with heterogeneity and the goal to provide a unified view as well as unified