Internet as a Business Place: Logistics and Challenges

Glen Kramer
University of California @ Davis
kramer@cs.ucdavis.edu

March 2001

1 Introduction

I work in R&D department in Alloptic, a company that designs and builds optical access networks. Our first product is capable of delivering hundreds of Megabytes of bandwidth to and from user. Future designs aim at Gigabit or tens of Gigabits speeds. However, the question I am asking myself for some time now is *what next*? Bandwidth will not be a scarce resource anymore. How network will evolve? What services will become available to users and what users will expect from the Internet?

In this study we will look at the current state of telecommunication industry and try to have an insight at what network evolution will bring us, what players will appear in the market and how business will be done over the Internet. We will look at what issues may arise in doing business over the Internet and what role cryptography may play in enabling e-commerce in an efficient and robust way.

In Section 2 we identify factors affecting network evolution, i.e., the push for new technologies and the driving forces (or absence of those) behind the creation of new services on the Internet. We also discuss what necessary level of infrastructure should be available to have an incentive to create new services.

Section 3 is a more futuristic section where we attempt to identify services that will be available and more important how users will interact with the Internet. We also identify four major institutions that will play role in e-commerce: Content or Service Providers (CSPs), Internet Service Providers (ISPs), Network Operators (NOs), and Customers.

Section 4 further discusses complex interactions between the above institutions and shows the need for security mechanisms if the Internet is to become a business place.

In Section 5 we discuss a possible protocol, which while being very simple, still provides the necessary guarantees for the issues we identify in Section 4.

Section 6 concludes this study.

2 Networking evolution

The year 2000 was the year when most network operators agreed that the data traffic in their public networks exceeded the voice traffic. While this event passed unnoticed by general population, the network operators realized that now their networks designed and build to carry voice traffic are not scaleable and not efficient for the data traffic. The burstiness of data traffic warrants the bandwidth over-provisioning and, therefore, very low utilization.

The major problem, however, is that about 90% of revenues of network carriers is generated by voice traffic, mostly by long-distance services. There are well-established billing practices for voice telephone services that users are used to and accept.

On the other hand, data access is usually priced at a fixed rate. There are no mechanisms to do a usage-based billing and most of the services that customers get over the Internet are free. Since the bandwidth getting cheaper and cheaper, the revenues the companies get from data traffic are decreasing. Correspondingly, many carriers complain that to accommodate the increased traffic on their networks they have to upgrade their networks often, and thus their capital expenses increase, but the revenue remains flat or even decrease.

And as if that was not enough, there are considerable technological advances in allowing new Competitive Local Exchange Carriers (CLECs) to carry voice as IP traffic, thus bypassing the Public Switched Telephone network (PSTN) with its billing system. Instead of going to a telephone switch, the voice traffic is being carried by a DSL or cable modem line to an IP router. Then it is routed through the backbone, finally reaching the router closest to the recipient and enters phone network to complete the call. Alternatively, call may originate as local call, be converted to IP traffic, traverse half the globe as data, and finally be converted to a local call in the country of destination. Thus, the call will be charged as local call (most often it will be free) instead of a long-distance toll. All that has negative impact on revenues generated by carriers.

With it comes the painful realization that neither selling raw bandwidth, nor selling telephone services is a profitable business anymore.

What can generate profits are new services that will become popular and ubiquitous. However, developers are reluctant to create such new applications for the simple reason that there is no proper network infrastructure to support the new services. Such infrastructure should be able to provide delay and delay variation guarantees, protection and restoration schemes, be scalable, but most important, it should provide many-fold higher bandwidth than what is available today [1]. Figure 1 presents new service types that will become available with growth of bandwidth available to a user.

