Journal of Universal Computer Science, vol. 1, no. 4 (1995), 206-220
submitted: 13/8/94, accepted: 3/10/94, appeared: 28/4/95, ©Springer Pub. Co.

The Hyper-G Network Information System

Keith Andrews, Frank Kappe, and Hermann Maurer
(Institute for Information Processing and Computer Supported New Media (IICM)
Graz University of Technology
A-8010 Graz, Austria.
{kandrews,fkappe,hmaurer}@iicm.tu-graz.ac.at)

Abstract As the Internet continues to experience exponential rates of growth, attention is shifting away from mainstream network services such as electronic mail and file transfer to more interactive information services. Current network information systems, whilst extremely successful, run into problems of fragmentation, consistency, scalability, and loss of orientation.
The development of "second generation" network information systems such as Hyper-G can help overcome these limitations. Of particular note are Hyper-G's tightly-coupled structuring, linking, and search facilities, its projection of a seamless information space across server boundaries with respect to each of these facilities, and its support for multiple languages. The Harmony client for Hyper-G utilises two and three-dimensional visualisations of the information space and couples location feedback to search and link browsing operations, in order to reduce the likelihood of disorientation. This paper presents a comprehensive overview of Hyper-G and Harmony.
Key Words: Hypermedia, information system, information visualisation, graphical interaction, Internet.
Category: H.5

1 Introduction

The Internet, the world-wide computer network, now connects more than 3.8 million individual computers (October 1994) with a growth rate of between 10–15% per month [Network Wizards 1994]. More than 18 terabytes of information traversed the NSFNET, the main US backbone of the Internet, in the month of September 1994 [Merit NIC 1994]. Whilst traditional services such as electronic mail, remote login, and file transfer (FTP) still account for the bulk of Internet traffic, by far the fastest growth is being experienced by network information systems like WAIS, Gopher, and the World-Wide Web.

These three information systems have transformed the way people perceive and interact with information resources on the net. All three are client-server applications using the Internet's underlying TCP/IP protocol: end-users run client software (available for a wide variety of platforms) which communicates over the network with servers managing access to a vast amount of extremely diverse information.

WAIS (Wide Area Information Servers), began in 1989 as a joint development of Thinking Machines, Apple Computer, and Dow Jones to provide on-line access to the Wall Street Journal [Kahle et al. 1992]. WAIS supports powerful content-based search of (previously) indexed databases, including *relevance feedback* by which (parts of) text documents returned by a search and deemed to be particularly relevant by the user are used as input for a subsequent search, in effect refining the search by asking WAIS to look for further similar documents. WAIS is purely a search engine, it supports neither associative browsing (hyperlinks) nor any structuring of its information content.

Gopher was started in 1991 as a campus-wide information system at the University of Minnesota [McCahill and Anklesaria 1995, Anklesaria et al. 1993]. It provides menu-like access to information resources. Although Gopher space is in fact a graph containing many loops, the menu presentation gives the impression of a tree. Typically, users begin navigation at the top of the tree and traverse down to leaf nodes containing actual data. One sub-tree on each server usually contains a menu of other Gopher servers. Gopher has no integrated search facilities of its own, but provides access to optional add-on search engines such as WAIS; it has no provision for hyperlinks.

The World-Wide Web project (WWW, W3, or simply "The Web") was initiated at CERN, Geneva, in 1989, originally as an information system for the particle physics community [Cailliau 1995, Berners-Lee et al. 1994]. W3 is a distributed *hypermedia* system: combining the concept of *hyperlinks* (associative browsing by following links to related information) with *multimedia* (text, image, audio, video, etc.). Through its URL (Universal Resource Locator) mechanism, W3 can represent links to any document on any W3, Gopher, or FTP server worldwide. The CGI script interface allows W3 servers to start arbitrary application programs, for example linking into external databases or implementing complex search algorithms. Simple, intuitive W3 clients such as Mosaic and Netscape have contributed to its tremendous current popularity.

However, W3 has a number of limitations. It does not provide any information structuring facilities beyond hyperlinks; its links are one-way (there is no way of determining which other documents refer to a particular document, leading to inconsistencies when documents are moved or deleted – the frequent "dangling links") and embedded within text documents (there are no links from other kinds of documents). Like Gopher, W3 has no native search facilities, but relies on external search engines such as WAIS, leading to patchy server-by-server provision of search facilities by individual sites and no real-time cross-server searches (searches in previously generated cross-server indices are available). The flexibility provided by CGI is achieved at great cost: the uniformity of the interface disappears, different W3 servers behave differently – resulting in the "Balkanisation" (to quote Ted Nelson) of the Web into independent "W3 Empires". Also, there is little support for the maintenance of large datasets, so it is not uncommon to see several W3 servers within a single organisation, each a fundamentally separate interactive context. The Web today is very much "read-only", in the sense that information providers prepare data sets in which information consumers can generally only browse. Finally, although its URL mechanism endows W3 with scalability in terms of number of servers, it is not scalable in terms of number of users. Extremely popular W3 servers such as Sun Microsystems' World Cup USA '94 [Sun Microsystems] site can often become overwhelmed by tens of thousands of users, necessitating their physical mirroring to many alternative sites.

