

Towards Flexible Metering & Charging For Information Services

Cliff Redmond and Vincent Wade
Department of Computer Science,
Trinity College,
Dublin 2,
Ireland.

e-mail: {Cliff.Redmond,Vincent.Wade}@cs.tcd.ie

1. Introduction

Electronic publishing is evolving to overlap with the world of on-line multimedia services. While there are many issues still to resolve, including copyright protection and content presentation, it is generally accepted that the Internet, and specifically the World Wide Web, provides a glimpse at a possible future electronic publishing industry.

The Internet itself has witnessed incredibly rapid growth in the last five years, thanks to the popularity and ubiquity of user-friendly Web clients like Netscape's Navigator and Microsoft's Internet Explorer. Many organisations now consider a presence on the World Wide Web to be vital to their commercial interests. However, due to the absence of a suitable remuneration mechanism, most Internet publishing is funded indirectly by advertisements that adjoin popular articles and sites, or through public subsidy of universities and research institutions. With the exception of some fixed-fee subscription based services with dynamic content, e.g. newspapers, very little content, or information, pays for itself at the moment.

Part of the reluctance of conventional publishers to embrace the World Wide Web is also due to the lack of a copy control mechanism; currently the Web encourages information and content replication, with caches and mirror sites becoming part of the everyday user's common lexicon. Together with the Web's architectural difficulties, the reluctance to publish using this new medium must be in part due to the Internet ethos - a throwback to the early days of the Internet, which has as a basic tenet that information should be free. This ethos still pervades the Web and its influence is present in the assumption, by many, that content made available on the Web is by its very existence public-domain and copyable [Inet97]. This assumption is left unchallenged because of the lack of viable content protection mechanisms; although some document protection devices are available they have failed to pervade the common clients to provide the ubiquity needed to support information commerce [Digibox].

This paper outlines a new paradigm for content provision by information service providers - one that backs away from physical marketplace models of commerce. Instead, it takes advantage of one of the unique qualities of on-line services: the information generated as a by-product of information usage. Also examined in this paper are the issues which are relevant to devising flexible usage and charging schemes to support information publishing on the Internet.

The information service providers mentioned will be among the publishers of the next generation. This is currently highlighted by traditional publishers battling to embrace the Web, even if they are not too sure how best to use it; most at present only support, or give information about, titles they sell through other media, e.g. paper and CD-ROMs.

This paper first outlines some general approaches to charging for information services and then examines the properties of digital information which can effect metering and charging for such information services. The paper details the service architecture, design and implementation of a generic accounting management system suitable for many services, but applied to a Web based information service.

2. Charging for multimedia services

This section discusses some of the basic charging schemes employed by service providers in the Internet and telecommunications domains.

Service providers are inclined to adopt charging schemes based on two extremes of what could be considered a continuous spectrum of possibilities [Sarkar]. Providers tend to adopt either a flat-fee (usage insensitive) or usage based (usage sensitive) charging scheme, with variations of the latter including responsive charging, content-aware (information sensitive) charging and content-blind (information insensitive) charging.

The flat-fee charging scheme is commonly used to pay for Internet connectivity by both individual and institutional users. For users it yields a predictable charge with little difficulty in understanding the bill, whereas for service providers it provides a predictable income with little accounting overhead. It must also be noted that flat-fee charging does not penalise information sharing; this sharing is referred to as “public-spirited behavior” [RFC1272], and provides the background for Internet usage and accounting. It also raises some difficulties for Internet based information service providers and content providers. Unfortunately this scheme typically leads to heavy network usage because there is no penalty (other than the resultant network congestion) for users consuming high levels of bandwidth.

To deal with congestion, usage-sensitive pricing schemes have been suggested, with common examples being statistical measures of effective bandwidth and charges for volumes of packets or cells [CANCAN]. MacKie-Mason and Varian [RPI] suggested a responsive usage-sensitive pricing mechanism that could actively deter congestion on networks. They proposed a ‘smart market’ where users indicate a maximum charge they would be willing to pay for bandwidth in a tight loop feedback mechanism. However, this scheme has had its detractors. For example Shenker *et al.* in [PCN] argue that the costs of implementing and policing such a policy might not actually exceed the monies recovered. It is also questionable whether such a mechanism is actually implementable because the costs of congestion are very difficult to determine accurately. Instead, Shenker *et al.* suggest that “It is important to allow prices to be based on some approximation of congestion costs, but it is important to not force them to be equal to those congestion costs”. Notably, they state that while support for pricing might be included, no pricing policy should be embedded in the network architecture. This would allow different connectivity providers a choice in how they price a service.

