

Cyberspace

Mark D. Pesce, Peter Kennard and Anthony S. Parisi

*Labyrinth Group
45 Henry Street #2
San Francisco, CA 94114*

mpesce@netcom.com peterk@netcom.com dagobert@netcom.com

Introduction

The emergence, in 1991, of the World Wide Web, added a new dimension of accessibility and functionality to Internet. For the first time, both users and programmers of Internet could access all of the various types of Internet services (FTP, Gopher, Telnet, etc.) through a consistent and abstract mechanism. In addition, WWW added two new services, HTTP, the Hypertext Transaction Protocol, which provides a rapid file-transfer mechanism; and the Uniform Resource Locator, or URL, which defines a universal locator mechanism for a data set resident anywhere within Internet's domain.

The first major consequence of the presence of WWW on Internet has manifested itself in an explosion in the usability of data sets within it. This is directly correlatable to the navigability of these data sets: in other words, Internet is useful (and will be used) to the degree it is capable of conforming to requests made of it. WWW has made Internet navigable, where it was not before, except in the most occult and hermetic manner. Furthermore, it added a universal organization to the data within it; through WWW, all *four million* Internet hosts can be treated as a single, unified data source, and all of the data can be treated as a single, albeit complexly structured, document.

It would appear that WWW, as a phenomenon, has induced two other processes to begin. The first is an upswing in the amount of traffic on Internet (1993 WWW traffic was 3000x greater than in 1992!); the second is a process of organization: the data available on Internet is being restructured, tailored to fit within WWW. (This is a clear example of "the medium is the message", as the presence of a new medium, WWW, forces a reconfiguration of all pre-existing media into it.) This organization is occurring at right angles to the previous form of organization; that is to say that, previously, Internet appeared as a linear source, a unidimensional stream, while now, an arbitrary linkage of documents, in at least two dimensions (generally defined as "pages"), is possible. As fitting the organization skills most common in Western Civilization, this structure is often hierarchical, with occasional exceptions. (Most rare are anti-hierarchical documents which are *not* intrinsically confusing.)

Navigability in a purely symbolic domain has limits. The amount of "depth" present in a subject before it exceeds human capacity for comprehension (and hence, navigation) is finite and relatively limited. Humans, however, are superb visualizers, holding within their craniums the most powerful visualization tool known. Human beings navigate in three dimensions; we are born to it, and, except in the case of severe organic damage, exceed the abilities of any other species in our comprehensive ability to spatio-locate and spatio-organize.

It seems reasonable to propose that WWW should be extended, bringing its conceptual model from two dimensions, out, at a right angle, into three. To do this, two things are required; extensions to HTML to describe both geometry and space; and a unified representation of "space" across Internet. This work proposes solutions to both of these issues, and describes a WWW client built upon them, called "Labyrinth", which visualizes WWW as a space.

Visualization and HTML

As of this writing, HTML is capable of expressing both textual and pictorial data, and can provide some limited formatting features for each of them; beyond this it provides a linkage mechanism to express the connection between data sets. HTML's roots are in text; its parent, SGML, specifies a format for printed media, a expression which is intrinsically two-dimensional. For this reason, we

have stepped “outside” of HTML in our language specifications for geometry and place, defining a simple, easily parsed scripting language for the generation of objects and spaces.

The basic functionality for any three-dimensional language interface to WWW can be broken into three parts; object definitions, which include the definitions of the geometric representations for these objects; scene definitions, which define “placement” of these objects inside of a larger context; and a mechanism which “binds” a URL to an object within a scene. The Labyrinth Scripting Language (Appendix A), while unsophisticated, does fulfill all of these requirements, and therefore provides all of the basic functionality required in a fully visualized WWW client.

As currently defined, Labyrinth “world” files are a unique data type, like MPEG or AIFF, and must be integrated with MIME in order to launch a companion “viewer”. This is not an optimal solution; rather, it should be possible to extend HTML to encapsulate “spatial” data types; these, then, could be visualized or ignored given the capabilities of the WWW client. The OpenGL or HOOPS specifications could form a basis, insofar as object definitions are concerned, for HTML extensions, and should be examined as a possible (and well-supported) solution to this issue. Our scripting language should serve as a starting example, rather than a proposal for an all-inclusive solution.

