Opportunities for E-commerce in Networking

David G. Messerschmitt
Department of Electrical Engineering and Computer Sciences
University of California at Berkeley (USA)

Jean-Pierre Hubaux
Institut pour les Communications Informatiques et leurs Applications (ICA)
Ecole Polytechnique Fédérale de Lausanne (Switzerland)

****

Abstract

There are numerous opportunities to apply e-commerce technologies to networking. These include the assembly, pricing and payments for complementary infrastructure resources, and the selection of and payment for value-added collaboration and information access services. E-commerce can support the separate provision and coordination of these elements, or allow them to be bundled by a customer-care organization. These opportunities and options are discussed.

1. Introduction

Rudimentary examples of e-commerce in networking include the "dial around" services to make per-call choices of competitive carriers and Web-based clearinghouses for excess capacity[1]. Opportunities for more sophisticated uses of e-commerce—particularly when terminals harbor intelligence and a rich user-interface—are discussed in this paper.

2. Some Opportunities

There are many motivations for e-commerce in networking. The consumer faces a choice of providers, and two or more providers may be needed for end-to-end connectivity. Networks will provide flexibility in quality-of-service (QoS) attributes (like rate and delay) and a great diversity in end-user functionality, like information access (e.g. the Web), e-commerce transactions (e.g. stock purchases), and collaboration (e.g. remote conferencing and collaborative authoring). These may be bundled with infrastructure services (with added revenues for providers) or purchased separately (requiring coordination). Increasingly terminals will be nomadic (and even mobile) with a location-dependent provider. E-commerce can match consumers with suppliers and providers, compare competitive options, coordinate complementary purchases, and make payments.

---

[1] Examples include Arbinet (www.arbinet.com), Band-X (www.band-x.com), and RateXchange (www.ratexchange.com).

Internet telephony client
software application
Web browser

PTN telephony
Remote conferencing
Information access

User functionality

- -

Infrastructure

Terminal equipment
Operating system

Bearer service
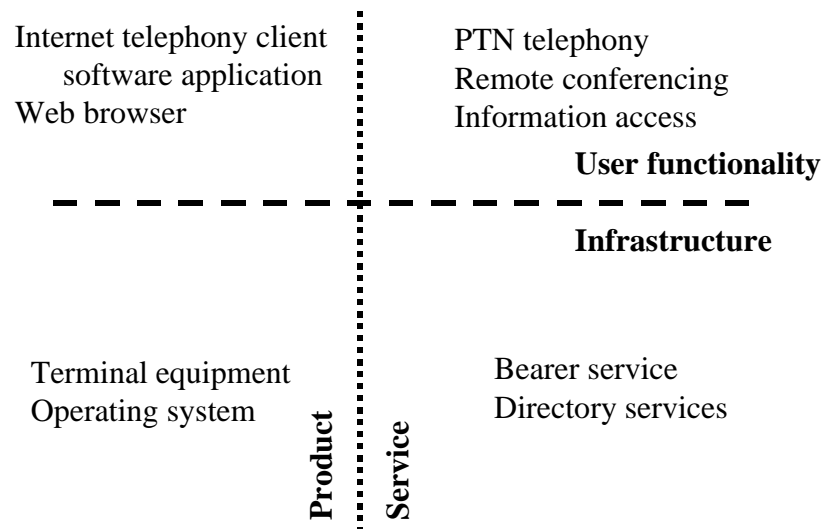Directory services

Product **Service**

**Figure 1. Two-dimensional taxonomy of goods
and services in networking, with examples.**

2. Some Terminology

Networking and computing use inconsistent terminology [1,2,3], so we define industry-neutral terminology. The *user* is an individual, a group, or an organization that purchases products and services. A taxonomy is shown in Figure 1. *Infrastructure* benefits a variety of uses, whereas *user functionality* leverages an infrastructure for specialized capabilities directly benefiting users. Either may be purchased as a *product* (operated by the user for his own benefit) or a *service* (operated by a *service provider* for the benefit of many simultaneous users). For example, telephony can be purchased as a product (software application supporting Internet telephony on a user-owned terminal) or a service (from a public telephone network service provider). An important distinction is that a service carries an explicit provider responsibility for quality and availability. The most fundamental infrastructure is *connectivity*, and a specific connectivity is *bearer service*, which includes QoS attributes such as availability, rate, loss, reliability, and delay.