WAIS, Gopher, and W3 belong to the first generation of information systems on the Internet [Fenn and Maurer 1994]. They work well in particular contexts but run into difficulties when applied to hundreds of thousands of documents distributed over many thousands of servers. They provide no graphical navigation aids, only rudimentary access control, little support for automatic database maintenance, no scalable document replication mechanisms, and little support for multiple languages. In the rest of this paper we describe Hyper-G, a second generation system designed to transcend some of the limitations of existing environments.

2 Hyper-G

Hyper-G is a multi-user, multi-protocol, structured, hypermedia information system, which runs as a client-server application on the Internet [Andrews and Kappe 1994, Andrews et al. 1994, Kappe 1993].

2.1 Design Goals

Based on an analysis of the strengths and weaknesses of existing systems, the following primary design goals were formulated:

1. Provide orientational and navigational aids.
2. Provide automatic structuring and maintenance.
3. Reduce fragmentation across servers.
4. Support user identification and access control.
5. Support multilinguality.
6. Maintain interoperability with existing systems.

To help alleviate the disorientation associated with becoming "lost in hyperspace", Hyper-G provides three closely coupled, orthogonal navigational mechanisms: structuring, hyperlinks, and search. The tight coupling allows clients to correlate search results, link maps, and structure overviews, providing a powerful aid to navigation.

Users of W3 are soon confronted with the "dangling link" syndrome inherent in W3's data model: when a document is moved or deleted, there is no way of automatically updating or deleting links to that document – following such links produces an error message. In Hyper-G, the consistency of links is maintained automatically by the server, as are Hyper-G's aggregate structures.

To counter the Balkanisation of databases, Hyper-G provides much functionality integrated into it (and hence uniform in nature) which has to be implemented on top of W3 or Gopher (and hence potentially differing from site to site), in particular powerful search and retrieval facilities which can be performed simultaneously and seamlessly *across* server borders. Information can also be *structured* both within and across server borders.

In order to differentiate between users and provide tailored views of the available information, support at the system level for user accounts, user groups, access rights, and access modes is essential. They are also a prerequisite for the implementation of charging facilities.

All of the currently popular Internet information systems are more or less implicitly bound to a single language: English. Although the Gopher+ and W3 protocols in theory support multi-language versions of documents, most clients do not use these facilities. Similarly, elements in the user interface are relatively easily translated to other languages, but most clients support only a single interface language; the same applies to full text search in multiple languages. Hyper-G and its clients, on the other hand, were specifically designed for multilingual use.

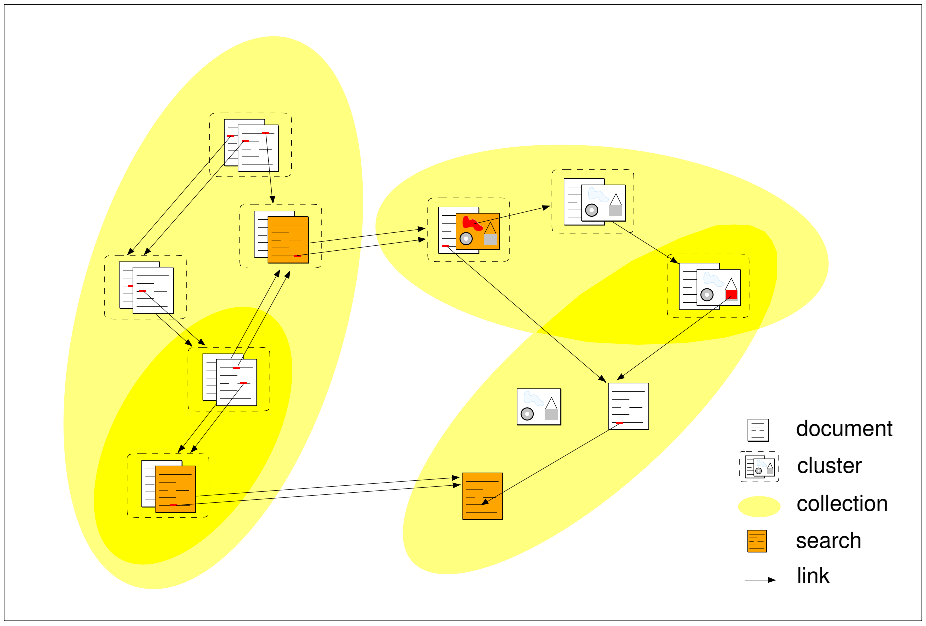

Figure 1: The Hyper-G Data Model

Hyper-G was designed to overcome the limitations of existing tools like WAIS, Gopher, and W3, while at the same time maintaining interoperability with them — interoperability was considered crucial to the acceptance of a system in today's multi-protocol Internet.

2.2 Design Features

The most fundamental design decision was to provide support for orthogonal yet closely coupled structuring, linking, and search facilities at the database level (as shown in Figure 1):

- *Structuring*: Documents may be grouped into aggregate *collections*, which may themselves belong to other collections. Every document must belong to at least one collection. Navigation may be performed down through the collection hierarchy (the collection "hierarchy" is, strictly speaking, a directed acyclic graph), access rights assigned on a collection-by-collection basis, and the scope of searches restricted to particular sets of collections. Collections may span multiple Hyper-G servers, providing a unified view of distributed resources.
- *Linking*: Hyperlinks connect a *source anchor* within one document to either a *destination anchor* within another document, an entire document, or a collection. Links are not stored within documents (as in W3) but in a separate link database (as pioneered by Intermedia [Haan et al. 1992]). This has a number of important advantages:

    - Source and destination anchors are not limited to text documents, but can be attached to any kind of media (image, audio, film, 3D scene, formatted PostScript document, etc.).
    - Links can be attached to otherwise read-only documents (for example documents on CD-ROM or with read-only access rights).
    - Links can be followed backwards.
    - A local map (fish-eye view) can be readily computed and visualised using the link database.
    - Consistency constraints can be met more easily (for example when moving or deleting a document, it is important to know which other documents contain links to it).
- *Attribute and Content Search*: Documents and collections have an associated set of attributes (author, title, keywords, etc.) which may be searched for, including boolean combinations and term truncation. Full text (content) search facilities include vector and fuzzy boolean queries. Every document and collection is automatically indexed upon insertion into the database – no extra indexing steps are required. The scope of a search may be focussed to one or more collections on one or more servers or may be as wide as all collections on all Hyper-G servers worldwide.

Each of these three orthogonal features work *seamlessly* across server boundaries, reducing fragmentation while at the same time promoting consistency. To achieve scalability in terms of number of users, Hyper-G servers replicate and cache remote objects; updates are propagated to other servers using an efficient, scalable flooding algorithm [Kappe 1995].

A special kind of collection, a *cluster*, groups documents into logical entities. Clusters are used both to define multimedia aggregates (for example, a text and an associated image or video clip) which are presented together and multilingual aggregates (for example, English and German translations of a text, audio clip, structure diagram, or any combination of these).

Other design features supported by Hyper-G and not found in comparable systems include:

- Anonymous and identified user identification modes.
- A scheme of user groups and subgroups maintained by the server.
- Access rights for users and user groups on a document or collection basis.
- "Home collections", personal information spaces for identified users (kept on the server) used to organise personal documents and pointers to resources.
- Language preferences, applied both to document retrieval and to the user interface.
- An underlying object-oriented database, which guarantees the consistency and integrity of data (for example the updating of links when a document is moved or deleted).

2.3 Architecture

Figure 2 shows the architecture of Hyper-G. Note the interoperability of Hyper-G with Gopher and W3 clients and servers. When accessed by a Gopher client, the Hyper-G server maps the collection hierarchy into a Gopher menu tree (hyperlinks cannot be represented in Gopher). A synthetic search item is

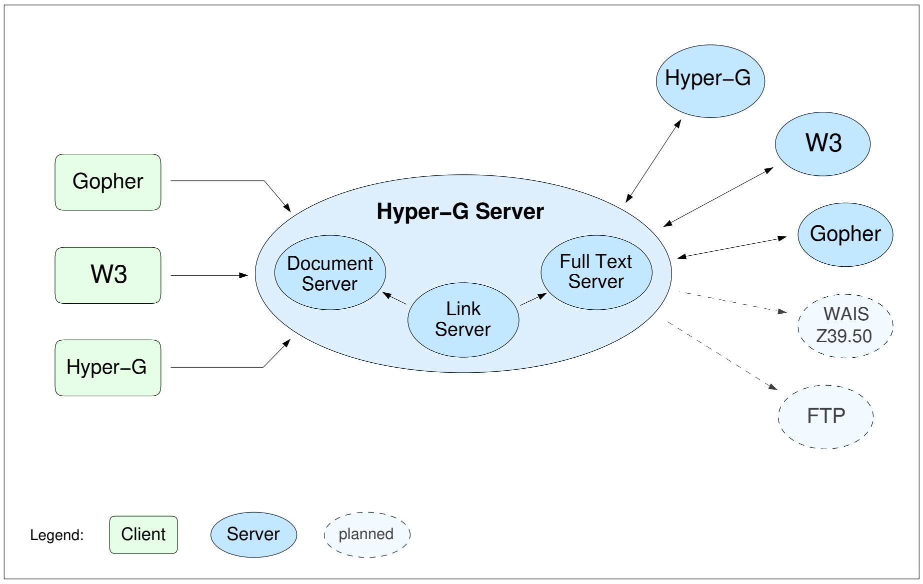

Figure 2: The Architecture of Hyper-G

generated at the foot of each Gopher menu to allow searching the corresponding collection. When accessed by a W3 client, each level of the collection hierarchy is converted to an HTML [Berners-Lee and Conolly 1993] document containing a menu of links to its members. Hyper-G text documents are transformed on-the-fly into HTML documents, including any links they might have. Additional Hyper-G functionality such as user identification, language preference selection, and searching are implemented via HTML forms and are accessible at any time.

The Hyper-G server is able to store pointers to remote objects on Gopher and W3 servers. This allows the incorporation of information on remote non-Hyper-G servers (almost) seamlessly: Gopher menus are transformed into Hyper-G collections and W3 text documents into Hyper-G text documents. Interoperability with WAIS (Z39.50) and FTP servers is planned.

Unlike Gopher or W3 clients which connect to many servers during a typical session, Hyper-G clients talk to a single Hyper-G server for the entire session. Should information from a remote server be needed, the local server acts as a proxy, i.e. it fetches the object and passes it on to the client. This approach has the following advantages:

- Clients are kept simple, the Hyper-G server handles external protocols.
- Remote information can be cached in the local server.
- User accounts and access rights have only to be maintained on the local server (the user has to identify to one server only).
- Statistics and user profile information can be gathered on a per-session basis.