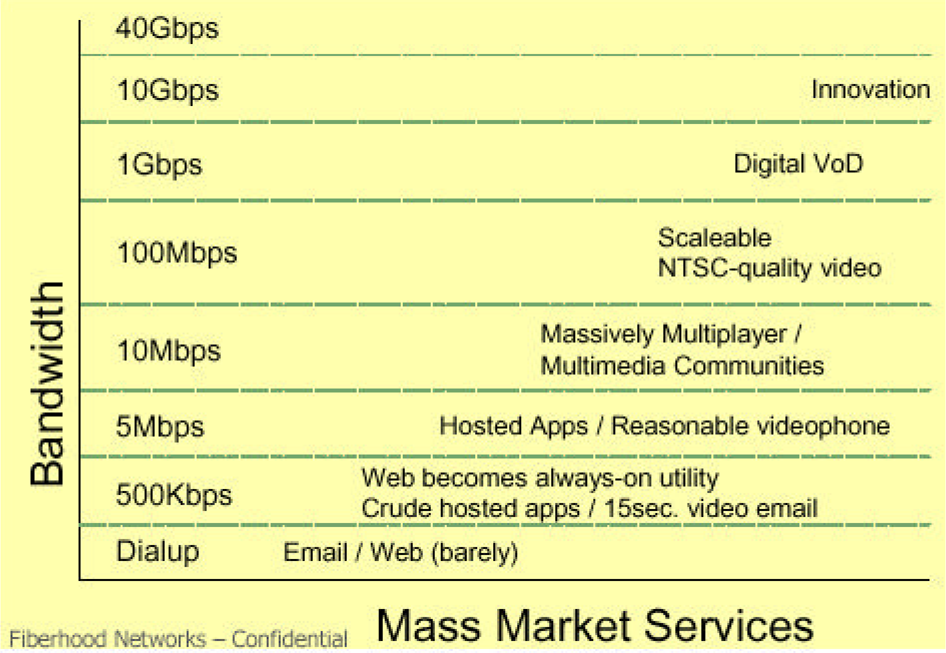

Figure 1. Per-user bandwidth requirements for new services
(with permission by Fiberhood Networks.)

But network operators are not rushing to upgrade the networks. Without guaranteed return on investments there are no economic reasons to upgrade the network. And, of course, customers will not pay for higher bandwidth if there is no use for it. That all creates a vicious circle in which telecommunication industry finds itself today (see Figure 2.)

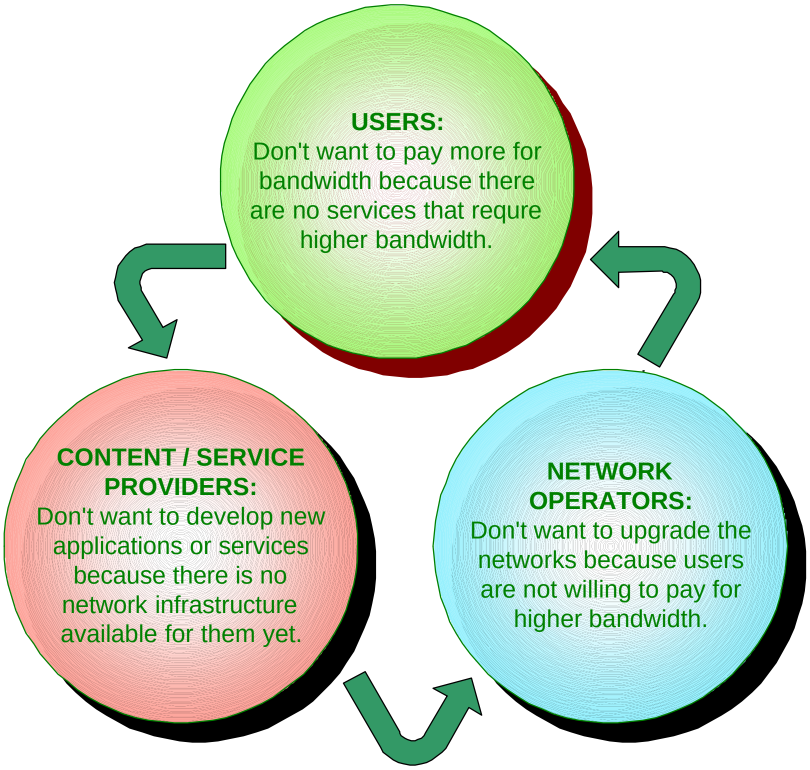
USERS:
Don't want to pay more for bandwidth because there are no services that requre higher bandwidth.
CONTENT / SERVICE PROVIDERS:
Don't want to develop new applications or services because there is no network infrastructure available for them yet.
NETWORK OPERATORS:
Don't want to upgrade the networks because users are not willing to pay for higher bandwidth.

