Aspects of a Modern Multi-Media Information System

Dissertation
for the Award of the Academic Degree
Doctor of Technical Sciences
at the
Graz University of Technology
submitted by

Frank M. Kappe

Institute for
Foundations of Information Processing and
Computer Supported New Media (IICM),
Graz University of Technology

Graz, June 1991

First Reader: O. Univ.-Prof. Dr. Hermann A. Maurer
Second Reader: A. O. Univ. Prof. Dr. Volkmar Haase

Aspects of a Modern Multi-Media Information System

Frank M. Kappe

Abstract

This thesis describes a new, large-scale Hypermedia project ("Hyper-G") currently being developed at the Institute for Foundations of Information Processing and Computer Supported New Media (Head: Professor Hermann Maurer) of the Technical University of Graz. Experience gathered from modern Hypermedia systems, large-scale information systems, computer aided instruction and user interface design considerations led to a number of ideas, features, and examples of applications of Hyper-G.

They were condensed, put into a logical relationship, and used to formulate a set of requirements. The requirements, additional design decisions, and a discussion of implementation-related issues are part of this thesis.

Also, a new concept for the creation of real-time, interactive animation is presented. It is essentially a combination of Computer Animation and Hypermedia technologies, therefore it is called "Hyper-Animation".

This concept is also the basis of some of the more advanced applications of Hyper-G that are described in this thesis. Applications range from information systems and electronic publishing to exhibits that may be found in a virtual museum or exhibition environment.

Acknowledgments

The support of parts of the research reported in this thesis by the Austrian Federal Ministry for Science and Research is gratefully acknowledged. In addition, a number of individuals supported me in research and preparation of the thesis:

Prof. Hermann Maurer, thesis supervisor, introduced me to Hypermedia and suggested the topic(s) of the thesis. Numerous fruitful discussions with him and others, including Prof. Ivan Tomek from Acadia University, Nova Scotia, Prof. Dieter Fellner from the Memorial University of Newfoundland and institute members Helmut Mülner, Gerald Pani, Peter Sammer, and Robert Stubenrauch led to ideas contained herein. Prof. Haase's fast reading allowed last-minute changes and completion within schedule.

Last, but not least, I want to thank my wife Sigrid for her love and patience during the preparation of this thesis.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

What is this thesis about?

To put is simple: This thesis is about Hypermedia.

1.1 What is Hypermedia?

This question cannot be answered trivially. The origin of the Hypermedia idea can be traced back to 1945, although it was not called Hypermedia at that time. The idea had no big impact then, because it could not be implemented in a satisfactory way with the available hardware and software.

However, Hypermedia has more than one root. The term itself is derived from Hypertext and Multimedia, but it seems that a number of different disciplines of computer science (and others) – including Information Retrieval, Computer Based Learning, Computer Graphics, Human-Computer Interactions, User Interface Design, Electronic Publishing, Communication Systems, and Cognitive Psychology – have evolved in such a way that eventually they met at Hypermedia.

This is why now, after almost 50 years, there is so much interest in Hypermedia. The time of Hypermedia has come. And it is also the reason why Hypermedia means different things to different people:

- For some people (typically those concerned with information retrieval), Hypertext is a way of organizing and retrieving information. The typical Hypertext application for this group is encyclopedia-like information, with no other structure other than cross-references (links) imposed on the data. Hypermedia is the obvious extension to this concept (pictures, sound, etc.).

- Computer Aided Instruction (CAI) experts use Hypermedia as a medium for presentation of courseware. However, the structure of lessons is often "pseudo-linear", i.e. linear or strictly hierarchical, in order to not confuse the users (students) and guarantee that they find their way through the material, see all the material, etc.

- Computer Graphics experts tired of defining 2D and 3D graphics standards and producing realistic images have found a new toy: Multimedia. For graphics people, Multimedia means basically digital video and/or real-time animation ("virtual reality"). To them, Hypermedia is basically Multimedia plus links.

- User interface designers may argue that Hypermedia is just a user interface technique. At last, there is a concept for software that lets them forget about keyboards and makes use of new user interface devices like touch screens, position sensors, eye trackers etc., and heavily relies on graphical user interfaces.

- Cognitive Psychology argues that the concept is more suited to the associative nature of human thinking than the conventional referential access to electronic data (e.g. having to know in what directory and file the information is in order to find it).

