

Research Opportunities in the Virtual Agora: Market Aspects of Open Resource Coalitions

Position Paper for EDSER-1

Mary Shaw

Computer Science Department
Carnegie Mellon University
Pittsburgh PA 15213

March 17, 1999

Mary.Shaw@cs.cmu.edu

Voice: 412-268-2589

Fax: 412-268-5576

March 17, 1999

Abstract

For at least a decade, we have looked forward to an open market for software components and services. Recent advances in networking have created in the Internet a fertile environment for the creation and distribution of components, and we see a rich potential for value-added as well as primary goods. However, the economic infrastructure for an open market in software resources is still lacking, and support for automation of the marketplace is scant indeed. The prospect of automatically identifying, procuring, and integrating network-based resources intensifies the need to provide a robust economic infrastructure. We sketch the requirements for an electronic marketplace in software goods and services, explores some of the remaining shortfalls, and identify research opportunities.

Keywords: software marketplace, distributed open systems, resource coalitions

Research Opportunities in the Virtual Agora: Market Aspects of Open Resource Coalitions

Mary Shaw

Computer Science Department
Carnegie Mellon University
Pittsburgh PA 15213

March 17, 1999

Abstract

For at least a decade, we have looked forward to an open market for software components and services. Recent advances in networking have created in the Internet a fertile environment for the creation and distribution of components, and we see a rich potential for value-added as well as primary goods. However, the economic infrastructure for an open market in software resources is still lacking, and support for automation of the marketplace is scant indeed. The prospect of automatically identifying, procuring, and integrating network-based resources intensifies the need to provide a robust economic infrastructure. We sketch the requirements for an electronic marketplace in software goods and services, explores some of the remaining shortfalls, and identify research opportunities.

Introduction

Over a decade ago, Miller and Drexler [MiDr88] proposed a framework, *agoric systems*, for applying market mechanisms to software. They viewed economics as the study of the consequences of pursuing goals within the constraints of limited knowledge and resources. They saw market mechanisms as a means of mediating what even then were increasingly decentralized elements of computation. They defined an agoric system as a software system using market mechanisms, based on foundations that provide for encapsulation and communication of information, access and resources among objects. They anticipated the need for software agents, the utility of sufficient specification information to support automatic product comparison, the central role of a currency, and other issues that remain current even today.

Technology has change dramatically over the past decade. Some elements of an open market are in place, though most still require direct oversight and manual confirmation. (Would you turn loose a software agent to locate candidates for a \$5000 software component? Of course. Would you turn loose a software agent to *purchase* that component unattended? Almost certainly not.)

Traditional software systems are closed— they are controlled by single institutions. The constituent components may be acquired from external sources, but when incorporated in a system they come under control of the system designer. Analysis of these systems uses traditional closed-system models for correctness that emphasize a priori reasoning and analysis.

Successful large systems can also emerge without central control. Widespread use of the Internet is enabling a fundamentally new approach to software development: *open resource coalitions*, in which computing is carried out through dynamically formed, task-specific, coalitions of distributed autonomous resources [Shaw99]. The resources may be information, calculation, communication, control, or services. Unlike traditional software systems, which have a central point of direct control over the incorporated resources, these coalitions bring together collections, often ad hoc, of resources that are independently created and managed: The proprietors of the resources can modify, and remove them at will; the resources may not be specifically aware of the

way they are being used, or even whether they are being used; the selection and composition is likely to be done anew for each task as resources appear, change and disappear.

Economic Infrastructure for the Virtual Agora

The classical Greek agora was an open area that hosted a marketplace together with public gatherings and open discussions. Open resource coalitions will require tools with new degrees of autonomy and automation in order to identify, procure, compose, and track the resources. A virtual agora must encourage creation of diverse online resources; provide automated evaluation, procurement, and integration of the resources to form a coalition; and support the subsequent operation of the coalition. Each of these tasks offers challenges (and hence research opportunities) for the economic infrastructure.

Encouraging creation of resources

Resources of interest include primary information, computation in the form of proprietary models or of applications that can be invoked with real-time data feeds, control that coordinates the use of resources, and services such as secondary (derived) information, simulation, editorial selection, or evaluation results. An economically viable reward structure will be required to establish a rich population of available resources. Notwithstanding the success of social rewards in the open source software community, much or most of the agora is likely to be based on an explicit currency. In conventional markets, many of these added values are packaged with the original product and priced as markups (e.g. the retail markup covers selection of products as well as operating the store); others are provided by long-term subscription. When selection can be based on the automatic synthesis of multiple individually-selected sources, there must be a means of compensating the sources incrementally.

Currently most software is sold in discrete transactions of sufficient value to justify ordinary credit mechanisms. Some documents and services are provided by subscription or by micro-debits against an established account.

Research opportunity: Micro-accounting and innovative models for electronic payment

In a market where value encompasses not only products but also information about those products, the cost of a resource reflects contributions from many parties and may be different for each purchaser. The virtual agora needs the ability to charge according to different models from current norms. This requires some sort of micro-accounting. It also implies authenticated components/properties, agents, third-party sources of metainformation, and reasonable measures for protecting and rewarding intellectual property.

A viable micro-accounting system should provide a fertile climate for third-party value-added services. These services might provide editorial services in selecting information, evaluation services for resources, brokering of needs and resources, and so on. Storage, bandwidth, and compute cycles are free or cheap. Value now resides in information, search/retrieval, computation, editorial services, evaluations, oversight, credibility, and requirements.