2.1 Multimedia and connectivity service providers

While network connectivity and access providers are currently only able to charge for Internet connectivity on a flat-fee basis, it is clear that any usage-based scheme could potentially yield higher revenues. Usage based charging also has the positive side-effects of enabling network policing and facilitating quality of service guarantees (which are often mooted as beneficial, independent of the mechanisms needed to realise them). As a basic transport architecture

adapts and adopts mechanisms to enable usage-based accounting schemes for multimedia services, the connectivity providers will in turn use these mechanisms.

Information service providers ought to be wary of too low-level a measure of usage (e.g. bytes transferred) as this could be to the detriment of providers of high value, high bandwidth content, e.g. still and video image content providers. From the content provider's point of view, some bytes are worth more than others!

Too heavy a dependence on low-level usage based accounting by connectivity providers could constrain the development of content providers, and their supporting information service providers - the electronic publishers of tomorrow. For this reason, promotion of competition at the connectivity provider level is vital. The provision of content-, rather than application-, aware architectures for communication could also benefit service providers at layers above the connectivity service.

At the other extreme, keeping the current status quo could detract from the overall service provided by an information service provider; flat-fee schemes do not discourage high bandwidth usage, which can in turn produce network congestion and frustrated users. It is clear that some bandwidth policing mechanisms are also required by information service providers.

It is apparent that a compromise solution is necessary. For example a combination of flat-fee and usage-based pricing that exists within the continuous pricing spectrum suggested by [Sarkar] may be more appropriate. Such a charging scheme was examined in a trial undertaken by the Prospect project (a European research projects funded under the ACTS program). This project involved the provision of an educational service by a value-added service provider (a combination of a content provider and a multimedia service provider) using a separate underlying connectivity service provider. This scheme enabled value-added service providers to pay for bandwidth requirements in an otherwise content-blind network, while users paid the value-added providers directly for the services delivered, in this case for (educational) information services.

2.2 Hybrid usage-based charging for information services

When attempting to build any charging mechanism, it is necessary to have a means to meter the use of a service's resources, store any data generated by this metering activity, and generate the charges based on appropriate tariffs. Once generated, these charges can be aggregated and used as a basis of a bill for the service user.

The Telecommunications Information Networking Architecture Consortium, TINA-C, has proposed a software-based service architecture to support the rapid development and provisioning of a wide range of telecommunications services. It defined a common approach for the design of services and service management which allowed the reuse of software components. It also defined a number of common components for service accounting, subscription and session control, and proposed the use of a distributed computing environment within which to implement these components.

The distributed processing environment defined by TINA closely resembles, and would provide similar functionality to, the Open Management Group's [OMG] Common Object Request Broker Architecture, CORBA. This resemblance led the Prospect project to begin to develop its systems using this common architecture. CORBA supports an abstract interface definition language (IDL) which allows components written on different platforms, in

different programming languages, to inter-operate. This facilitates service component and software integration and reuse.

Prospect developed subscription and accounting management services based on the TINA consortium specifications on top of a CORBA component platform. Central to the TINA-C Architecture is the concept of a session. TINA sessions allow identification and authorisation of users (access sessions), service usage, control and management (service sessions) and communication resource management (communication sessions). Fundamentally a session provides a context for relating a user's interactions with a service; the implementation of this context enables service management and facilitates a hybrid charging scheme that is based on a combination of flat fees and service usage charges. The project concentrated on the development of access and service session components because the Internet Protocol (IP) based services used in the project do not require active communication resource management.

The implementation of common accounting components is described in section 5.2. The use of these components necessitated the mapping of the service session concept onto the inherently stateless protocol interactions of IP services. The service implementation section outlines how this was achieved for a WWW based information service, and describes the interfaces to the common accounting components. The following sections outline special considerations for the content that information services support, and describe a charging mechanism and policy suitable for such services.