- Hypertext and Hypermedia may also be seen as a concept of electronic publishing. When compared to a conventional book, electronic Hypertext and Hypermedia versions offer much greater functionality. Electronic publishers would probably say "Hypermedia is what you will find on a typical CD ROM".

- Hypermedia should enhance communication and collaboration of users. In fact, early systems emphasized this aspect more than today's modern implementations. "Groupware", "Computer Supported Collaborative Work", "Electronic Classroom", and "Electronic Meeting Systems" are new challenges for Hypermedia systems.

- The casual computer users who are no expert in any computer science field will probably already have seen one incarnation of Hypertext: The context-sensitive on-line help that is built into most modern personal computer software. So they also (think to) know what is Hypertext.

Now what is Hypermedia, again?

My opinion is that Hypermedia is all of the above, or at least has the potential to be. This thesis describes the design and some applications of a new Hypermedia system called Hyper-G. The system is developed at the Technical University of Graz, and is designed in such a way that it serves as a solid foundation on which to build diverse Hypermedia applications.

1.2 Previous Work

This work relies on experience gathered at the Institutes of Information Processing of the Graz University of Technology under the leadership of Professor Hermann Maurer over the last decade or so, as briefly described below.

First, there is rich experience with Bildschirmtext – the Austrian videotex system, that can be considered a simple Hypertext system. Information is structured in "pages" that are organized hierarchically; however, arbitrary links between pages can be defined. The pages can also contain (simple) graphics, so it can almost be considered a Hypermedia system.

However, the user interface is cumbersome: Links are activated by pressing the keys "0" to "9", there is no full-text search, transmission is slow, and so on. Users report the typical problems with Hypertext: It takes them too long to find information, they have difficulty finding it again in a later session, get lost in the information, etc.

Encyclopedias have been prepared in Hypertext form and are now available to the general public. On-line communication of videotex users and off-line bulletin boards, a sort of E-mail, and access to other databases (e.g. Homebanking, phone directory) are the communication facilities already implemented in the Austrian videotex system. It is a large-scale, multi-user information system and many the problems of such systems apply and have been investigated.

Second, experience with Computer Aided Instruction (CAI) – more than 700 lessons on various topics are available – has shown us some of the user interface related problems of presentation-type CAI systems. Also, the system employed can be regarded as simple Hypertext (Hypermedia) system. There is some experience with question-answer dialogs not found in similar systems, and this will be incorporated into Hyper-G.

Third, my personal experience is in the field of computer graphics, graphical user interfaces, and computer animation. This is why chapters 3 and 4 are a bit graphics- and user-interface-biased.

1.3 Overview of Thesis

Many of the ideas that can now be found in the design of Hyper-G are condensed from discussions with Hermann Maurer, Ivan Tomek, Robert Stubenrauch, Helmut Mülner, and others. There have been previous attempts to come up with a clean design of the overall system, based on (sometimes vague) ideas of what might be possible applications of Hyper-G, but none was entirely successful in separating requirement specification and implementation details.

I tried to grasp those ideas, structure them into a set of requirements, and to strictly separate the specification of requirements from additional design decisions and the description of the implementation model. This can be found in chapter 2, which certainly constitutes the main part of the thesis. In its clear separation of the various aspects and its elegant design using a relatively small number of orthogonal concepts this can be considered an important and novel contribution to Hypermedia research.

Chapter 3 introduces a new concept for the creation of real-time, interactive animation. It is essentially a combination of Computer Animation and Hypermedia technologies, therefore I call it *Hyper-Animation.* A key point is that it can be naturally integrated within the Hyper-G environment.

This concept is also the basis of some of the more advanced applications of Hyper-G described in chapter 4. Typical applications will be information systems, electronic publishing, and exhibits that may be found in a museum or exhibition environment. Chapter 4 also closes the circle of this thesis, as the applications described there led to the original requirements contained in chapter 2.

Finally, a word on the structure of the text. In order to make this introductionary chapter as readable as possible, I have deliberately avoided cross-references, references to figures and tables, citations, and footnotes.