A user may purchase separately and integrate complementary products and services, or may purchase a *bundle* (e.g. a terminal, bearer service, and stock quotes, all from an information provider), requiring behind-the-scenes coordination and revenue splitting. In a sealed-bid auction, either competitive sellers or potential buyers make bids, which are selected by the other.

3. E-Commerce for Infrastructure Services

Consumer E-Commerce

Traditionally users enter long-term contracts with infrastructure service providers at published prices. Past networks have provided a single or small set of service options, but future networks will offer very wide ranges of bit rates (from zero to hundreds or Mb/s and higher) and QoS options (such as differentiated services in the Internet). E-commerce could provide means for the user to explore possibilities, signal choices, and pay accordingly. One opportunity is for dynamic pricing of services per-session, rather than by long-term subscription.


Figure 2. Complementary networks are required to provision end-to-end infrastructure services.

Inter-Provider E-Commerce

End-to-end bearer service often requires complementary resources from different providers—for example, a wireless access and long-distance carrier as shown in Figure 2. Infrastructure providers may want to establish per-session QoS contracts and traffic attributes with users and configure internal network resources accordingly. For example, with future broadband access links, configuration and pricing based on rate attributes may be appropriate. (In the past, this has been approximated by fixed prices based on restricted access speed options.) Wireless access subnetworks with the flexibility required for a variety of user functionality may need configuration and pricing based on bit error rate, delay, and availability objectives.

Achieving a specified QoS requires coordination to configure resources, set an aggregate price, and split revenue. Traditionally this has been handled by static interconnect and settlement agreements. With a plethora of service options, such agreements may become

too cumbersome. Inter-provider e-commerce could substitute dynamic mechanisms for coordination and pricing, and, to keep sellers honest, competitive seller bidding to users.
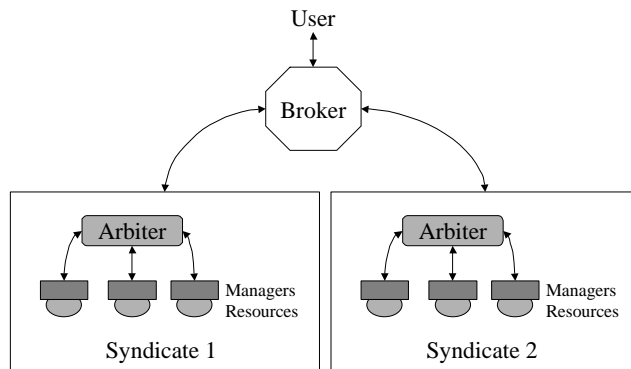

Figure 3. An approach to competitive sealed bidding for infrastructure services.

This might work as illustrated in Figure 3. The user designates a *broker* to obtain competitive bids and configure a service option. Upon a request from the broker, sets of providers controlling complementary resources (like subnetworks forming end-to-end connectivity) form a temporary *syndicate* to generate a bid and provision and charge if the bid is accepted. An *arbiter* designated by each syndicate coordinates individual *resource managers*, gathers price bids from the managers, and aggregates the results as a bid to the broker. For example, the broker may specify end-to-end delay objectives, and the arbiter coordinates complementary subnetworks to partition that delay among subnetworks and gather the resulting prices. The broker, arbiters, and resource managers may be software agents transported to a single host for negotiation in minimum time [4,5,6]. Competitive bidding is not essential to this model.

Scenarios similar to Figure 3 are an interesting topic for research. Each syndicate wants to minimize its price bid to maximize the chance of winning. If end-to-end QoS attributes must be disaggregated to individual subnetworks, minimizing price requires an additional negotiation among resource managers—for example, wireless access links and congested subnetworks would be allocated the greatest impairments. Mechanized negotiations of this type have been studied by economists.