3. Properties of information relevant to charging

There are several important considerations to take into account when charging for information in digital form. Some of the more important concern the nature of the information in this form. For example, this digital information is:

- **medium-free;** this means that information exists essentially independently of its medium; its expression to the ultimate user can be transformed. For example, e-mail discussion lists are regularly archived and made into an information resource, with threads, or links between e-mails used to enable the user to follow the discussion; still images can be taken from movies, and altered; recorded sound can be edited to change the order of speakers, their voices or background noises. The multimedia publishing industry is most immediately affected by this problem, and is faced with trying to license information for uses for which it was never intended, often finding great difficulty even tracing the original producer. This has even led for calls to set up a single 'clearing house' for multimedia publishers [ACM196], [Bangemann].
- **easily copied;** digital information is "...something that can be stolen without depriving the owner of it..." [ElecPir] because instant and perfect copies are possible. Existing economic theories, based on pay-per-copy, are being stretched by this property [Cox]. Copyright - the traditional link between ideas, their expression and exploitation - is also suffering, with many newly framed laws that cover copies of digital information (even copies transient in nature) being criticised for upsetting the traditional balance between the owner of copyright and the public good that such law was originally designed to protect [ACM1294].

A solution to these problems might be to discourage users from copying information, other than in a transitory way. Thus multimedia producers could refer to the original form when transforming information, and copies of documents could refer back to their source when being read. There have been schemes that attempted to enforce this, including electronic document management systems [DigiBox] and the infamous Xanadu® system [Xanadu] with its attendant notions of a new permissions system called transcopyright and transclusion ('quoting without copying') [CoXanadu].

Other important considerations that are relevant to charging for any form of information include:

- the **cost** of producing the information, including rewards for the primary producers (e.g. authors, film production company) and the publishers.
- the **nature** of the information; whether it is relatively dynamic (perhaps news normally found on television or in newspapers, or stock market quotes), or relatively static (perhaps reference material, e.g. a family encyclopaedia, or novels).
- the **user** using the information; while information is traditionally sold on a mass marketing basis, an interactive information system offers information producers much more feedback from their customers, on a per person basis, and will possibly lead to personalised information sources [Persapress],[PinaWeb] and much more niche marketing and production of information: "Think niche. It's the net's greatest strength" [O'Reilly].

The publishing industry has a number of roles in the "information society" [Toffler84], [Bangemann], including those of editing and selecting information, which are highlighted by another two properties of information. These properties are linked:

- information is **generative**; "If you use a piece of information, I can use it too...if we both use it, the chances are improved that we will produce more information [hence] it is generative" [Toffler84]. This property is evidenced by the explosive growth of sites on the World Wide Web in the past five years [Netcraft] and their inter-referencing.
- **clutter**; too much information is, or can be, bad; "Up to a certain point more choice is better. Then... the user gets overwhelmed, and less is more." [O'Reilly].

Thus publishing could continue overseeing both content and context, where providing context (or meta-information) could be as simple as providing a selection of links to favourite information sources, or as complicated as a sophisticated catalogue, e.g.[Yahoo]. Many on-line services already offer a 'best of the web' service to their subscribers, for example Prodigy, America-Online and Compuserve. Such meta-information points can be more valuable because of their selection of other information resources, while still referencing the original source of information which allows more in-depth user exploration. This could give the original information producers an opportunity to gain in any usage-based charging scheme.

Web-based information services are most commonly financed by one of several indirect means: by subsidy, in the case of universities and research institutions; by advertising, in the case of search engines and other meta-information points; and by subscription, or a combination of subscription and advertising, in the case of on-line editions of newspapers or magazines (*The Economist*, *The New York Times*). As yet, very few (if any) sites charge for information accessed or used.

Two usage policies could be supported by the existing Web infrastructure. These policies depend on information in the protocol headers detailing the age of a document (the 'Last-modified' date of the HTTP header). They are:

- **Paying per access**; this is equivalent to paying for a new copy of a document on each access, while not necessarily fetching it. Even though the user might use a local cached copy instead of a copy from the original server, the original server is still informed of its use by a conditional request for the document (an "If-Not-Modified" header).
- **Paying per copy**; this ensures a new copy of the document is down-loaded each time it's requested, which is a potential waste of bandwidth. This can be supported by giving a document a "pre-expired" last-modified date, or none at all.