However, technical texts like this thesis are difficult to write as linear text. In the forthcoming chapters, you will see a lot of constructs like “ *(section 4.3.1.8, page 136)*”, “*see figure 2.1 on page 17*”, “[1]”, and “*[136]*”, which are typical elements of Hypertext (“links”). Because the text contains also some figures, you might regard it as a (poor) implementation of Hypermedia.

As a consequence, the text is difficult to read on paper. I have included page numbers in the “long jump” references to reduce “access time” when following the link. On some occasions, I have avoided cross-references and duplicated information (e.g. where it is only one sentence). This introduces a bit of redundancy, but eases reading.

The text was written using LaTeX[90, 91], which lets you specify the structure of the text instead of the appearance. It is easy to decompose the text into its individual sections, subsections, and paragraphs and then create overview documents that contain just links to those pieces of text [57]. References are not explicit but use symbolic names of the target (e.g. Referencing figure 2.1 is done by specifying `\ref{figCoreSystem}`; at the figure there is a `\label{figCoreSystem}`). Such

[1] This is a footnote (or annotation in Hypertext terminology).

references can easily be converted to Hypertext links to the actual figure. The same is true for cross-references in the text, footnotes and citations (if the cited literature is also available as Hypertext).

As a conclusion, I hope that soon documents like this one can be automatically converted to Hyper-G documents and links, and be easier to read and use as a reference.

Chapter 2

Hypermedia

This chapter constitutes the main part of this thesis. It contains both the specification of requirements and an overall system design of Hyper-G. Great care has been taken to strictly separate the specification of requirements from additional design decisions and the description of the implementation model.

While most of today's Hypermedia systems are designed for a specific purpose, Hyper-G is designed as a general-purpose Hypermedia system that can be used for a number of applications. Only relatively few 'orthogonal' concepts are combined in a clear design to achieve different features found in different special-purpose Hypermedia systems.

E.g., the concept of '*Associative Link Following*' of static links resembles the idea of bi-directional links found in the Intermedia system. '*Dynamic Links*' integrate database queries within the Hypermedia metaphor. '*Associative Link Following*' and '*Dynamic Links*' combined yield the '*full text search*' facility found in some Hypermedia systems related to electronic publishing, without the need for another independent concept. The '*Remote Database*' concept combined with '*Dynamic Links*' allows to impose Hypermedia structure on non-Hypermedia, remote databases, which is a step towards the global Hypermedia and electronic publishing system envisioned by Ted Nelson [136]. In contrast, almost all current Hypermedia systems are limited to working with local data [139] that is under control of the Hypermedia system.

The fairly clear separation between the user interface and the underlying Hypermedia engine which I have achieved allows to run Hyper-G with a number of user interface metaphors, which is desirable for a general-purpose Hypermedia system that is to be used in diverse environments by different kinds of users. To my knowledge, all current Hypermedia systems stick to a specific user interface metaphor, and can therefore be considered more or less special-purpose systems.

I specifically designed Hyper-G as a large-scale Hypermedia system. High-level navigation tools and search strategies, a multi-lingual Hypermedia concept, new user identification notions and communication facilities ease the use of Hyper-G in an international, multi-user environment. My implementation model allows to implement Hyper-G in a network of standard UNIX workstations, and is specifically designed for extensibility of the system (e.g., new document types, new user interface metaphors).

2.1 Overview

Generally, the term *Hypermedia* refers to a combination of two other terms that have been used quite loosely in the past 20 years for many different collections of features: *Hypertext* and *Multimedia*.

The basic idea of Hypertext is not exactly a new one: The term, coined by Ted Nelson in the 1960's [135], describes a vast network of text fragments linked together, an electronic writing and reading system that uses the power of the computer for more than editing and display. However, the idea can be traced back to the 1940's. Vannevar Bush, President Roosevelt's director of the Office of Scientific Research of Development, is usually credited with first describing a Hypertext system in his 1945 article "As we may think" [16][1]. He envisaged a machine for browsing and making notes, and linking information by "associative trails", modeled after the associative way of human thinking. In general, Hypertext is not so much a new idea as an evolving concept of the possible applications of the computer.

In europe, Sam Fedida – the "father" of videotex – had similar visions concerning large-scale, publicly available information systems [46].

Implementations of such large-scale information systems have been pioneered at institutions like Brown University (*Intermedia* [129, 189]), the University of Maryland (*Hyperties* [163]), Xerox PARC (*NoteCards* [60]) and by Ted Nelson (*Xanadu* [136]). More complete surveys can be found in [30], [52], [139] and [176].