Hyper-G clients connect to a Hyper-G server using the assigned port number 418 by default. Port 418 is used for control information, documents are

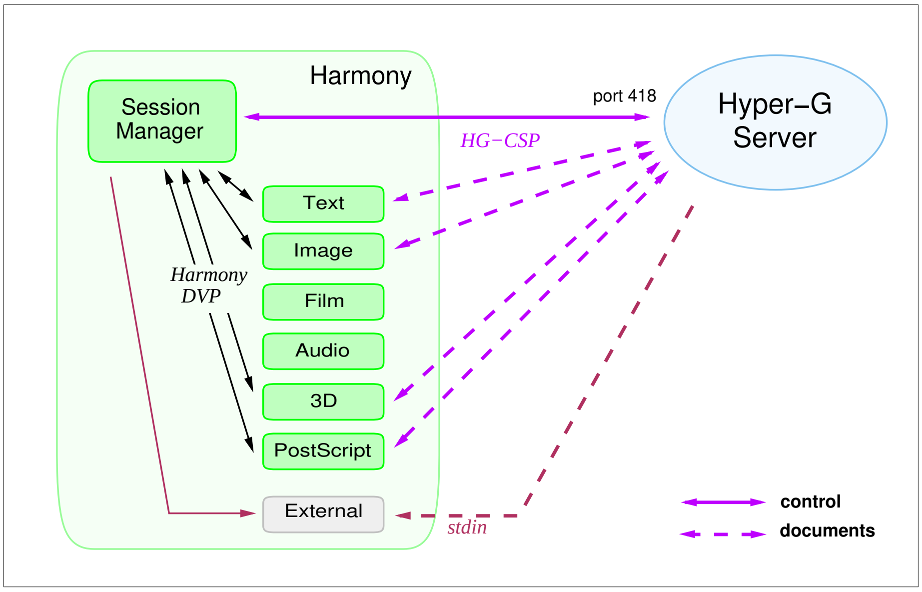

Figure 3: The Architecture of the Harmony Client for Hyper-G

generally sent (simultaneously in the case of multimedia clusters) using dynamically assigned port numbers. The precise mechanism is described in the Hyper-G Client/Server Protocol (HG-CSP) specification [Kappe and Pani 1994].

3 The Harmony Client for Hyper-G

Harmony is the native Hyper-G client for X Windows on Unix platforms. It takes advantage of Hyper-G's underlying facilities to provide intuitive navigational tools and informative feedback about the location of information.

As can be seen in Figure 3, Harmony is a multi-process application: the primary process, the *Session Manager*, communicates with the Hyper-G server, provides navigational facilities, and coordinates all other activities. The session manager starts secondary processes, *document viewers*, as necessary to display particular documents. Native Harmony document viewers conform to the Harmony Document Viewer Protocol (DVP) [Andrews et al. 1995], which defines various browsing, editing, and link functions. There currently exist native document viewers for text, images, MPEG films, audio, 3D scenes, and PostScript. Harmony may be configured to run external programs in place of any native viewer and also for unsupported document types (the document is piped to standard input), however without provision for link activation and editing.

Figure 4 shows a typical Harmony session. The Session Manager (top left) provides navigation through the collection structure, search facilities, and various general functions such as user identification and language selection. Collections

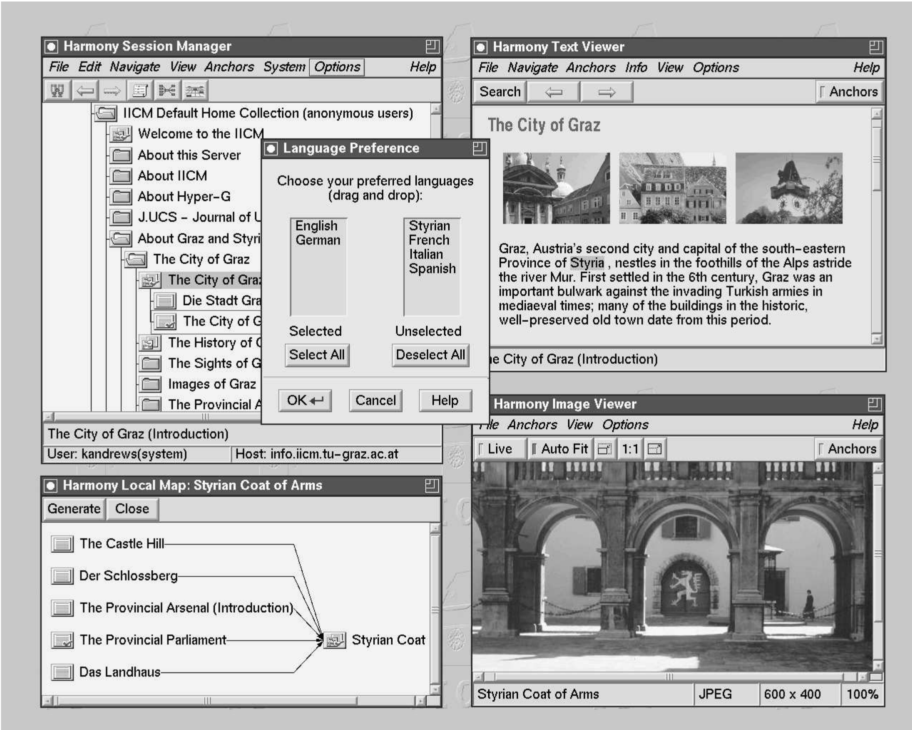

Figure 4: Harmony – The Hyper-G Client for X Windows

may be opened and closed and clusters or individual documents activated within the graphical collection display by double-clicking. Collections, clusters, or documents which have already been visited are marked with a tick. In this example, a descriptive text and an image about the city of Graz have been accessed.