Figure 2. Vicious circle.

How can we break that circle? It looks like the network operators have most of all to lose and thus will be interested in changing things. It is worth noting that things have been changing: bandwidth capacity of the backbone has increased tremendously over the last decade. But one can foresee even more backbone growth considering the recent advances such as dynamic wavelength routing, optical packet switching, etc.

However little has changed in the local exchange. “Last mile” still remains the bottleneck between a high capacity LAN or home network and a backbone. Part of the reason why the residential access remains bottleneck is that CLECs don’t have “right-of-way” – the right to dig the ground and lay down new cables. ILECs, on the other hand already have infrastructure in place and don’t want to upgrade it without ability to get additional returns on investments or even cover the cost of upgrades. And those upgrades can be costly: considering bit rates and the distances from the switching office to the residences, it becomes clear that only optical fiber with its huge bandwidth capacity would be the right medium for residential networks.

2.1 Optical access – Bandwidth of Plenty

There are however, different players emerging in the local exchange market: power utility companies. Power companies have been known to lay down fiber cables along the power cables in the long haul. That fiber plant is later leased to network operators as a dark fiber. Recent advances in fiber manufacturing and abundance of silica (sand) from which fiber is made resulted in fiber being cheaper than copper. According to [2] most of the cost of laying new fiber cables is labor cost. But power companies are undertaking this cost anyway, so why not to lay down few strands of fiber along the cable to each new residence? They may lease it to local exchange carriers and have additional revenue without considerable up-front capital expenses.

You ask why they were not doing that before, with copper cables? The reason is that electro-magnetic fields from power cables would induce an unacceptably high noise in the copper. And besides, it is not safe to have two conductors close to each other (flooding and cable aging may result in high voltage leaking into phone lines). Fiber, on the other hand is made of dielectric material. It is tolerant to the electromagnetic interference and it is safe to use along the power cable. Similarly, fiber can be attached to an air-dropped power cable or even manufactured as a hybrid power-fiber cable.

There is one more reason why power companies should be interested in having their users "wired", or at least to have their electricity meters connected to the net. That reason is truck-roll cost – the necessity to send a person to each power meter once a month to record the amount of Watts spent. Consider the mistakes and necessary staff to resolve them and it adds up to a very considerable amount. On-line power meter, on the other hand, may be queried remotely. It will report its current reading and the bill will be issued – all automatically.

Additional benefit for power companies, as well as users, in having on-line power meters is rate-based billing. California power crisis suggests that there may be a need for power utilities to adopt a new billing model similar to one used by phone companies: a Watt of energy during peak-hour should not cost users the same as Watt of energy during the night. Again, that is easy to accomplish with power meters always on-line.

The power utility company will probably sell or lease the fiber to a network operator.

3 Ultimate goal – Internet as a business place

Let us for a moment fantasize and imagine what services would be available to customers if bandwidth was not a limiting factor.

A. Videophones take place of a most widely used communication medium. Different connection options will be aimed at various markets: one-to-one connection is a simple video-call. Many-to-many connection is a video-conferencing tool employed mostly by businesses but available to anyone as well. One-to-many connections will be used by colleges and various courses as a distance-learning tool.

B. Newscast on demand. Many independent reporters will bring the latest news from around the world. It is different story how to choose one.

C. Entertainment. The notion of TV schedule or primetime would become obsolete. All programs are available anytime – choose what you like and when you like it. Users will pay only for programs thay watched and will also choose the amount of advertising they might tolerate for the benefit of lower fees. What is more important, users themselves will be able to select the topic of advertisements they interested in. Someone is shopping for a new car. Someone is planning a vacation. So why they both should stare at a mattress sale commercial? Blockbuster would love to be able not only rent their videocassettes online, but to get rid of them and its more than 7000 stores altogether. Instead it will stream the movies directly to the customers from its video servers. No more hassles for Blockbuster to guarantee the availability of new releases, no more hassles for users to return the cassettes back before the midnight. That is a service: rent any movie, anytime, anywhere!