Multimedia adds new facets to Hypertext. Although the term Multimedia is used by a number of persons (and companies) with a number of different meanings, the following definition seems to be the greatest common denominator:

> A Multimedia system supports at least three out of this list of data types:
> Text, still image, video, and audio information.

[1]You probably won't be able to get the original. However, the article was reprinted in [15] and [136].

Hypermedia can be seen as the Hypertext paradigm plus Multimedia features. In the last years, significant research has been undertaken in this field. E.g., [123] lists some 1400 articles and books on related topics. Hypermedia research integrates methods and ideas from information retrieval, user interface design, computer graphics, and cognitive science. Although significant progress has been made, a number of open problems still remain.

Today's Hypermedia systems can roughly be divided into two groups: The large, multi-user information retrieval systems (like the ones mentioned above), that rely on special hardware and are not commercially available. And there are small but popular single-user "Hypermedia" systems (e.g. *HyperCard* [185]), that lack a number of essential features, but are commercially available. As pointed out in [116], these systems can not truly be regarded as Hypermedia systems, due to the lack of communication features and the limitations on database size. Also, the user interface metaphor of such systems is more targeted towards browsing knowledge in a way predetermined by an author than to finding facts [112].

2.2 The Design of Hyper-G

Hyper-G is the name of an ambitious Hypermedia project currently being carried out as a joint effort by a number of institutes of the IIG (Institutes for Information-Processing Graz) of the Technical University of Graz[2] and the Austrian Computer Society. Part of this research is carried out in cooperation with other research institutions, including Joanneum Research (Graz), the Memorial University of Newfoundland (St. John's), Acadia University of Nova Scotia (Wolfville), and the University of Colorado (Boulder).

Implementation of the full system is scheduled for 1995, but a prototype will be implemented by end of 1992[3]. A very simple prototype called *Hypermedia Base System (HBS)* has been implemented on both PC and UNIX platforms.

Also, data acquisition is performed parallel to the design and implementation stages, in order to fill the system with real-world data from the beginning on, as one of the aims of the Hyper-G development is to test user interface metaphors and search strategies in very large Hypermedia systems.

[2] Yes, the 'G' in Hyper-G stands for Graz.

[3] Funding of this project by the Austrian Ministry of Science, grant 613.531/5-26/90, is gratefully acknowledged.

2.2.1 Design Strategy

The forthcoming sections describe the Hyper-G system as well as the prototype system in some detail. But they do not contain just a descriptive specification, but also the requirements and design decisions that led to certain system features. A great effort has been undertaken to not mix up requirement and specification.

The rest of this chapter is structured as follows:

1. Section 2.2.2 lists 'axiomatic' system requirements. These requirements were gathered by comparing other Hypermedia systems and analysis (or requirement engineering) of their strengths and shortcomings [116]. From the beginning on, requirements are classified in five groups and hierarchically structured.
2. This set of requirements is then refined in section 2.2.4 on the basis of a number of additional design decisions made in section 2.2.3.
3. This large set of requirements logically leads to a specification of Hyper-G. The specification itself is split into:
   (a) Functional descriptions of the processes involved and their intercommunication (section 2.2.5).
   (b) Definitions of the data items the system will handle and the databases needed to find them (section 2.2.6).
   (c) The user's view of the system (section 2.2.7).
   (d) The author's view (section 2.2.8).
   (e) The system administration's view (section 2.2.9).
4. Eventually, a specification of a Hyper-G prototype, based on priorities given to the requirements, may be found in section 2.3.

However, it is not necessary to read this document in the order presented. For example, if you are interested in the features of the prototype, but not in the requirements and design decisions that led to them, you may skip the corresponding sections.

Of course, this strict division of requirements, design decisions, and descriptive specification introduces some redundancy to the text. E.g., you may read about access rights in sections 2.2.2.4, 2.2.3.16, 2.2.4.3, 2.2.6.1, and more.

Part of the information is duplicated (e.g., where it is just one sentence), and a large number of cross-references is contained within the text. Unfortunately, paper does not offer the same functionality as Hypertext (just click on the references), so page numbers have been added to the cross-references where it seemed reasonable.