However, both of these charging schemes would tend to encourage copying of information; the user is effectively penalised for using the same information more than once. A policy outlined in the following sections might be properly called pay-once-per-use, and could act as an encouragement not to copy.

4. An architecture for multimedia service accounting systems

The Prospect accounting model and components are based on the concepts and models outlined in the TINA Consortium's Service Architecture [TINA95]. Accounting management components and information objects were defined in a three-stage model of usage metering, charging and billing, with tariffs providing a mapping between usage data and charging, and charging and billing, at different stages of the accounting process. The accounting components were designed to be used with many multimedia services, with service-specific components enabling different services to reuse the same components.

4.1 An information-service-usage metering system

The first stage of the three-stage model, the usage-metering mechanism, was designed to operate on a per-session basis. Sessions were modelled using protocol state information in HTTP requests to a modified web server [Apache]. The Web server was based on Apache release 1.1.1 code - ported to C++ - with additions to the authentication module and a complete re-write of the user-tracking (or cookie) module. This included a small, lightweight, in-built CORBA client to enable remote usage monitoring and statistics gathering. These additions resulted in new directives that could alter the server's overall configuration and that could change its behaviour on a per-directory basis.

In the TINA architecture a user is required to explicitly start a service session. This is achieved using Java/CORBA desktop components in Prospect, where a URL with a reference to the service specific MIME-type¹ is returned as the initial service reference. The service user then has a fixed session starting point contained in the first URL that they request from the server. This is used to identify the user to the service; the protocol state information is generated from the contents of this URL in a service-specific MIME-type handler. The state information in a request acts as a context in which a request is made, and therefore as an

¹ A MIME-type describes the content of a requested file and can invoke a specific handler in the Web server

application layer session identifier. The user cannot retrieve any information from the server without a valid session identifier.

4.2 Charging and billing

The usage data generated for each user by the modified Web server is stored. This usage data is then used to generate charges based on a tariff with a charging schedule. These charges are in turn used to generate a bill for service usage, based on another tariff. This two stage charging and billing mechanism allows additional fixed charges and the discounting of basic charging information.

At the moment the majority of service charges are based on fixed tariffs which include costs per session and a flat per use fee for each piece of information used within that session. This scheme has some in-built policy decisions, and does not discourage the copying of information, as it is still a basic pay-per-access scheme. The following section discusses a more innovative charging approach for information services which could support electronic publishing.

5. A proposed dynamic charging policy based on information usage

Enforcing service access through a session mechanism allows the comparison of service usage in different sessions. A new charging policy could try to account for users requesting some information more than once and not penalising them for this, since this could show they are still using the original source of the information, not a local copy. A single cumulative charge for the use of a piece of information over all sessions could be used, so that a user might actually be charged less for a piece of information if their usage profile fell within defined parameters, e.g. users might only pay once for a piece of information, and are rewarded by smaller reimbursements if their profiles show they have not copied information. In order to compare a user's session, two basic indexes were devised; an inter-session repeat index, or ISRI, which gives the number of repeat requests for information in two consecutive sessions (where consecutive sessions were thought adequate on a brief examination of data) and an intra-session repeat index, or IaSRI, which monitors the number of repeat requests for a piece of information within the same session.

In order to have an adaptive pricing mechanism it is necessary to have measures to estimate a typical user's repeat-request average (ISRI) and have a policing mechanism to stop users trying to take advantage of the system (IaSRI). A policy-based scheme which tries to discourage copying by referring to the original source of the information can only be effective if users are educated as to the adaptive nature of the pricing. Unfortunately there will always be the user that will try to 'beat the system' by requesting a document more than once within the same service session in the mistaken belief that the price of the piece of information will fall to its lowest level. The IaSRI is in place to attempt to police this needless waste of bandwidth and could be used to penalise such attempts.