The major advantage that we all are yet to comprehend is that the information dissemination will migrate from push technology to pull technology. This is the one thing that will have the most impact on the way we work, communicate, and spent our leisure time. Broadly we can differentiate the two technologies by looking at who initiates the transfer of information. If the transfer is initiated by the information originator, it is a push technology (information is being pushed at recipient). If a user asked for information this is pull technology. The push technology by its very nature cannot be always relevant or timely. Various schemes were developed to improve the push technology. Some schemes used sophisticated user profiling engines in an attempt to disseminate the information more selectively. None of them proved very successful either due to privacy concerns or just simple inability to track the extreme complexity and dynamic nature of user-network interaction.

We can foresee four major groups of players in the new e-marketplace: Content or Service Providers (CSPs), Internet Service Providers (ISPs), Network Operators (NOs), and the driving force of it all – Customers. Below we explain the specific roles each group plays in new Internet.

3.1 Content and Service Providers (CSP)

It is obvious that only content and services can generate revenue on the Internet. The sufficient bandwidth, traffic engineering, and quality or service are just enablers. Once they are in place, there will be a proliferation of content and service providers.

They will range from private one-person occasionally-on enterprises to corporations with always-on services. They will be interactive or not, targeting mass audience or a single customer, free of charge or rate-based, etc. The entry barrier for becoming a CSP will be very low. There will be an explosive growth of the number of providers – the second coming of Dot-Coms.

Some of the services are already available and free. Providers of those services will probably meet certain challenges in trying to charge for them. Most likely the revenues for such services will be generated by advertisements. On the other hand, users are used to pay for movies or concerts. Therefore, there probably will not be any difficulties to migrating content like that to the Web. There, of course, will be some issues to resolve. For example, in an interactive talk show some viewers will opt for no-commercials version at higher cost. The other group may choose free or cheaper version with commercial breaks. The version with breaks will be delayed by the advertisement time and thus users will not be able to interact with the host in real-time. One solution is to have a split-screen advertisement without the show interruption.

The profit that CSPs generate will directly depend on usefulness of the content or services they provide. To lower their expenses, many providers will opt for selling the content or services to Internet Service Providers (ISPs) at wholesale prices thus eliminating the problems of creating and maintaining user accounts, performing metering and billing. All that will be performed by ISPs.

3.2 Internet Service Providers

Correspondingly, ISPs will play more important role than just managing delivery of data. Rather then being just enablers for e-commerce, they will become active participants in it, the connecting link between content providers and customers. ISPs will buy content from content providers at wholesale prices and deliver it to users that signed for it. To stay competitive the ISPs will have to attract richer content while being able to provide it at a lower cost to users. In that sense ISP will take on functions of Internet portals where customers will be able to choose the content they like, configure the services they want, and pay for these services, all in one place.

In fact, the convenience of payment in one place will be the predominant factor for users to stay with an ISP rather then negotiating with content providers directly. Consider a typical family that subscribed to few hundred services. Should they set few hundred accounts with few hundred providers and receive few hundred monthly statements at the end of each month? Or should they have just one account with their ISP? In fact, they may choose to sign up with multiple ISPs as the services provided by the ISPs may differ in price or variety. Say ISP A has greater choice of educational programs while ISP B provides video-telephony at lower prices. Even now users make this kind of decisions when signing up for different carriers for local and long-distance telephone calls. As ISP will compete for customers they will constantly expand their offering and thus the number of accounts users need to establish with various ISPs to get the desired variety of services is not expected to blow out of proportions. I believe one to three accounts will be typical.

Besides metering and billing functions, ISP will have to manage the delivery of the information. Different types of traffic require different performance parameters from the networks, such as maximum latency, jitter (delay variation), buffering capacities, bandwidth and throughput guarantees. Thus, ISP will have to classify the traffic in their edge routers and make routing decisions based on the type of data. ISPs routers will be virtually interconnected in the mesh networks using circuits, virtual circuits, or lightpaths provided by network operators.