Still the document remains difficult to read. One personal experience I got from writing this text is that ordinary text just isn't good enough to write technical specifications. With all these cross-references, figures, and tables it would have been better to write is as a Hypertext from the beginning on. If I have not convinced you by now, I hope that the forthcoming specification of a fine Hypermedia system will do so. Let's start!

2.2.1.1 Definition of Terms

Before we begin, however, table 2.1 defines some terms that will be commonly used throughout this chapter, so we can rely on understood basic definitions. Other definitions are supplied where needed.

System	Will almost always mean the Hyper-G system. Includes all the hardware and software of the run-time system (i.e. the system the user sees), plus additional hard- and software for authors and system administrations (tools).
User	User of the system is the person that utilizes the run-time system, e.g. to retrieve information.
Author	Author is the person that creates information that is to be included and may be accessed by the run-time system. It is the 'electronic' author, the 'owner' of an object (e.g. in the case of a picture of the 'Mona Lisa', the author would be the person that inserted that picture into the System, rather than Leonardo da Vinci).
Editing	Will mean preparation of information by the author. Information may include textual, pictorial, sound, and other data.
Terminal	The hardware the user is working on directly.
Session	The period of time the user works with the system on a terminal. Usually, the users have to log in (identify themselves to the system) when starting a session, and log out to stop a session.
Document	An item of information that the user may interact with (e.g. look at it). Throughout the system, information is structured in documents.
User State	Conceptually, the system sees the users' work as changes of their user states. The user state records such information on what documents the user is looking at, etc.
Database	A storage mechanism for documents and other data items the system needs.
Attribute	A piece of information attached to documents. Used to find documents with certain attributes in the database.
Collection	A set of documents that have the same set of attributes attached to them. In some way, these documents belong together.
Link	A connection between two documents. Links play an important role in Hypermedia systems.
Annotation	Sometimes, a user wants to (temporarily) become an author and add information (private or public) to the system, attached to other documents.
Tour	A concept that allows to visit documents in a context supplied by a tour author.

Table 2.1: Definition of Terms

2.2.2 Basic Requirements

This section contains a list of 'first-generation' requirements for Hyper-G. They are structured in five groups: *General, User, Author, System* and *Implementation*. For ease of reference, a unique label is assigned to each requirement. This set of basic requirements will be refined in section 2.2.4 on the basis of additional design decisions made in section 2.2.3.

2.2.2.1 General Requirements

Req. G.1:	The system should allow access to all kinds of information one can think of (*Multimedia Axiom*).

Let us refine this to a set of more specific requirements:

Req. G.1.1:	Information is organized in *documents* of a specific *document type*. Document types include (but are not limited to):
Req. G.1.1.1:	Textual Information
Req. G.1.1.2:	Technical Drawings
Req. G.1.1.3:	Raster Images
Req. G.1.1.4:	Digitized Sound
Req. G.1.1.5:	Digital Movies
Req. G.1.1.6:	Animation
Req. G.1.1.7:	Maps
Req. G.1.1.8:	Electronic Mail, Bulletin Boards, Computer Conferencing
Req. G.1.1.9:	Dialogs (e.g. Menus, Forms)

A system able to handle above document types is generally referred to as a *multimedia system*. To put it simple, above requirements state that Hyper-G is such a multimedia system. Above list is by no means complete; e.g., document types like synthesized voice and music may be added in the future.

Req. G.2:	Documents (**Req. G.1.1**) may be cross-referenced and accessed by *links* (*Hypertext Axiom*).

This means that Hyper-G adopts the well-known Hypertext paradigm. Requirements **Req. G.1** together with **Req. G.2** state that Hyper-G is what is known as a *Hypermedia system*. More specific versions of **Req. G.2** will be given in section 2.2.4.1 (page 35).

Req. G.3:	Hyper-G provides access to very large amounts of data.

As stated in [35], Hyper-G is targeted to mega-quantities of documents and terra-quantities of bytes. This leads to interesting problems not only for system design, but – more important – for the users and the authors of such a large system [86]. An immediate consequence is

Req. G.3.1:	Documents and links are maintained automatically.

Link maintenance involves such problems as deleting all links from or to a document that is to be deleted, in order to avoid dangling references and to ensure database integrity. Also, the system should assist the system administration in protecting the database against obsolete or outdated documents and/or links.