5.1 Applying charging policies in tele-education

The ISRI is a general index, and will probably vary widely depending on the nature of the information being served, i.e. how relatively dynamic it is. Figure 1 shows a graph of session overlap (ISRI) versus inter-session times for educational resources used in a web-based

educational course that was used by final-year university students in 1996/1997. This graph shows a marked peak in the number of sessions with two to three repeat requests for a period of 24-48 hours after the previous session. One might expect that in an educational course the number of repeat requests for particular pages (information resources) would be small in consecutive sessions, but would be larger when separated by a period in the region of one week. That is to say students would not need to refresh their memory on a daily basis, but would rather require reuse of information over a weekly period. As shown, the number of repeat requests peaks within a period of 48 hours after the previous access of the information. There also appears to be quite a low number of repeat requests in subsequent sessions (this would have shown as a pronounced diagonal ridge on the graph). This illustrates that the user usage pattern for information resources may not be intuitively predictable. Thus in order to define reasonable usage profiles/patterns, non-commercial trials must be used to generate acceptable patterns. A trial period is already common practice with many Internet based subscription services, to allow the subscribers to judge the product; by recording the usage patterns in such test periods the service provider could gain useful information too.

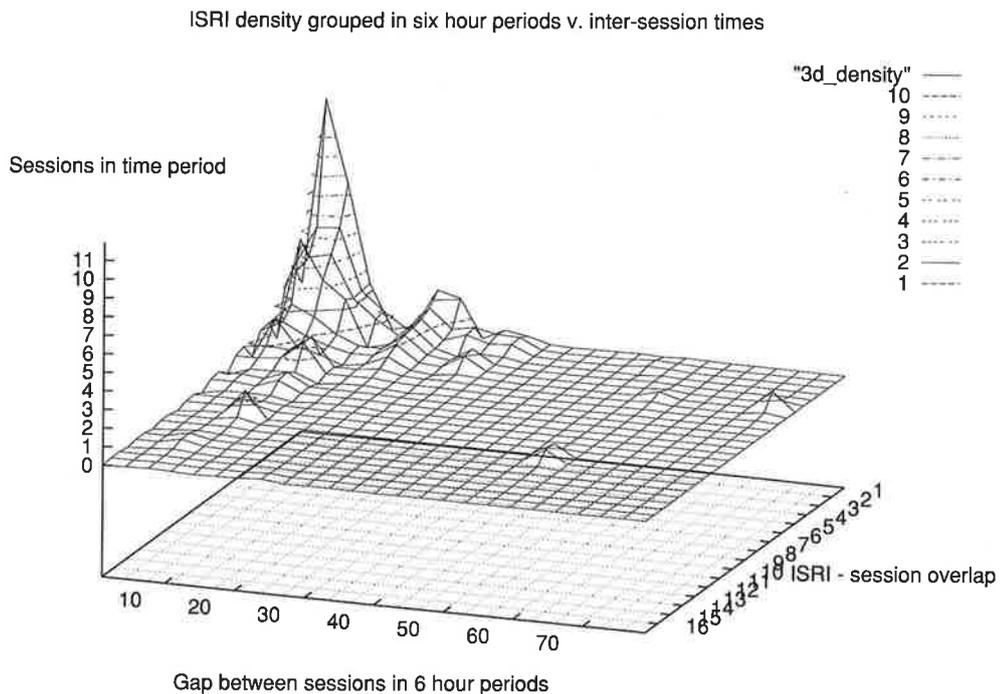


Figure 1.

The information made available was course materials, but for other information types, e.g. general reference material, quite a different ISRI profile might be expected. It must also be noted that most users used the course quite quickly, which could also explain the higher number of sessions with repeat requests within the first two days of use.

Given the ISRI for a pool of information, and any user's ISRI, a policy could be implemented that decreased the cost of information to that user if their ISRI fell within the expected range for the inter-session time. The presence of too high an IaSRI could be used to penalise a user for needlessly using a provider's bandwidth, it could be used as a basis for a charge for that user. Too high an ISRI might also be penalised, because this may be due to the usage of more

than one user. The research on these indexes is ongoing, and further trials and experiments are planned.

6. Service implementation

Sessions are modelled using HTTP protocol state information, commonly called 'cookies'. This state information was originally introduced by one browser developer, but was implemented as a *de facto* standard by many web server developers as it offered web-sites a useful means of user tracking. It has subsequently come onto the Internet standards track [RFC 2109], and is now supported by most browser implementations.