3.3 Network Operators

Networks will be owned and maintained by network operators (NO). NOs own the equipment such as multiplexors, crossconnects, or switches that allow them to establish any virtual topology over the existing physical topology of their network. The virtual links of the specified capacity and with corresponding protection or restoration schemes is what network operators will sell to ISPs. Network operators never need to look at types of traffic to deliver it, nor they need to perform traffic-engineering functions – ISP's IP routers and label-switching routers (LSR) at the end of each virtual link will do all that. It is possible that one company will combine the NO and ISP functions, however, that will not be a common case as the network operators will be interested in working with multiple ISPs in order to increase the network usage. Likewise, the virtual topology used by an ISP may actually span multiple physical networks owned by different network operators, i.e., virtual links may be provided by different operators.

3.4 Customers

The true era of e-commerce will arrive when most households switch to Internet services as means of communication, receiving information and entertainment. It is crucial that Internet must stop being the domain of specialists and enthusiasts and become a background technology that is transparent to the customers. The use of new services as video on demand, video-telephony or anything yet to appear should not be less intuitive then, say, making a phone call or operating a set-top box.

Customer's choice will ultimately determine the popularity of different CSPs. When choosing a reporter for a news coverage or an e-doctor for a medical advise, customers will base their decision on ratings collected by that provider. As we will discuss below, there should be mechanisms in place to allow and encourage customers to generate the feedback. That feedback will take a form of ratings and will be available to future customers through agencies like e-Billboard or Chamber of e-Commerce. The content and service providers themselves are most interested in more customers generating the feedback. It is very likely that CSPs will provide incentives to customers for generating the feedback. However, the feedback should be unbiased and thus, the incentives should not depend on the feedback itself.

Information originates with the content or service providers. Internet service providers deliver it to customers based on their subscriptions or requests (Figure 3.a). The cash flow will generally propagate in the opposite direction from the flow of information. Payments are initiated by the customers. Some services may require a per-session payment; some will have a flat monthly fee. In any case, the transaction is done between customer and ISP (Figure 3.b). ISP will distribute (immediately or eventually) the portions of this payment to the service or content provider and network operator(s).

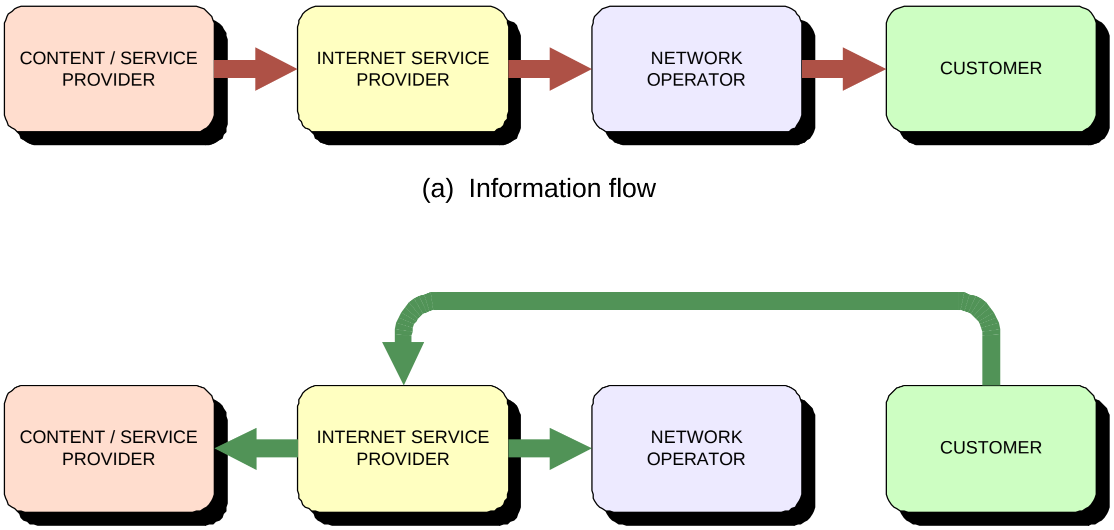
CONTENT / SERVICE PROVIDER
INTERNET SERVICE PROVIDER
NETWORK OPERATOR
CUSTOMER
CONTENT / SERVICE PROVIDER
INTERNET SERVICE PROVIDER
NETWORK OPERATOR
CUSTOMER

(a) Information flow

(b) Cash flow

Figure 3. Flows of information and cash on the Internet

4 New Role of Internet: To Serve and Protect

In previous section we described four classes of institutions that play role in making e-commerce a ubiquitous feature of our live. In this section we take a look at complex set of interactions between those players. Specifically, we will look at what cryptography mechanisms may be employed to allow the entities to perform e-commerce transactions in the absence of mutual trust, how provider's authenticity and customer anonymity can be guaranteed and how the content can be protected.