Req. G.3.2:	The mayority of the links is generated automatically.

Especially for text documents, a number of links can be generated automatically either when inserting the document into the database, or at runtime. More on this in section 2.2.3.4 (page 22).

A user related consequence (**Req. U.3**) of **Req. G.3** is given in section 2.2.2.2 (page 13).

An author related consequence (**Req. A.2**) of **Req. G.3** is found in section 2.2.2.3 (page 15).

A system related consequence (**Req. S.3**) of **Req. G.3** is found in section 2.2.2.4 (page 15).

Req. G.4:	The system should be easily extensible.

This is especially true for the set of supported document types and user interfaces (**Req. U.1.1**) and has some effect on the implementation model (sections 2.2.5, 2.2.6).

2.2.2.2 User-related Requirements

Req. U.1:	The system should be usable by untrained users as well as expert users.

In general, Hyper-G is likely to see a number of scenarios [87, 35, 121, 126]. One is a university information system (section 4.1) based on Hyper-G where users are expected to communicate with computers on a regular basis, another one is a so-called *viewseum* (section 4.3) available to the general public. Hyper-G should adapt to the needs of all user groups. Implications of this requirement are discussed in section 2.2.4.2 (page 39).

Req. U.2:	The system should offer a choice of user languages.

This not only requires the user interface to work with a number of languages, but also means support of multi-lingual documents (e.g. translations of the same document provided by the author, or simple computer translations).

Req. U.3:	The user is supported by high-level navigation tools.

Above requirement may also be seen as a consequence of **Req. G.3** (large number of documents). High-level navigation tools include use of user interface metaphors, database queries, collections, tours and scripts. These concepts are discussed in detail in sections 2.2.3 (page 17 and 2.2.4.2 (page 39).

Req. U.4:	The system is configurable on a per-user basis. It remembers a user's preferences between sessions.

This means that users may configure the system to adapt to their special needs (e.g., language preferences (see **Req. U.2**).). Moreover, this configuration is retained by the system until the next session. Of course, this implies that the system is able to distinguish between users by some login procedure and keeps configuration files. It also suggests use of user groups, access rights and so on. However, these are not basic requirements and discussion is therefore postponed until section 2.2.4.2.

Req. U.5:	At any time, the user may return to a previous user state.

This applies to all the previous states of an active session, which yields sort of an *undo* function. As a special case, when entering a new session, the user may return to where the previous session was left. **Req. U.5** requires the system to remember user states, keep user logs etc. Again, the details are related to the actual implementation and are discussed in section 2.2.4.2.

Req. U.6:	The system can be used anonymously.

An information system that is to be used by a large number of users on a regular basis raises a data security problem. If users are known by name, the system might be used to find out who looks for certain information, monitoring users working time, etc. To avoid such problems from the beginning on, users may login anonymously, as in the Austrian videotex system [118]. The modes *semi-identified* and *anonymously identified* (section 2.2.3.15, page 32) still allow to perform most operations of the system.

Req. U.7:	Certain system features and information items are chargeable.

For example, printing a screen dump or saving it to disk may cost some (small) amount of money. In addition, information providers may charge something for the access to certain pieces of information, remote databases, etc. In conjunction with **Req. U.6** this implies that a rather tricky charging mechanism has to be implemented.

2.2.2.3 Author-related Requirements

Req. A.1:	Any user is a potential author.

The database may be manipulated by potentially every user, either by manipulating documents or links (*annotation*). Of course, access rights control the user's ability to change existing data items or create new ones (section 2.2.3.16, page 34).

Req. A.2:	The author is supported by high-level information editing tools.

These information editing tools divide into document preparation tools and link editing tools. They are not part of the Hyper-G runtime system, so the details are postponed to section 2.2.8.

2.2.2.4 System Requirements

This section describes basic requirements related to system administration.

Req. S.1:	A number of *terminals* is attached to the system. On any terminal a Hyper-G *session* may be started. The capabilities of the terminals may differ.

That is, the terminals are not all the same. There may be terminals with special user interfaces, communication features, hard/softcopy capabilities etc. As a consequence, Hyper-G must not rely on the presence of specific terminal facilities.

Req. S.2:	Access rights control whether a given user on a given terminal may perform a certain function on an object.