Evaluating resources

A major current stumbling block to an automated open market is the severe lack of trustworthy information about available resources. It's hard enough to discover automatically what the functionality and packaging of a resource may be, let alone to gather enough information to compare the quality or price of multiple candidates. Even though perfect information is not achievable in practice, better information about resources can lead to better choices. Such metainformation must not only exist, it must also be organized in an accessible form, and it must

include information about the credibility of sources [Shaw96]. Further, the market for the evaluation information is somewhat decoupled from the market for the primary resource, because different purchasers care about different properties and about the same property to different degrees.

Metafiles in current Geographic Information Systems provide rich information about the content, including not only essential information such as format but also information about sources and intermediate transformations. Search services for on-line shoppers currently provide little more than text search on the web pages that advertise the products. A few third-party value-added services exist, such as the certificate services for software downloads from Verisign (which only provide 1 bit of authenticity information).

Research opportunity: Third party value-added services

Consider the possibility of setting up an on-line certification business. It could act in cooperation with vendors of resources (as Underwriters Laboratory does) or independently (as Consumers Union does). It would certify specific properties of resources, which it would validate by some combination of test, analysis, experiment, audit, monitoring, and other techniques. It would report these properties, perhaps in response to direct queries, perhaps by attaching certifications indelibly to the product. What is a plausible business model for such a service? Who would pay for the information, and how much? How would it deploy its internal resources among validation techniques [BCJRS99]? How would it market, especially to automatic tools? How would it deliver its results? What assurances would it provide about the results, and would it provide guarantees?

Procuring and integrating resources to form a coalition

Integrating diverse independent resources poses a significant technical challenge. Validating the result presents an even greater challenge, requiring new models of adequacy (“Is this good enough for the intended task?”) rather than classical full correctness. This is truly a matter of drawing conclusions under limitations of knowledge and resources.

Traditional closed-form models for correctness must be replaced with models for “good-enough”-ness. There are human limitations on the correctness of requirements, especially on understanding the full generality and implications of requirements and solutions. Specifications can be no better, so it’s not realistic to expect large systems to be fully verified. It follows that complex systems will require oversight, checking, and recovery. In resource coalitions, where the individual resources may change or even disappear, the problem is particularly intense. We need a new criterion to replace absolute correctness— a criterion that allows us to bring together various sources of information to establish that a system is sufficiently good for the task at hand, and to do this analysis at a reasonable cost. Since we derive information about the system incrementally, this criterion must be incremental, progressive, and approximate.

Research opportunity: Reasoning about software under uncertainty

What’s the right way to think about whether a resource coalition is good enough for its intended purpose? Consider the matter of finding or of purchasing the \$5000 component. Certainly the level of assurance required for the former is less than for the latter. Analysis similar to that required to underwrite insurance (see below) may help decide how much to invest in validation. However, many validation techniques are available. Portfolio analysis may help to decide which combination of techniques to use, and in what proportion [BCJRS99].

Operation of a coalition

One of the current impediments to software reuse of any kind is assignment of responsibility and liability. The typical software “warranty” warrants little more than the readability of the disc that contains the software, and the typical limit of liability is the purchase price. That additional uncertainties associated with using independently-managed resources in an open market would seem to present major obstacles to such use unless the liability can be reassigned or otherwise managed. The conventional means of doing this is through contracts (to assign and reward responsibility) and insurance (to mitigate risk).

Research opportunity: Risk reduction by underwriting: (warranties and insurance)

The insurance industry has centuries of experience with quantifying risk. The most familiar forms of insurance depend on actuarial tables derived from experience. However, underwriters also have techniques for pricing special-risk insurance for unique or improbably events. Suppose that we are able to find means of reasoning about sufficient correctness, and suppose that those results include measures of confidence and risk. How could we use that (admittedly unreliable) information as a basis for underwriting software operations warranties and insurance?

Acknowledgements

These views have emerged over several years and crystallized during my recent term as a Fellow of the Center for Innovation in Learning at Carnegie Mellon. Previous discussions with colleagues in the Carnegie Mellon School of Computer Science, especially David Garlan and other members of the Composable Software Group, have laid the groundwork. Vic Vyssotsky encouraged me to think about large software as more like city planning than like product manufacture. Mark Miller started me thinking about agoric systems about 10 years ago.

Bibliography

[BCJRS99] S. Butler, P. Chalasani, S. Jha, O. Raz, M. Shaw. “The Potential of Portfolio Analysis in Guiding Software Decisions.” Position paper submitted to First Workshop on Economics-Driven Software Engineering Research, March 1999.

[MiDr88] Mark S. Miller and K. Eric Drexler. “Markets and Computation: Agoric Open Systems”. In B.A. Huberman (ed), *The Ecology of Computation*, Elsevier Science Publishers, 1988.

[Shaw96] Mary Shaw. “Truth vs Knowledge: The Difference Between What a Component Does and What We Know It Does” *Proc. 8th International Workshop on Software Specification and Design*, March 1996.

[Shaw99] Mary Shaw. “Architectural Requirements for Computing with Coalitions of Resources.” Position paper for WICSA1: First Working IFIP Conference on Software Architecture. February 1999.

http://www.cs.cmu.edu/afs/cs.cmu.edu/project/vit/www/paper_abstracts/Shaw-Coalitions.html