Economic Considerations

An alternative to dynamic pricing is extensions to current fixed interconnect and settlement processes. The user may contract with a local provider (which, with nomadicity, may depend on location), who has predefined fixed-price contracts with complementary providers. E-commerce could allow pricing to be dynamic—at the expense

of increased (although hopefully modest) transactions costs—and also enable differentiated QoS guarantees with associated pricing. Providers could vary prices at will, based for example on requested QoS, current traffic conditions, and indications of user willingness to pay.

Transaction costs can be reduced by over-provisioning facilities to provide superb QoS to all users (although this is unlikely in wireless access networks). To evaluate these alternatives, transaction costs should be compared to the value added to both users and providers. They are worthwhile if that value is greater, as measured by user willingness to pay and/or economic advantages to the provider. This value proposition is now discussed, first for users and then for providers. The many issues raised here are complex and poorly understood—there is no "right" answer at the present state of knowledge [7].

Value to the User

QoS configuration accommodates differentiation by user functionality and by user preference. As an extreme example, telesurgery has more stringent requirements than email. If the network provisions differentiated QoS, associated pricing is inevitable, since otherwise a rational user would always choose the highest QoS option.

During congestion—primarily an issue limited to wireless access subnetworks in the future—resources must be rationed, and pricing is the usual mechanism for matching supply and demand. As a form of congestion control, willingness to pay voluntarily distinguishes compelling needs from those that can be deferred. All other forms of congestion control are involuntary—a user may be unable to obtain desired resources no matter how compelling the need. QoS also shifts risk from user to provider, increasing value and willingness to pay.

Competitive bidding for services increases competition and reduces price. It also obviates user lock-in to providers due to switching costs (which allows providers to increase prices [8]). Bidding also allows easier market entry, increasing competition.

Value to the Providers

For providers, dynamic optimization of resources based on network conditions and user demand would result in efficient resource allocation, for example by allocating greater impairments to congested facilities. Dynamic pricing would reserve resources for users with the highest willingness to pay, increasing revenue. Facilities would be more fully utilized—for example by incrementally reducing prices until filled—also increasing revenue. Congestion pricing is a source of revenue for needed facility expansions.

Networks have large economies of scale—high fixed costs of construction and low operational costs—which introduces special challenges. Absent congestion, the *marginal*

cost of provisioning a new service request is low. Competition tends to drive prices toward this low marginal cost, resulting in operating losses. There are also social costs: Scale economies create winner-take-all effects, as a dominant provider can undercut prices of others. Dynamic pricing supported by e-commerce offer some solutions to these dilemmas.

A key to dealing with scale economies is *value pricing*—tying prices to the user value (expressed by willingness to pay) rather than costs [8]. Willingness to pay varies widely, so value pricing requires *differential pricing*—charging different prices to different users. One established technique is *quality differentiation* or *versioning*, in which users self-select from among versions: Those with a high willingness to pay typically choose more expensive, high quality versions. Options and prices can be designed so users with a high willingness to pay are not tempted by cheaper versions. (These strategies are familiar in the airline industry, which experiences similar economic challenges.) QoS provides a natural dimension for versioning and value pricing strategies. Alternatively, buyer auctions—an extreme form of value pricing—are advantageous to the provider, since users directly signal willingness to pay.

Policy Issues

Versioning can address *universal service* (expanding service to a wide segment of the population) by a market mechanism. With fixed prices, providers tend to ignore users with less willingness to pay in the pursuit of maximum revenues. It is socially desirable for higher quality versions to be offered without precluding lower quality options for those able to pay less.

Inter-provider QoS provisioning is important to preserve competition in the industry, and may as a result be mandated. To see this, assume there are two service providers. Each user derives greater value when connected to a larger universe of users, including those of the other provider—this is called a *network externality*—so the providers will likely make interconnect arrangements. However, if QoS is more predictable staying within a single provider's facilities, users prefer the network with the larger user base, biasing the market toward a dominant provider.

4. E-Commerce for User Functionality Services

The greatest opportunities for provider profits are in value-added user functionality services.