Any conceptualization of space contains within it, implicitly, the quality of number; i.e., “how much” or “how far” is contained within the simple expression of existence. Space, in its electronic representation, is numbered, and, if it is to be shared by billions of simultaneous participants, it must be *consistent, unique, and very large/dense*. Despite this, it is rarely necessary for a WWW client to deal with the totality of space; operations occur local to the position of the WWW viewer, and this local description of space is nearly always a great deal more constrained than the entire spatial representation.

It is necessary for HTML to define a numbering system for visualization which conforms to the three principles outlined above. Another section of this work describes such a system.

Cyberspace

For the purposes of continuity in navigation, it is necessary to create a unified conceptualization of space spanning the entire Internet, a spatial equivalent of WWW. This has been called “Cyberspace”, in the sense that it has at least three dimensions, but exists only as a “consensual hallucination” on the part of the hosts and users which participate within it. There is only *one* cyberspace, just as there is only one WWW; to imply multiplicity is to defeat the objective of unity.

At its fundamental level, *cyberspace is a map that is maintained between a regular spatial topology and an irregular network topology*. The continuity of cyberspace implies nothing about the internetwork upon which it exists. Cyberspace is complete abstraction, divorced at every point from concrete representation.

All of the examples used in the following explanation of the algorithmic nature of cyberspace are derived from our implementation of a system that conforms to this basic principle, a system developed for TCP/IP and Internet.

Metrics in Cyberspace

Internet defines an address “space” for its hosts, specifying these addresses as 32-bit numbers, expressed in dotted octet notation, where the general form is $\{s.t.u.v\}$. Into this unidimensional address space, cyberspace places a map of N dimensions ($N = 3$ in the canonical, “Gibsonian” cyberspace under discussion here), so that any “place” can be uniquely identified by the tuple $\{x.y.z\}$.

In order to ensure sufficient volume and density within cyberspace, it is necessary to use a numbering system which has a truly vast dynamic range. We have developed a system of “address elements” where each element contains a specific portion of the entire expressible dynamic range in the form:

$\{p.x.y.z\}$

where p is the place value, and x , y , and z are the metrics for each dimension.

The address element is currently implemented as a 32-bit construct, so the range of p is 127, and x , y , and z , are unsigned octets. Address elements may be concatenated to any level of resolution desired; as most operations in cyberspace occur within a constrained context, 32, or at most, 64 bits is sufficient to express the vast majority of interactions. This gives the numbering system the twin benefits of wide dynamic range and compactness; compactness is an essential quality in a networked environment.

This is only one possible numbering scheme; others may be developed which conform to the principles as given, perhaps more effectively.

Cyberspace has now been given a universal, unique, dense numbering system; it is now possible to quantify it. The first quantification is that of existence (metrics); the second quantification is that of content. Content is not provided by cyberspace itself, but rather by the participants within it. *The only service cyberspace needs to provide is a binding between a spatial descriptor and a host address.* This can be described by the function:

$f(s) \rightarrow a$

where s is a spatial identifier, and a is an internetwork address. This is the essential mathematical construction of cyberspace.

Implementation of Cyberspace Protocol

If cyberspace is reducible to a simple function, it can be expressed through a transaction-based protocol, where every request yields a reply, even if that reply is NULL. In the implementation under examination, cyberspace protocol (CP) is implemented through a straightforward client-server mechanism, in which there are very few basic operations; registration, investigation, and deletion.

In the registration process, a cyberspace client announces to a server that it has populated a volume of space; in this sense, cyberspace does not exist until it is populated: this is a corollary to Benedikt's *Principle of Indifference*, which states, "absence from cyberspace will have a cost."

The investigation process will be discussed in detail later in this work. The basic transaction is simple: given a circumscribed volume of space, return a set of all hosts which contribute to it. The reply to such a transaction could be NULL or practically infinite (consider the case where the request specifies a volume which describes the entirety of cyberspace); this implies that level-of-detail must be implemented within the transaction (and hence, within registration), in order to optimize the process of investigation. Often, it is enough to know cyberspace is populated, nothing more, and many other times, it is enough to know only the gross features of the landscape, not the particularities of it. In this sense, level of detail is a quality intrinsic to cyberspace.