The Harmony Text Viewer (top right) uses a generic SGML parser to display marked-up text documents (both Hyper-G's HTF and W3's HTML formats), and has the usual facilities for scrolling, finding strings, selecting, etc. Hyperlinks within the text can be highlighted in a number of ways. Inline images in TIFF, GIF, and JPEG formats are supported.

The Harmony Image Viewer (bottom right) accepts raster images in a variety of common formats (TIFF, GIF, JPEG, etc.). Operations such as zooming and panning are available. Link anchors are rectangular, circular, or elliptical (soon also polygonal) areas, which are overlaid atop the image. All Harmony document viewers support both interactive link following and definition.

Harmony's Film Player (not shown) displays MPEG (MPEG-1, soon MPEG-2) video streams. It is possible to define a link anchor which follows an object of interest in the video: rectangular and circular anchor regions overlaying the film are simply defined for specific keyframes (they are interpolated in between keyframes). Anchors can be activated both during playback and while playback is paused. The Harmony Audio Player (not shown) can be configured to use either the Network Audio Server [Fulton and Renda 1994] or local audio commands to play audio files in a variety of common formats.

The Local Map facility (bottom left), provides a kind of short-range radar, generating on request (dynamically) a map of the link relationships of a chosen document, similar to the local map of Intermedia [Haan et al. 1992]. By default, two levels of incoming and outgoing hyperlinks are represented. One can navigate within the local map by selecting (single-clicking) another object toward the edge of the map and generating a new display. Objects can be activated by double-clicking.

The language preference dialogue (centre) allows the user to specify an ordered list of preferred languages. Harmony's user interface adjusts dynamically to the language of first choice (English and German interfaces are currently supported), documents available in multiple languages are selected in order of language preference, and searches are optionally language-dependent.

Central to the design of Harmony is the concept of *location feedback*. When a document or collection is visited, its location within the collection structure is automatically displayed in the Session Manager's collection browser (by opening up the path to it), regardless of whether the object was reached as the result of a search, by following a hyperlink, or via the local map. This unique feature of Harmony is a powerful instrument in the fight against becoming "lost in hyperspace" – users can orient themselves with reference to a fixed structural framework. In the case of search results and the local map, mere selection of an object initiates location feedback, providing users with a sense of the context of an object, prior to any decision to view it.

The Harmony search dialogue (not shown) provides an interface to the full range of Hyper-G searches: attribute and content, boolean and fuzzy. The scope of a search may be focussed to a single collection or set of specific collections (possibly spanning server boundaries), or may be as wide as all collections on the local server. Search results are presented as a ranked list. As noted above, Harmony applies the principle of location feedback to search results: selecting an object in the result list causes its location in the collection structure to be displayed, hence users can make informed choices before committing to fetch particular documents.

A further innovative feature of Harmony is its use of three-dimensional visualisations, both hand-crafted and automatically generated [Andrews 1993, Andrews and Pichler 1994]. Model description files representing arbitrarily complex scenes or objects are displayed by the Harmony 3D Scene Viewer. Figure 5 shows the scene viewer displaying a model of the Great Hall of the Austrian National Library (top left) and a staute of Kaiser Karl VI (bottom left). Users typically view a model of a scene by moving themselves (walk, fly, fly to, heads-up) and view a model of an object by moving the model (translate, rotate, zoom). 3D models are fully-fledged hypermedia documents: hyperlinks may be attached to individual objects within a scene or to groups of polygons within an object.

Harmony's Information Landscape, also shown in Figure 5, is an interactive, three-dimensional visualisation of the collection structure, tightly coupled to the Session Manager's two-dimensional collection browser display (changes in one are reflected in the other). The collection hierarchy is mapped out onto a plane, documents within a collection are arranged on top of the corresponding block;