The web server, or daemon, used in the implementation was modified to allow it to be configured to refuse to serve all but one service-specific MIME-type without the presence of the correct state information in the client's HTTP request. Without this state information, normally set by the server in the headers of the first response to the client, a user using a browser could only successfully request the service-specific MIME-type from the server. This modification included the addition of new configuration directives, acting on a per virtual server and per directory basis, that influence how the server handles requests

Given a request for the service-specific MIME-type, the handler for this mime-type decodes the URL that was requested and uses the information contained in it to form and initiate the cookie using a 'Set-Cookie' HTTP response header.

If the user subsequently requests a file of any MIME-type from the server, the request will always include the state information associated with the first interaction. When a request arrives at the server, the requested file and user information, along with a code signifying the success or otherwise of the request, are forwarded to a standardised accounting subsystem via a CORBA one-way call to a multithreaded User Session Manager (USM). (The TINA based USM component is described in more detail in the next section.) The addition of these CORBA invocations does cause some overhead, with a small, but constant, delay of approximately 28 milliseconds per request (e.g. this resulted in a reduction of server throughput from 1.57 MB/sec to 1.49 MB/sec for a 10k file)

The usage metering data is forwarded from the USM to an Usage Metering Data Computational object (UMDataCO), which collects usage metering data for each session. When a session ends, the data for the user's session is stored in another accounting component, the Usage Metering Log Computational object (UMLogCO), ready to be used as a basis for charging.

6.1 Component implementation

The components used to monitor and log usage, control tariffs and generate charges and bills are based on Computational Object specifications published by the Telecommunications Information Networking Architecture (TINA) Consortium in their 1995 Service Architecture.

The UMDataCO, or usage metering data computational object, is based on the X.742 Recommendation 'Usage Metering Function for Accounting Purposes'. A subset of the metering data information package outlined in the X.742 Recommendation is used. This consists of usage data blocks to report session registration, requests and session completion

(complete block). The implementation has concentrated on single user sessions, and the request block represents a request made inside a session, rather than a request to join a session.

The UMDDataCO accepts session registration and completion blocks from service-independent parts of a service-specific USM. As mentioned above, the USM receives usage information from the web server, via CORBA one-way calls to the Service dependent interface (Sdi). It can associate the information with a TINA service session and forward it to a UMDDataCO for that user's session.

Figure 2 outlines the service-dependent and - independent parts of the accounting subsystem for the web-based service. The notation used follows that used by Jacobson in [OOSE]. A detailed description of the interactions follows.