Registration contains within it the investigation process; before a volume can be registered successfully, "permission" must be received from cyberspace itself, and this must include an active collaboration and authentication process with whatever other hosts help to define the volume. This is an enforcement of the rule which forbids interpenetration of objects within the physical world; it need not be enforced, but unless it is observed in most situations, cyberspace will tend toward being intrinsically disorienting.

Finally, the deletion process is the logical inverse of the registration process, where a volume defined by a client is removed from cyberspace. These three basic transactions form the core of cyberspace protocol, as implemented between the client and the server.

Cyberspace Servers

Cyberspace is a unified whole; therefore, from a transaction-oriented point of view, every server must behave exactly like any other server (specifically with respect to investigation requests). The same requests should evoke the same responses. This would appear to imply that every server must comprehend the "totality" of cyberspace, a requirement which is functionally beyond any computer

yet conceived of, or it places a severe restriction on the total content of cyberspace. Both of these constraints are unacceptable, and a methodology to surmount these constraints must be incorporated into the cyberspace server implementation.

The cyberspace server is implemented as a three-dimensional database with at least three implemented operations; insertion, deletion, and search. These correspond to the registration, deletion, and investigation transactions. Each element within the database is composed of at least three items of data; the volumetric identifier of the space; the IP address of the host which “manifests” within that space; and the IP address of the cyberspace server through which it is registered. The investigation transaction is the core of the server implementation. Cyberspace servers use a repeated, refined query mechanism, which iteratively narrows the possible range of servers which are capable of affirmatively answering an investigation request until the set exactly conforms to the volumetric parameters of the request. This set of servers contains the entire possible list of hosts which collaborate in creating some volume of cyberspace, and will return a non-null reply to an investigation request for a given volume of space.

(Lots of pseudo-code going here...)

An assumption implicit in this algorithm is that investigative searches have “depth”, that investigation is not performed to its exhaustive limit, but to some limit determined by both client and server, based upon the “importance” of the request. Registrations, on the other hand, must be performed exhaustively, but can (and should) occur asynchronously.

The primary side-effect of this methodology is that cyberspace is not instantaneous, but is bounded by bandwidth, processor capacity, and level of detail, in the form:

where c is a constant, the “speed limit” of cyberspace (as c is the speed of light in physical space), l is the level of detail, b is bandwidth of the internetwork, p is processor “capacity”, and s is the “density” of a volume, v , of cyberspace. This expression is intended to describe the primary relationships, between the elements which create cyberspace, and is not mathematically rigorous, but can be deduced from Benidikt’s Law.

Finally, because cyberspace servers do not attempt to contain the entirety of cyberspace, but rather, search through it, based upon client transaction requests, it can be seen that the content of a cyberspace server is entirely determined by the requests made to it by its clients.

One way to visualize the operation of cyberspace servers is with the metaphor of Indra’s Net, from Vedanta Hinduism; finely woven of glittering jewels, each jewel reflecting every other.

Cyberspace and the World Wide Web

Having defined, specified, and implemented an architecture which provides a binding between spatio-location and data set location, this architecture needs to be integrated with the existing WWW libraries so that their functionality can be similarly extended. As “location” is being augmented by the addition of CP to WWW, it is the Universal Resource Locator which must be extended to incorporate these new capabilities.

The URL, in its present definition, has three parts: an access identifier (type of service), a host name (specified either as an IP address or DNS-resolvable name), and a “filename”, which is really more of

a message passed along to the host at the point of service. Cyberspace Protocol fits well into this model, with two exceptions; multiple hosts which collaborate on a space, and the identification of a “filename” associated with a registered volume of space.

We propose a new URL of the following form:

cs://{{pa.x.y.z}}{pb.x.y.z}.../filename

where *{pn...}* is a set of CP address elements.

Resolution of this URL into a data set is a two-stage process: first the client CP mechanism must be used to translate the given spatio-location into a host address, then the request *must* be sent to the host address. Two issues arise here; multiple host addresses, as mentioned previously, and a default access mechanism for CP. If a set of host addresses are returned by CP, a request must be sent to each specified host; otherwise, the description of the space will be incomplete. Ideally, all visualized WWW clients will implement a threaded execution mechanism (with re-entrant WWW libraries) so that these requests can occur simultaneously and asynchronously.