4.1 CSP-ISP interactions

Consider two entities: Alice – a content provider and Bob – an Internet Service provider. Alice wrote and recorded a song, which she decided to sell to Bob at a cost of 1¢ per one play of the song. Bob, in turn will distribute the content among its servers and allow its customers listen to it for 2 cents. The charge for that song will appear on their monthly bill.

Of course, Bob doesn't know how popular this song will be and what revenue he will generate from it. Thus, Bob will not pay Alice up-front. Instead, he will negotiate the contract to pay Alice at the end of month based on number of plays of that song, or after each play of the song. Since Alice doesn't trust Bob, how can she be sure that Bob will pay the proper amount rather then hiding the true number of requests for the song?

On the other hand, how Bob can guarantee that Alice assumes her true identity, and not just pirated a popular name, or even more, how Bob can guarantee that Alice is not reselling someone-else song?

4.2 ISP-Customer interactions

The interaction between ISP and customers are long-term in nature: customer establishes an account with the ISP, chooses and configures services to receive, and selects payment options. ISP trusts its customers because in majority of cases, the ISP performs user authentication before connection is established. Customers trust ISP because their Internet connections are configured to go to the gateways and edge routers managed by that ISP.

Having an ISP between content providers and customers will allow the customers to remain anonymous to content providers. For users it is difficult to establish trust relationships with multitude of content providers, but trust relationship with ISP is given.

However, some information ISP would like to hide from customers. For example, ISP Bob would not like to disclose to customers its contract with CSP Alice. That may prompt some customers to attempt to negotiate with Alice directly [3].

4.3 CSP-Customer interactions

The relationship between CSP and Customers are not well defined. Some content may have a very short lifespan (traffic report, for example) and thus the relationship with given provider inherently is short-lived. Considering the multitude of CSPs and the dynamic changes in CSP population (known CSPs disappear and new CSPs come up all the time), it is clear that the relationship between customers and content providers are not trustful. Customers would like to hide their true identity, as they don't know how this information may be used by CSP. On the other hand CSP would like to know how many customers viewed its content and it also would like to accumulate rating points.

5 Three-way business transactions

Below, we propose a scheme that allows a business transaction to occur while preserving the specific interests of all the participants. Those interests for CSP include the ability to verify the true number of customers that requested her content or services, and thus, the ability to get full profit. Specifically, even though the CSP doesn't know how exactly the content being distributed (i.e., the concept of super-distribution [4] is employed), it should know the number of users of the content. For customers, their specific interests include the guarantee that their identity will remain unknown to content providers.

1. ISP Bob distributes the content provided by content provider Alice and encrypted by symmetric key K known only to Alice. When customer Carol after viewing a free short sample decides to order the content she forms a request in a following way:
    a. Generate random key R
    b. Encrypt R using Alice's public key PUB_A
    c. Digitally sign the message
2. ISP Bob verifies Carol's signature and if it's correct, modifies the message by removing Carol's signature and appending his own signature and *randomly* generated message ID number. Bob then sends this message to Alice. Internally Bob remembers the ID and that the request came from Carol. When he receives response from Alice with the same ID, he will know to forward it back to Carol. In that sense, Bob acts similarly to Anonymous Internet Proxy (AIP) described in [5].

3. Alice receives the message from Bob and verifies his signature. Since Bob is one of ISPs with whom Alice has a content distribution contract, she accepts the request. She performs the following routine:
    a. Decrypt random key R using its private key.
    b. Encrypt content key K using key R received from Carol.
    c. Send R-encrypted key K back to Bob with the same message ID that it received from Bob.

4. Bob performs table lookup with the received ID and determines that the message should be forwarded to Carol. He then forwards the message, stripping it of ID number. When Carol receives this message she decrypts it using key R which it generated earlier. At this moment Carol obtains key K and can start viewing the content.