After describing some design decisions related to access rights (section 2.2.3.16), this concept is elaborated in more detail in section 2.2.4.3.

Req. S.3:	The system should avoid uncontrolled growth of unimportant or outdated information.

The system administration should be supported in finding such information and deleting it at regular intervals.

2.2.2.5 Implementation Requirements

Req. I.1:	Access to information should be reasonably fast.

This requirement is somehow contradictory to **Req. G.3** (large amount of data). However, it is well-known that a response time larger than about 250 milliseconds confuses users [19]. A lower response-time variance seems to be more important than absolute response times [23]. Also, users are willing to accept, e.g., to wait a few seconds for a database query, if the system supplies them with immediate feedback of acceptance of the query. We have seen systems (e.g. videotex [46, 114, 115]) whose acceptance by the public has been dampened by less than optimal response time.

Above requirement has significant impact on the design of the databases (speed vs. space tradeoffs) and the computers used. More on the database design may be found in sections 2.2.4.4 (page 50) and 2.2.6.2.1 (page 70).

2.2.3 Design Decisions

2.2.3.1 Distributed System

Requirements **Req. G.3** (large amount of data) and **Req. S.1** (number of terminals) suggest to implement Hyper-G as a distributed system running concurrently on a heterogeneous network of computers. To achieve high performance, Hyper-G is split into concurrent processes (section 2.2.5, page 52) that may execute on a number of machines simultaneously.

Also, the data is distributed among the machines (section 2.2.6, page 60). Data includes documents of various types, links, log files etc. The use of object-oriented design eases the task of distribution of data and processes.

2.2.3.1.1 Core System

The *Core System* of Hyper-G will consist of a high-speed local area network (LAN) of Hyper-G Servers and Workstations (see figure 2.1).