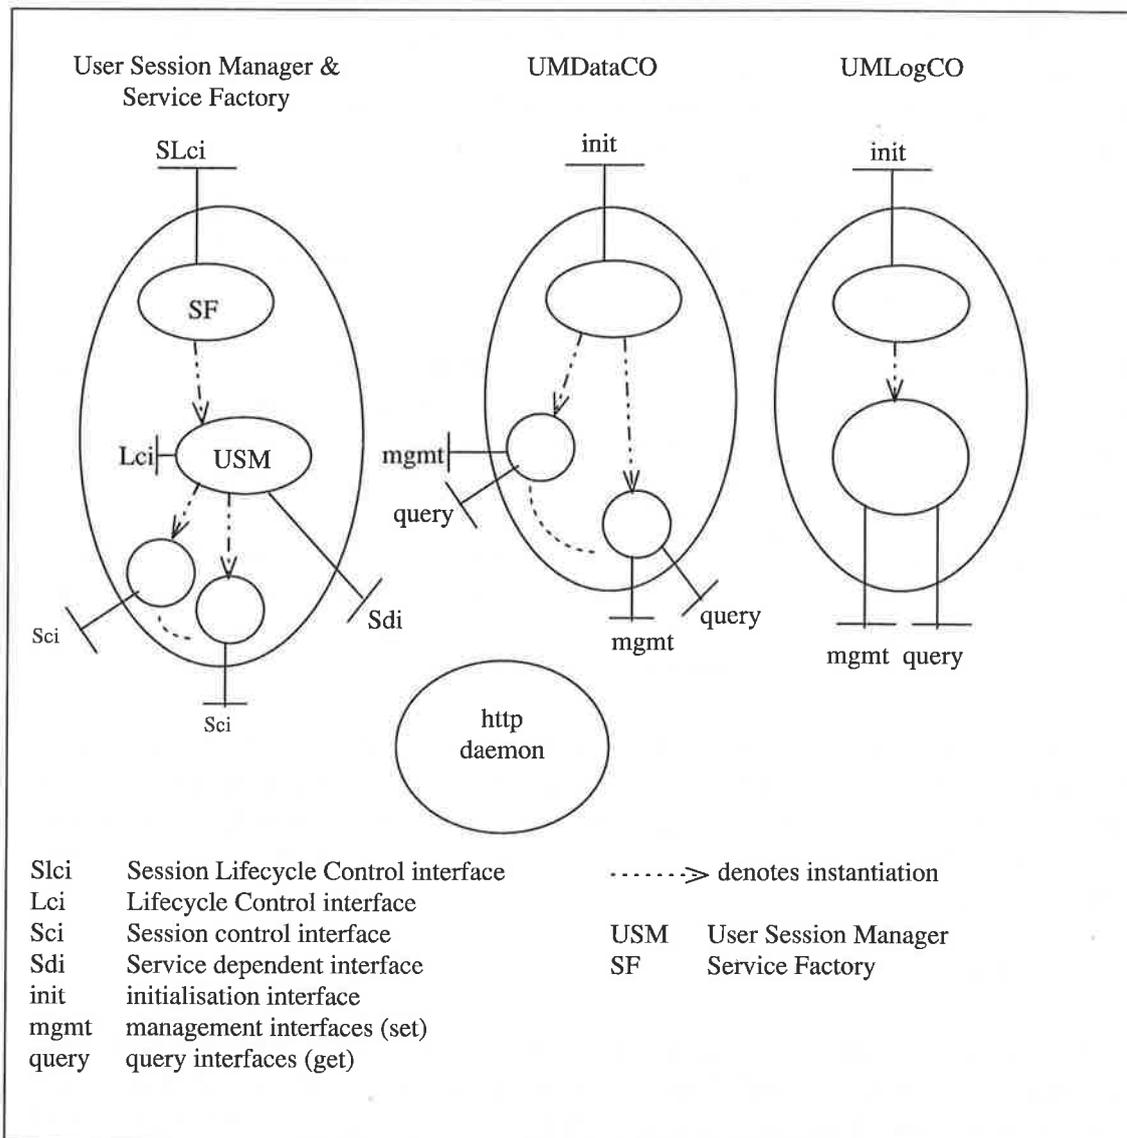


Figure 2.

There are two factory interfaces shown. The object that implements the initialisation interface of the UMDDataCO produces objects that meter a session's usage, and the object that implements the USM's Lifecycle Control interface (Lci) produces Session control interface

(Sci) objects. Each session control interface object has a corresponding object within the UMDDataCO that meters that session's usage. When a session is started the Sci object forwards a registration block to the appropriate UMDDataCO management interface (i.e. the management interface for that user), and when the session ends a complete block is forwarded.

Requests for resources are forwarded through client code embedded in the HTTP daemon and arrive at the Service dependent interface of the user session manager. The information in these calls signifies the session identifier of the requester and the information requested; the session identifier is used to find the corresponding management interface of the UMDDataCO, and the request information is forwarded there. The implementation of the Sdi object takes advantage of multi-threading to ensure collection of all usage information; a session dependent interface object per user wasn't viable because of a need to keep the client code in the HTTP daemon simple and fast.

When a session ends, a cascade of interactions occurs. The object representing a session, its Session control interface (Sci) object, is deleted by the Session Lifecycle control interface (SLci) object of the service factory (SF) via the USM's Lifecycle control interface (Lci). These interactions are internal to the SF/USM component. When the Sci object is deleted, this in turn deletes the corresponding object within the UMDDataCO via the UMDDataCO's init interface. When this happens the information associated with that session is stored using a call to the management interface of the UMLogCO component (the store operation). This information is essentially independent of any particular charging policy and can be used by many different charging schemes, some of which have been outlined earlier.