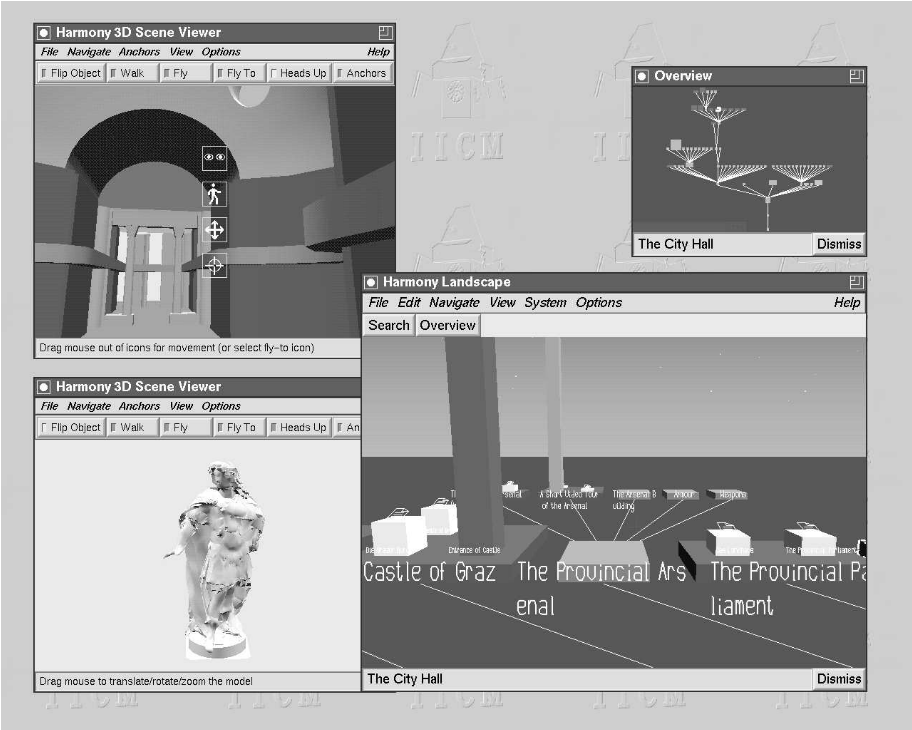

Figure 5: Harmony Landscape and 3D Scene Viewer

colour and height are used to encode document type and size respectively. Users can “fly” over the landscape looking for salient features, like flying over a file system with FSN [Tesler and Strasnick 1992]. A flat overview window (upper right) provides a further aid to orientation. Through their ability to compactly display many thousands of objects, 3D visualisations are perhaps the only effective means of browsing in and judging the extent of large, dynamic information spaces.

The Harmony PostScript Viewer (Figure 6) displays arbitrary documents in PostScript format; the documents are typically stored and transmitted in compressed form and uncompressed locally by the viewer. Rectangular link anchors are supported.

As was mentioned in the opening discussion, Hyper-G clients can be used to *edit* as well as browse the contents of a Hyper-G server, in so far as the user is identified and has appropriate access rights. Figure 7 shows the Harmony Insert Dialogue being used to upload a text document about Graz from the local file system into the user’s home collection. Figure 8 shows a link being created from the word Styria in the text about Graz to a multimedia cluster about the Austrian Province of Styria.

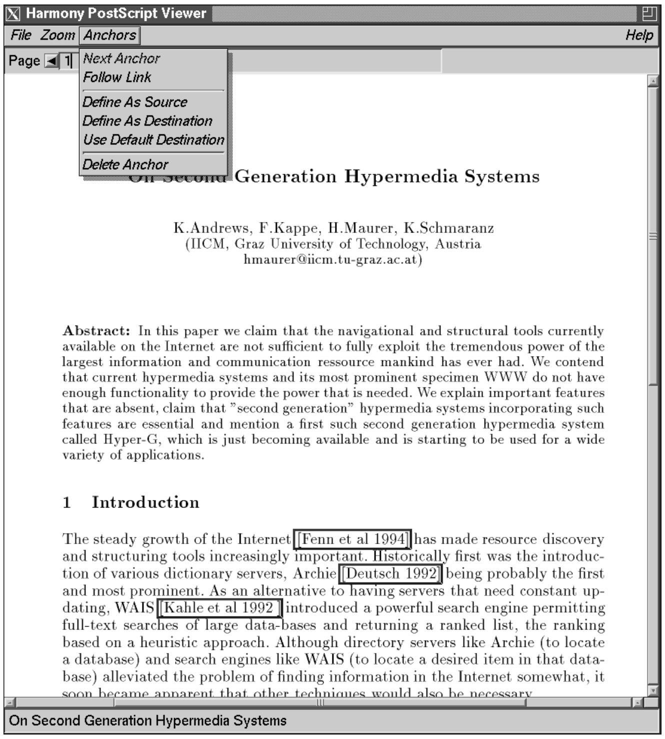

Figure 6: The Harmony PostScript Viewer

Finally, Harmony's History Browser (not shown) offers a timeline of past interactive waypoints, including previous search panels – another means of orienting in hyperspace.

Features to be implemented in Harmony over the coming months include interactive forms, a drag-and-drop interface to the local file system (allowing documents to be simply pulled into collections on the server), integrated electronic mail facilities, semi-automatic link generation, and three-dimensional representations of hyperlink relationships and search results. An immersive (virtual reality) interface is also planned.

4 Concluding Remarks

We have presented the design rationale and the current development status of Hyper-G and its Harmony client for X Windows. Of particular note are Hyper-G's tightly-coupled collection, link, and search facilities, its projection of a seamless information space across server boundaries, and its support for multiple languages. Harmony makes innovative use of location feedback and two