Collaborative Services

Collaborative services, such as telephony, remote conferencing, and collaborative authoring [1,2], benefit most from QoS configuration. It is widely assumed (as described earlier) that users will control bearer service QoS directly. However, a user is interested in perceptual measures directly related to the user functionality, such as audio or video quality or interactive responsiveness.

It is natural for user functionality to present perceptual quality options and associated prices (for example, representative examples of video quality) and relate them to QoS and infrastructure service pricing [9]. It understands the contextual relationship between QoS and perceptual quality (for example, bit error rate and video quality for a specific coding) and can thus relate user and infrastructure concerns. It then subsumes the role of broker in Figure 3.
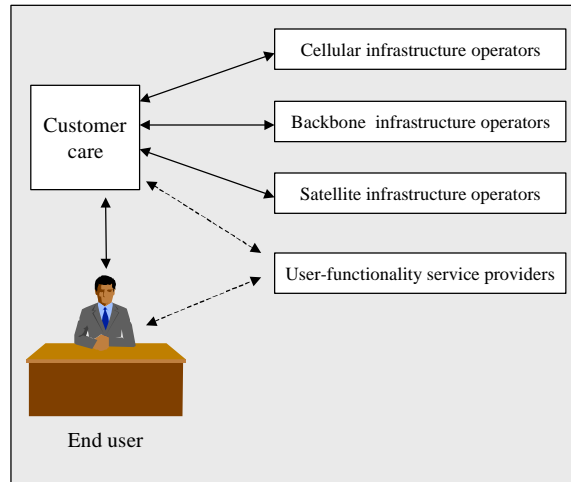
Information Services

A separate provider will frequently offer information services such as newspapers, financial information, and video on demand. Separate payment may be required for the infrastructure, the information service, and the information content. E-commerce can support the choice of providers and content and payments. This is particularly sensitive to transaction costs—the charge for individual pieces of information is normally small—and thus an interest in *micropayments*.

Bundling infrastructure and information is common today—for example, "900" service in the U.S.—but this is a legacy of the wireline telephone, where a terminal is permanently connected to a local provider with a subscription. A nomadic user may deal with many providers, making a bundle less attractive, and e-commerce can mitigate the resulting inconvenience. This is illustrated by the third-generation mobile system UMTS [10,11]. Tokens obtained from a UMTS infrastructure provider are passed from user to information service provider as proof of credit worthiness. The UMTS provider collects billing information from all information providers, aggregates them, and bills the user.

Figure 4: Connectivity is one service among many.

Examples of customer care functions:
- Trusted relationship with the end-user and with infrastructure providers
- One stop shopping
- Authentication and payment
- Optimized roaming agreements
- Integration of service mechanisms on a single smart card
- Selection and recommendation of user-functionality service providers
- Personal customization (language, physical impairment,…)
- National customization (regulations)
- Management of bids on behalf of the user
- Management of multiple-purchase transactions
- Conflict prevention and resolution

Customer care

Cellular infrastructure operators

Backbone infrastructure operators

Satellite infrastructure operators

User-functionality service providers

End user

Customer Care

As described, a user may obtain complementary products and services from a variety of sources, burdening him with the integration unless there is a third-party *customer care intermediary*. This organization is a single point of contact and responsibility, encompassing both infrastructure and user functionality, as shown in Figure 4. It could provide secure, evolvable and anonymous payment mechanisms [13], and subsume the role of the broker in Figure 3. It also addresses the proliferation of advanced user functionality and providers, much as credit card associations arose to intermediate among consumers, banks, and merchants. Finally, it relieves the user from worrying about different underlying network organizations and the hybrid character of services [12].

5. Other Issues

Exploiting a Common E-Commerce Infrastructure

Historically, infrastructure services have been provided by subscription with direct monthly aggregated billing. UMTS continues this tradition. E-commerce enables other business models such as immediate payment for services. Rather than each provider providing separate user billing, a common shared billing infrastructure, similar to the credit card associations, could be developed. This may be more flexible and cost-effective, and reduce credit risk to providers.