7. Future work and conclusions

Several charging paradigms were outlined, including content-blind usage-sensitive charging and content-sensitive charging for information. Also outlined were the difficulties associated with information charging and an analysis of user usage patterns for a web based course.

Current work is concentrating on incorporating some of the charging approaches discussed into the TINA accounting architecture, in which information services are currently not very well supported.

Future research will concentrate on re-requests of information from the point of view of the information being used, rather than the user using it; this would allow group usage schemes, which would be difficult to account for using the single user session focus of the analysis discussed above.

Future work will also accommodate HTTP caches into the charging scheme; users ought to be rewarded if they do not needlessly use a provider's bandwidth, so some discount might accrue when a browser makes a conditional (on the expiry date of the information) request. It might also be useful to account for badly administered or implemented caches that never verify information, by charging them a fixed price for information for all users in a group (or at a site) once one of them requests the information; here charging from the information's point of view, rather than the user's, could show its utility.

8. References

- [ACM1294] Pamela Samuelson, *The NII Intellectual Property Report*, Communications of the ACM, Vol. 37, No. 12. (December 1994).
- [ACM196] Pamela Samuelson, *Intellectual Property Rights and the Global Information Economy*, Communications of the ACM, Vol. 39, No. 1. (January 1996).
- [CANCAN] CANCAN Project (AC014), *ATM Charging Schemes*, CANCAN Deliverable 5.
- [COX] B. Cox, *Superdistribution*, Addison-Wesley, Reading MA, 1996.
- [PCN] Shenker et al., *Pricing in computer networks: reshaping the research agenda*, Telecommunications Policy, Vol. 20, No. 3, Pergamon, 1996.
- [SACP] J. MacKie-Mason et al., *Service Architecture and content provision*, Telecommunications Policy, Vol. 20, No. 3, Pergamon, 1996.
- [DigiBox] Olin Subert et al., *The DigiBox: A Self Protecting Container for Information Commerce*, Proceedings of the First Usenix Workshop of Electronic Commerce, Usenix Association, Berkeley, CA, USA, 1995.
- [ElecPir] John Chesterman and Andrew Lipman, *The electronic pirates : DIY crime of the century*, Routledge, London, 1988.
- [Inet97] Internet Magazine (UK), September 1997.
- [OOSE] Ivar Jacobson, *Object-oriented software engineering : a use case driven approach*, Addison-Wesley, Wokingham, 1992.
- [OPISN] Q. Wang et al., *Optimal Pricing for Integrated Services Networks*, Internet Economics, MIT Press, Cambridge MA & London, 1997.
- [O'Reilly] Tim O'Reilly, *Publishing Models for Internet Commerce*, Communications of the ACM, Vol. 39, No. 6. (June 1996).
- [RPI] J. MacKie-Mason et al., *Responsive Pricing in the Internet*, Internet Economics, MIT Press, Cambridge MA & London, 1997.
- [Sarkar] Mitabarun Sarkar, *Internet Pricing: a regulatory imperative*, Internet Economics, MIT Press, Cambridge MA & London, 1997.
- [TINA95] Telecommunications Information Networking Architecture Consortium, *Service Architecture*, Version 2, editors; Berndt & Minerva. (Available at [TINAC])
- [Toffler84] Alvin Toffler, *Previews & premises : an interview*, Pan Books, London, 1984.

8.1 WWW References

Apache	http://www.apache.org/docs/misc/API.html
Bangemann	http://www2.echo.lu/eudocs/en/bangemann.html
CoXanadu	http://www.wired.com/wired/3.06/features/xanadu.html
DigiBox	http://www.intertrust.com/architecture/overview.html
IPv6	http://ds1.internic.net/rfc/rfc1883.txt
NetCraft	http://www.netcraft.com/survey/index.html
OMG	http://www.omg.org/
Persapress	http://ksgwww.harvard.edu/~itbspp/proj16/home.htm
PinaWeb	http://www.cern.ch/WebOffice/Projects/Newspaper/index.html
RFC1272	http://ds1.internic.net/rfc/rfc1272.txt
TINAC	http://www.tinac.com/
YAHOO	http://www.yahoo.com/
Xanadu	http://www.sfc.keio.ac.jp/~ted/xanadu.and.osmic.html