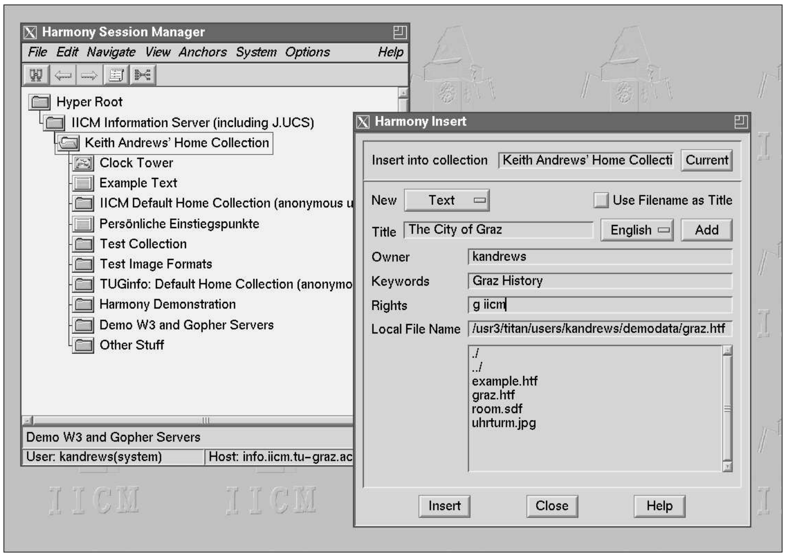

Figure 7: The Harmony Insert Dialogue

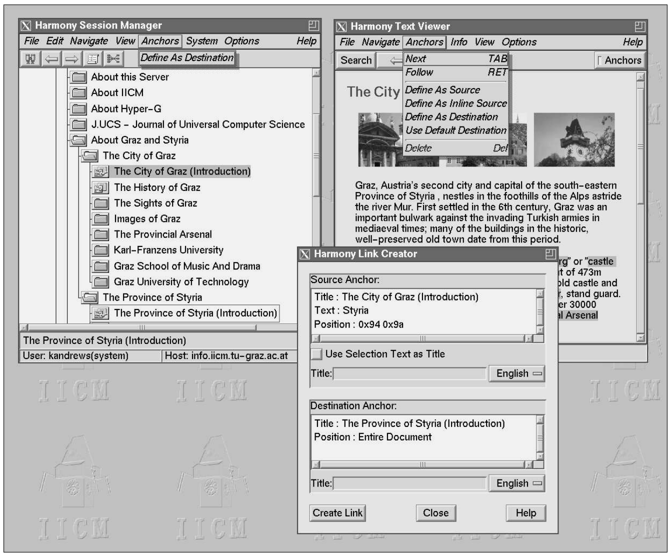

Figure 8: The Harmony Link Creator

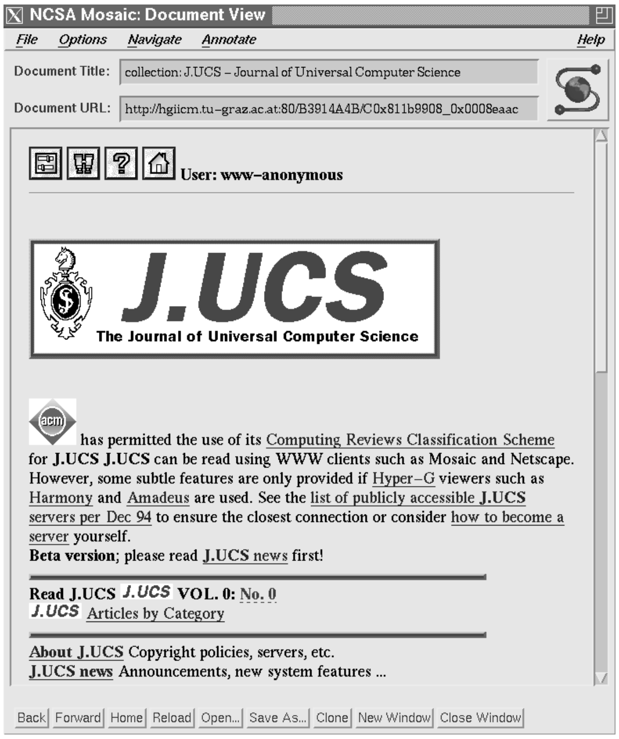

Figure 9: Accessing a Hyper-G Server with Mosaic

and three-dimensional visualisations to help users navigate and orient themselves within large, dynamic information spaces.

A key design goal of Hyper-G was to maintian interoperability with current information systems. The combination of the Hyper-G server's rich structuring and maintenance facilities and its ability to service requests from W3 and Gopher clients in addition to native Hyper-G clients make it ideal for use as a multi-protocol server. Figure 9 shows the Mosaic client for W3 being used to access the J.UCS collection on the IICM's Hyper-G server.

Although still in its infancy, Hyper-G has already gained considerable acceptance. The European Space Agency and the Museum of New Zealand have adopted Hyper-G for their own information systems, the Austrian Ministry of Science has adopted it as its information system of choice for all Austrian universities, and the German Mathematics Association (DMV) is setting up an information system spanning most German universities and colleges, to name just a few examples.

Hyper-G also is being used as the basis for a major new electronic publishing venture. The Journal of Universal Computer Science (J.UCS), supported by Springer Verlag, is among the first high-quality, fully-refereed, fully-citable scientific journals to depend primarily on Internet distribution [Maurer and Schmaranz 1994]. The pilot issue is already available at several sites, the first regular issue will be available world-wide at the end of January 1995.

In addition to the Harmony client described in this paper, native Hyper-G clients are available for Unix VT100-style terminals (hgtv) and MS-Windows (Amadeus), and a client is under development for the Macintosh. Further information about Hyper-G and Harmony and installation details may be retrieved by anonymous ftp from `ftp.iicm.tu-graz.ac.at` in directory `/pub/Hyper-G` or from the IICM Information Server under `http://info.iicm.tu-graz.ac.at/` or `gopher://info.iicm.tu-graz.ac.at/`.

Acknowledgements

Financial support of Hyper-G by the Austrian Ministry of Science, JOANNEUM RESEARCH, and the European Space Agency is gratefully acknowledged.