A default access mechanism for CP within WWW must be selected. The authors have chosen HTTP, for two reasons; it is efficient, and it is available at all WWW servers. Nonetheless, this is not a closed issue; it may make sense to allow for some variety of access mechanisms, or perhaps a fallback mechanism; if one service is not present at a host, another attempt, on another service, could be made.

(More stuff about filenames in cyberspace going here...)

Labyrinth

It is now possible, from the previous discussion, to describe the architecture and operation of a fully visualized WWW client. It is composed of several pieces; WWW libraries with an integrated CP client interface; an interpreter for an HTML-derived language which describes object geometry, placement, and linkage; and a user interface which presents a navigable “window on the web”.

The operation of the client is very straightforward, as is the case of the other WWW clients. After launching, the client queries the “space” at “home”, and loads the “world” as the “axis mundi” of the client’s view of the web. As a user moves through cyberspace, the client makes requests, through CP, to determine the content of all spaces passed through or looked upon. A great deal of design effort needs to be put into the development of look-ahead caching algorithms for cyberspace viewers; without them, the user will experience a discontinuous, “jerky” trip through cyberspace. The optimal design of these algorithms will be the subject of a subsequent work.

At this time, visualized objects in WWW have only two possible behaviors; no behavior at all, or linkage, through an attached URL, to another data set. This linkage could be to another “world” (actually another place in cyberspace), which is called a “portal”, or it could link to another data type, in which case the client must launch the appropriate viewer. Labyrinth is designed to *augment* the functionality of existing WWW viewers, such as NCSA Mosaic, rather than to supplant them, and therefore does not need a well-integrated facility for viewing other types of HTML documents.

Data Abstraction Protocols

Cyberspace Protocol is a specific implementation of a general theory, which has implications well beyond WWW. CP is the solution, in three dimensions, of an N-dimensional practice for data set location abstraction. Data abstraction places a referent between the “name” of a data set locator and the physical location, allowing physical data set location to become mutable.

If an implementation were to be developed for the case where $N = 1$, it would be an effective replacement Internet’s Domain Name Service, which maintains a static mapping of “names” to IP addresses. Any network which used a dynamic abstraction mechanism could mirror or reassign hosts on a continuous basis (assuming that all write-through mirroring could be maintained by the hosts themselves), so that the selection of a host for a transaction could be made *based upon criteria*

that would tend to *optimize the performance of the network* from the perspective of the transaction. It would also be easy to create a data set which could “follow” its user(s), adjusting its location dynamically in response to changes in the physical location or connectivity of the user. In an age of wireless, worldwide networking, this could be a very powerful methodology.

Conclusions

This work attempts to outline the requirements for architectures which can fully visualize WWW, and proposes solutions to the issues raised by these requirements. While much further study needs to be done, this work is meant to serve as a starting point for an understanding of the subtleties of wide-area, distributed, visualized data sets.

Labyrinth and Cyberspace Protocol are logical extensions to the World Wide Web and Internet. Indeed, without the existence of WWW, neither would be very useful immediately; they would operate, but lack content, and individuals would hardly be compelled either to use them or adapt their existing data sets to realize their new potentials. Used together, they work to make both WWW and Internet inherently more navigable, because they help to make Internet more human-centered, adapting data sets to human capabilities rather than vice versa. This, thus far, is the single largest contribution that “virtual reality” research has offered to the field of computing; a human-centered design approach that lowers or erases the barriers to usage by creating user-interface paradigms which serve humans to the full of their potential.

Finally, network visualization marks the end of the “first age” of networking, where protocols, services, and infrastructure dominated the discourse within the field. In the “second age” of networking, questions like data architecture and the inherent navigability of a well-designed data set become infinitely more important than “first age” questions; where “how do I find what I’m looking for?” becomes more relevant than “where did it come from?”

Acknowledgments

The authors would like to thank the following individuals, who contributed their own thoughts to the formation and development of the ideas expressed in this work: Michael J. Donahue, Owen Rowley, Clayton Graham, Christopher Morin, Neil Redding, James Curnow, Marina Berlin, and the *Fugawi* tribe.