Authentication and Privacy

A necessary ingredient of e-commerce, user authentication to prevent fraud, has a major downside for users in the loss of privacy. Authentication allows providers to log user activities, and data warehousing and mining potentially allow tracing and aggregation across a variety of suppliers and providers. Without definitive privacy policies, the development of a dynamic market may be stifled. A customer care organization can preserve user anonymity and enforce privacy policies.

5. Conclusion

E-commerce can support both infrastructure and user-functionality services, bundled or unbundled. These opportunities enable greater choice and competition, hopefully without inconvenience or transaction costs. Research is needed to fully qualify and quantify these opportunities.

References

[1] D. Messerschmitt, *Networked Applications: A Guide to the New Computing Infrastructure*. Morgan Kaufmann, Jan.1999.

[2] D. Messerschmitt, *Understanding Networked Applications: A First Course*. Morgan Kaufmann, Sept. 1999.

[3] D. Messerschmitt, "The convergence of telecommunications and computing: What are the implications today?", *IEEE Proceedings*, Aug. 1996.

[4] L. Kan, "Agent-based QoS Price Negotiation and Resource Reservation Protocol", Masters Project, University of California at Berkeley, 1998.

[5] W. Li and D.G.Messerschmitt, "Mobile Agent-based Network Signaling for Resource Negotiations", *Workshop on Resource Allocation Problems in Multimedia Systems, IEEE Real-Time Systems Symposium,* December 1996, Washington, DC.

[6] W. Li, "Agent-Based Service Signalling and Negotiation", Ph.D. dissertation, University of California at Berkeley, 1999.

[7] L. McKnight and J. Bailey, "Internet Economics: When Constituencies Collide in Cyberspace", *IEEE Internet Computing*, Dec. 1997.

[8] C.Shapiro and H.Varian, *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press, 1998.

[9] H. Metz, "Perceptually Based Network QoS Configuration", Masters Project, University of California at Berkeley, 1999.

[10] K. Martin et al., "Secure Billing for Mobile Information Services in UMTS", *ISN'98,* available at http://www.esa.kuleuven.ac.be/cosic/aspect/papers/isn98.ps.

[11] G. Horn and B. Preneel, "Authentication and Payment in Future Mobile Systems", *ESORICS'98*, available at http://www.esa.kuleuven.ac.be/cosic/aspect/papers/esorics98.ps.

[12] J.P. Hubaux, D. Nagel, B. Kieburtz, Eds, "Feature Topic Issue on the Provision of Communication Services over Hybrid Networks", *IEEE Communications Magazine*, July 1999.

[13] L. Buttyan, J.-P. Hubaux, "Accountable Service Usage in Mobile Communication Systems", EPFL/SSC Technical Report SSC/1999/016, June 1999, available at http://sscwww.epfl.ch.

Biographies

David G. Messerschmitt

David G. Messerschmitt is the Roger A. Strauch Professor of Electrical Engineering and Computer Sciences (www.eecs.berkeley.edu) at the University of California at Berkeley. From 1992-96 he served as Chair of EECS, and prior to 1977 he was with Bell Laboratories in Holmdel, N.J. He received a Ph.D. from the University of Michigan, is a Fellow of the IEEE, a member of the National Academy of Engineering, and is the 1999 recipient of the IEEE Alexander Graham Bell Medal recognizing "exceptional contributions to the advancement of communication sciences and engineering". Messerschmitt's current research interests include wireless access to packet networks, network management, the role of mobile code in network infrastructure, the convergence of computing and communications, and the economics of networks.

Jean-Pierre Hubaux

Jean-Pierre Hubaux is a Professor at the Swiss Federal Institute of Technology - Lausanne (EPFL), which he joined in 1990. Previously he spent 10 years in France with Alcatel, where he was involved in R&D activities in the area of switching systems architecture and software. Hubaux is co-founder and co-director of the Institute for Computer Communications and Applications (ICA, icawww.epfl.ch). His research activity is focused on service engineering, with a special emphasis on multimedia services; more recently this focus was extended to security and mobile applications. He has authored and co-authored more than 40 publications in this area and holds several related patents. He defined the new communication systems curriculum at EPFL (sscwww.epfl.ch) and will chair the Communication Systems Division starting in October 1999. Hubaux is a senior member of IEEE.