References

[Andrews 1993] Andrews, K.: "Constructing Cyberspace: Virtual Reality and Hypermedia". Presented at Virtual Reality Vienna '93. Available by anonymous ftp from `iicm.tu-graz.ac.at` in directory `pub/Hyper-G/papers` (1993).

[Andrews et al. 1995] Andrews, K., Faschingbauer, J., Schipflinger, J.: "The Harmony Document Viewer Protocol (DVP)". In preparation (1995).

[Andrews and Kappe 1994] Andrews, K., Kappe, F.: "Soaring Through Hyperspace: A Snapshot of Hyper-G and its Harmony Client". In Herzner, W., Kappe, F. (editors), Proc. of Eurographics Symposium on Multimedia/Hypermedia in Open Distributed Environments, Graz, Austria. Springer (1994), 181–191.

[Andrews et al. 1994] Andrews, K., Kappe, F., Maurer, H., Schmaranz, K.: "On Second Generation Hypermedia Systems". Journal of Universal Computer Science (Pilot Issue), 0, 0 (1994), 127–135. Available at `http://info.iicm.tu-graz.ac.at/Cjucs_root`.

[Andrews and Pichler 1994] Andrews, K., Pichler, M.: "Hooking Up 3-Space: Three-Dimensional Models as Fully-Fledged Hypermedia Documents". In Proc. of East-West International Conference on Multimedia, Hypermedia, and Virtual Reality (MHVR'94), Moscow, Russia (1994), 11–18.

[Anklesaria et al. 1993] Anklesaria, F., McCahill, M., Lindner, P., Johnson, D., Torrey, D., Alberti, B.: "RFC 1436: The Internet Gopher Protocol". Available by anonymous ftp from `ds.internic.net` as `/rfc/rfc1436.txt` (1993).

[Berners-Lee et al. 1994] Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H. F., Secret, A.: "The World-Wide Web". Communications of the ACM, 37, 8 (1994), 76–82.

[Berners-Lee and Conolly 1993] Berners-Lee, T., Conolly, D.: "Hypertext Markup Language (HTML)". Available on the WWW at URL `http://info.cern.ch/hypertext/WWW/Markup/HTML.html` (1993).

[Cailliau 1995] Cailliau, R.: "About WWW". Information Processing and Management (1995). Special Issue: Selected Proc. of Distributed Multimedia Systems '94, Graz, Austria, Nov. 1994.

[Fenn and Maurer 1994] Fenn, B., Maurer, H.: "Harmony on an Expanding Net". Interactions, 1, 4 (1994), 26–38.

[Fulton and Renda 1994] Fulton, J., Renda, G.: "The Network Audio System". Available by anon. ftp from `ftp://ftp.x.org/contrib/audio/nas/xcon94paper.ps.gz` (1994).

[Haan et al. 1992] Haan, B. J., Kahn, P., Riley, V. A., Coombs, J. H., Meyrowitz, N. K.: "IRIS Hypermedia Services". Communications of the ACM, 35, 1 (1992), 36–51.

[Kahle et al. 1992] Kahle, B., Morris, H., Davis, F., Tiene, K., Hart, C., Palmer, R.: "Wide Area Information Servers: An Executive Information System for Unstructured Files". Electronic Networking: Research, Applications and Policy, 2, 1 (1992), 59–68.

[Kappe 1993] Kappe, F.: "Hyper-G: A Distributed Hypermedia System". In Leiner, B. (editor), Proc. INET '93, San Francisco, California. Internet Society (1993), DCC–1–DCC–9.

[Kappe 1995] Kappe, F.: "A Scalable Architecture for Maintaining Referential Integrity in Distributed Information Systems". Available by anonymous ftp from `ftp.iicm.tu-graz.ac.at` as file `/pub/Hyper-G/papers/p-flood.ps`, Graz University of Technology, Austria (1995).

[Kappe and Pani 1994] Kappe, F., Pani, G.: "Hyper-G Client/Server Protocol (HG-CSP)". Technical report, IICM, Graz University of Technology (1994). Available by anonymous ftp from `ftp.iicm.tu-graz.ac.at` in file `/pub/Hyper-G/papers/Protocol.ps.gz`.

[Maurer and Schmaranz 1994] Maurer, H., Schmaranz, K.: "J.UCS – The Next Generation in Electronic Journal Publishing". Computer Networks for Research in Europe, 26 (1994), S63–S69. Supplement to Vol. 26 of Computer Networks and ISDN Systems.

[McCahill and Anklesaria 1995] McCahill, M. P., Anklesaria, F. X.: "Evolution of Internet Gopher". Information Processing and Management (1995). Special Issue: Selected Proc. of Distributed Multimedia Systems '94, Graz, Austria, Nov. 1994.

[Merit NIC 1994] Merit NIC: "NSFNET Backbone Statistics". Available by anonymous ftp from `nic.merit.edu` in `/nsfnet/statistics` (1994).

[Network Wizards 1994] Network Wizards: "Internet Domain Survey". Available on the WWW at URL `http://www.nw.com/zone/WWW/top.html` (1994).

[Sun Microsystems] Sun Microsystems: "World Cup USA '94". Available on the WWW at URL `http://sunsite.doc.ic.ac.uk/wc94/`.

[Tesler and Strasnick 1992] Tesler, J., Strasnick, S.: "FSN: The 3D File System Navigator". Silicon Graphics, Inc., Mountain View, CA (1992). Available by anonymous ftp from `sgi.sgi.com` in directory `sgi/fsn`.