5. When Carol is done viewing the content she will rate the content provider Alice using some known rating system. Carol's incentive to perform the rating is that Alice offered a small discount to customers that do that. To rate Alice, Carol fills the form *rating* (see Figure 3), attaches it to the same random key R she generated in step 1, encrypts everything with Alice's public key and appends its signature.

6. ISP Bob verifies Carol's signature and if it is correct, modifies the message by removing Carol's signature and appending his own signature and *randomly* generated message ID number. Bob then sends this message to Alice.

7. Upon receiving message from Bob, Alice verifies Bob signature, decrypts the rating, and then she must reply with a signed message granting a discount to customer (but without knowing who the customer is). Alice encrypts this message with key R and sends it back to Bob with the same ID number.

8. Similarly to step 4, Bob performs table lookup and forwards message to Carol after removing the message ID.

9. At last, Carol forms a payment message where she writes the amount she owes Bob reduced by the amount of credit received from Alice. Carol also sends the credit amount signed by Alice to Bob for accounting purposes.

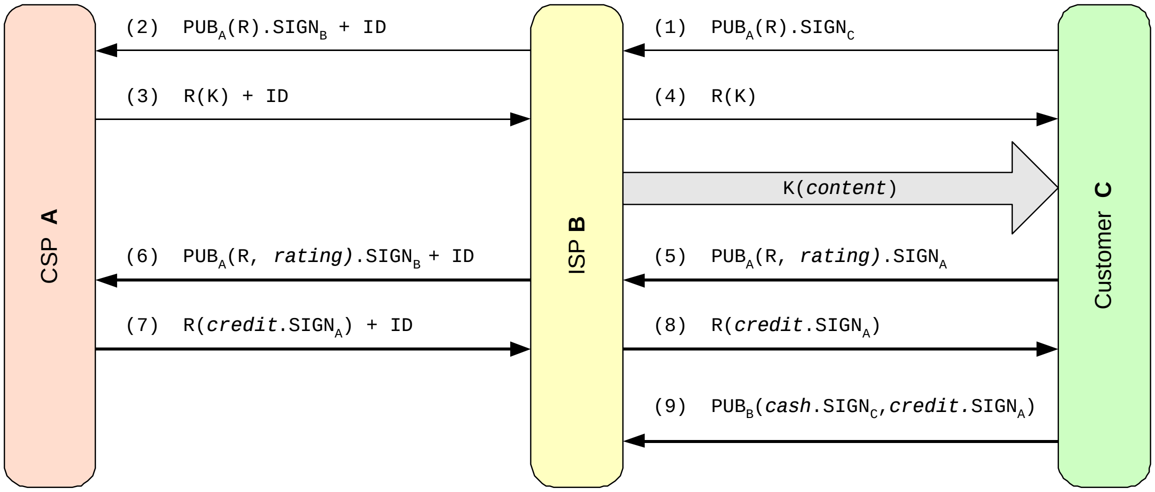
CSP A
ISP B
Customer C
(1) PUB_A(R).SIGN_C
(2) PUB_A(R).SIGN_B + ID
(3) R(K) + ID
(4) R(K)
K(content)
(5) PUB_A(R, rating).SIGN_A
(6) PUB_A(R, rating).SIGN_B + ID
(7) R(credit.SIGN_A) + ID
(8) R(credit.SIGN_A)
(9) PUB_B(cash.SIGN_C,credit.SIGN_A)

Figure 4. Protocol for three-way business transaction

As it is clear from the above protocol description, Alice never learns identity of the customer Carol. All Alice learns is the randomly generated key R and rating value, which reveals no information about Carol. What is important, Alice cannot keep activity log on the per-customer basis, because even if Carol requests the same content again, there will be different key R and Alice will not know that it is the same customer. On the other hand Alice knows exactly how many customers of Bob were viewing the content. By storing customer ratings signed by Bob, Alice can prove in court the amount Bob owes her, shall Bob underpay her. Thus, the problem of Alice not trusting Bob is solved.