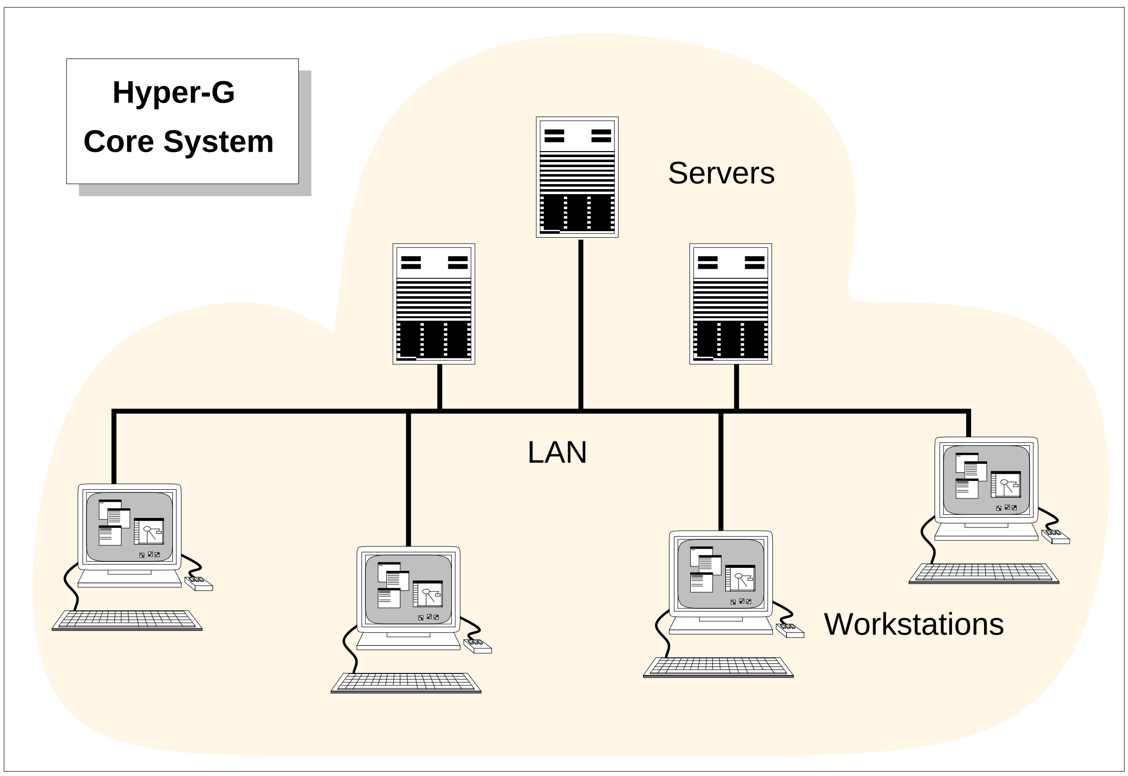

Figure 2.1: The Hyper-G Core System

Hyper-G Servers will typically be supplied with gigabytes of fast mass storage (harddisks), but there may be special-purpose Hyper-G servers equipped with, e.g, CD-ROMs, Video or Audio Discs or Tapes.

Workstations will not be all the same. Most of them will be standard graphic workstations, but some of them may have additional audio/video capabilities or special user interfaces.

2.2.3.1.2 Front-Ends

In addition, Hyper-G will be accessible from remote terminals over a wide area network (*WAN*, see figure 2.2). In most cases, this WAN will be slow (typically dial-up phone lines or telematic services networks; up to 64 kbits/s), and the facilities of the terminals will vary, so the access to certain types of documents will be limited.

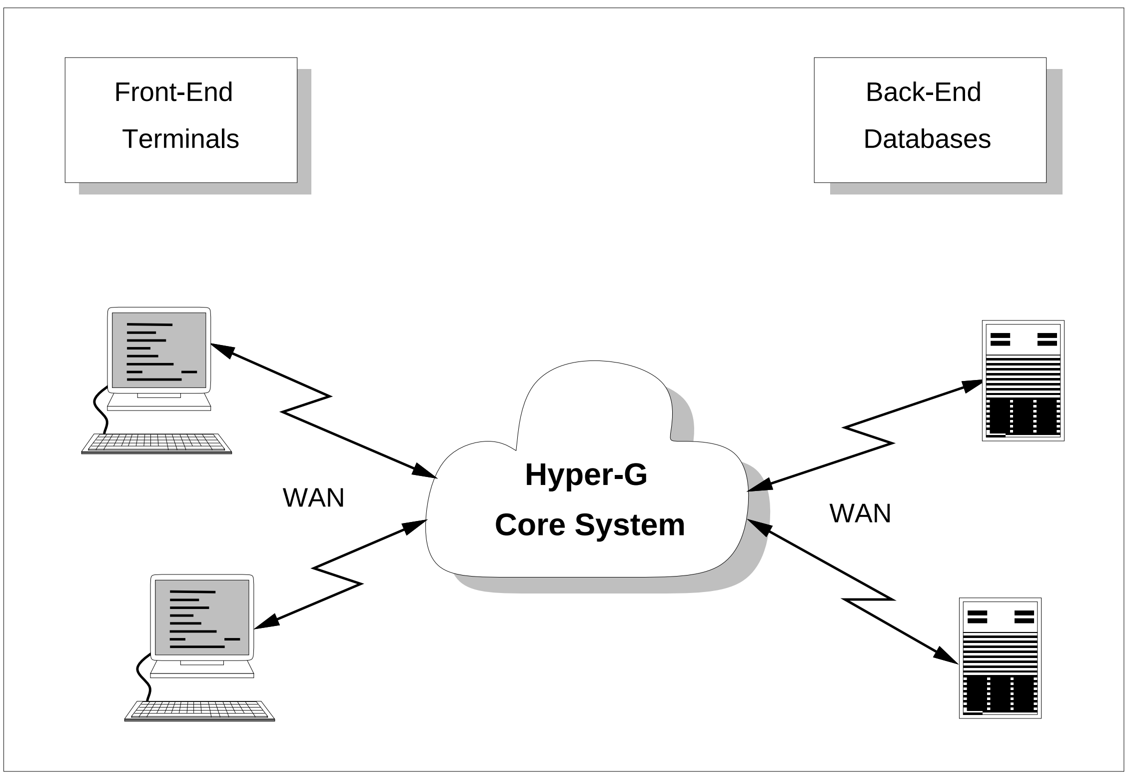

Figure 2.2: Hyper-G and the World

2.2.3.1.3 Back-Ends

Also, it will not be feasible to keep all the data in the core system. Therefore, Hyper-G will support access to remote databases (e.g. official airline guide, phone