Bob and Carol, on the other hand, have a trusted relationship in the fact that they know each other identities. This, however, creates problem for Carol because Bob may learn the content she requests and may misuse this information. To eliminate this problem, in the suggested protocol Bob never learns the content, unless he is simultaneously an ISP and a customer. Alice provides Bob the encrypted content. Of course, being an ISP Bob should know the type of the content, e.g., video, audio, data file, etc. Bob needs this information to be able to properly configure his network to provide the necessary quality of service for the given type of traffic. But Bob doesn't know the content itself.

It is interesting to see why Bob cannot succeed in trying to resell Alice's content if he obtains the key K. Indeed, Bob may pretend to be the customer of himself. He will purchase the key from Alice, but then he will distribute the content to other customers. This clearly will not work as the customers expect the key encoded with the key R, but Bob never learns R. Thus he cannot distribute key K to honest customers.

Dishonest Bob may then try to obtain more credits from Alice without actually distributing the content to more users. Assume that only 90% of viewers decided to obtain the discount and filled the rating form. Thus, if Bob can obtain the discounts for those remaining 10% he may pay Alice less than he actually owes her by the amount of those discounts. Then Bob will forge the message shown in step 6. Instead of relaying Carol's message, he will create a message of his own. The problem here is that Bob should provide key R that is the same as was requested in step 1 by some of the customers that didn't fill the survey. But Bob doesn't know it. If Bob creates his own random key R and validates it by performing step 1 (i.e., he will request a key K from Alice) then Alice will consider him to be just another customer and that means he will have to pay for the content as a regular user. Thus, Bob may end up paying more and surely he will not attempt to forge any messages.

And finally, what prevents Alice from being dishonest and never send the discount back after receiving the rating from Carol? In this case Carol will reveal its random key R to Bob. Bob will request Alice to resent the signed discount. At this moment Alice has two choices: either resend the discount or report that key R was never logged in her database, i.e., the key K was never issued for the presented key R. But in the latter case, Bob will reduce the amount he owes Alice by the cost of one content preview. Thus, Carol's payment to Bob will consist of only Bob's portion ant that will automatically satisfy Carol's discount (and even more). Alice will end up receiving less cash in this case.

Conclusion

In this study we attempted to illustrate how Internet may become a place for business. Specifically, we believe that most realistic scenario would be three-way transactions where content providers interact with Internet service providers and Internet service providers interact with customers. The major reasons for our argument is that ISPs are the necessary link between CSPs and customers, at least due to the fact that content delivery requires network provisioning which only ISPs can perform. Then, we argued that since ISP cannot be excluded from the distribution chain, it would undertake additional tasks to maximize its role and increase its profits from participating in the distribution chain. And indeed, there are specific need of customers and content providers that ISP is uniquely positioned to perform: privacy, billing and metering.

In Section 5 we illustrated a three-way transaction and described the protocol. While this protocol may not be very efficient (uses public-key encryption) it is very simple and does not require expensive special hardware or complicated deployment strategies. To minimize the computational overhead, the content itself is encrypted using the symmetric-key encryption. Public-key encryption is used only to obtain the symmetric key and to provide the feedback. The open question we left is how to prevent distribution of key K by a customer, or at least how to do traitor tracing. General schemes of traitor tracing [6] are not easily compatible with a concept of super-distribution. That is a topic of future research.

References

[1] M. Silveron, *"The Value of Bandwidth"*, Fiberhood Networks, http://www.fiberhood.net/riis/presentations/value_of_bandwidth.pdf.

[2] S. Hardy, *"Verizon staffers find fiber-to-the-home cheaper than copper,"* Lightwave, PennWell, vol. 17, no. 134 pp. 1, December 2000.

[3] G. Durfee and M. Franklin, "Distribution Chain Security", *Proc. ACM Conference on Computer and Communication Security*, Athens, 2000.

[4] Kaplan M.A., "*IBM Cryptolopes TM , SuperDistribution and Digital Rights Management*", IBM Corporation, December 1996, http://www.research.ibm.com/people/k/kaplan

[5] I. Goldberg, "Using the Internet Pseudonymously", RSA 2000, www.isaac.cs.berkeley.edu/rsa200_slides.ps

[6] B. Chor, A. Fiat, M. Naor and B. Pinkas. "*Tracing traitors***"**, IEEE Transactions on Information Theory 46, pp. 893-910, 2000.