DIVISION OF COMPUTER SCIENCE
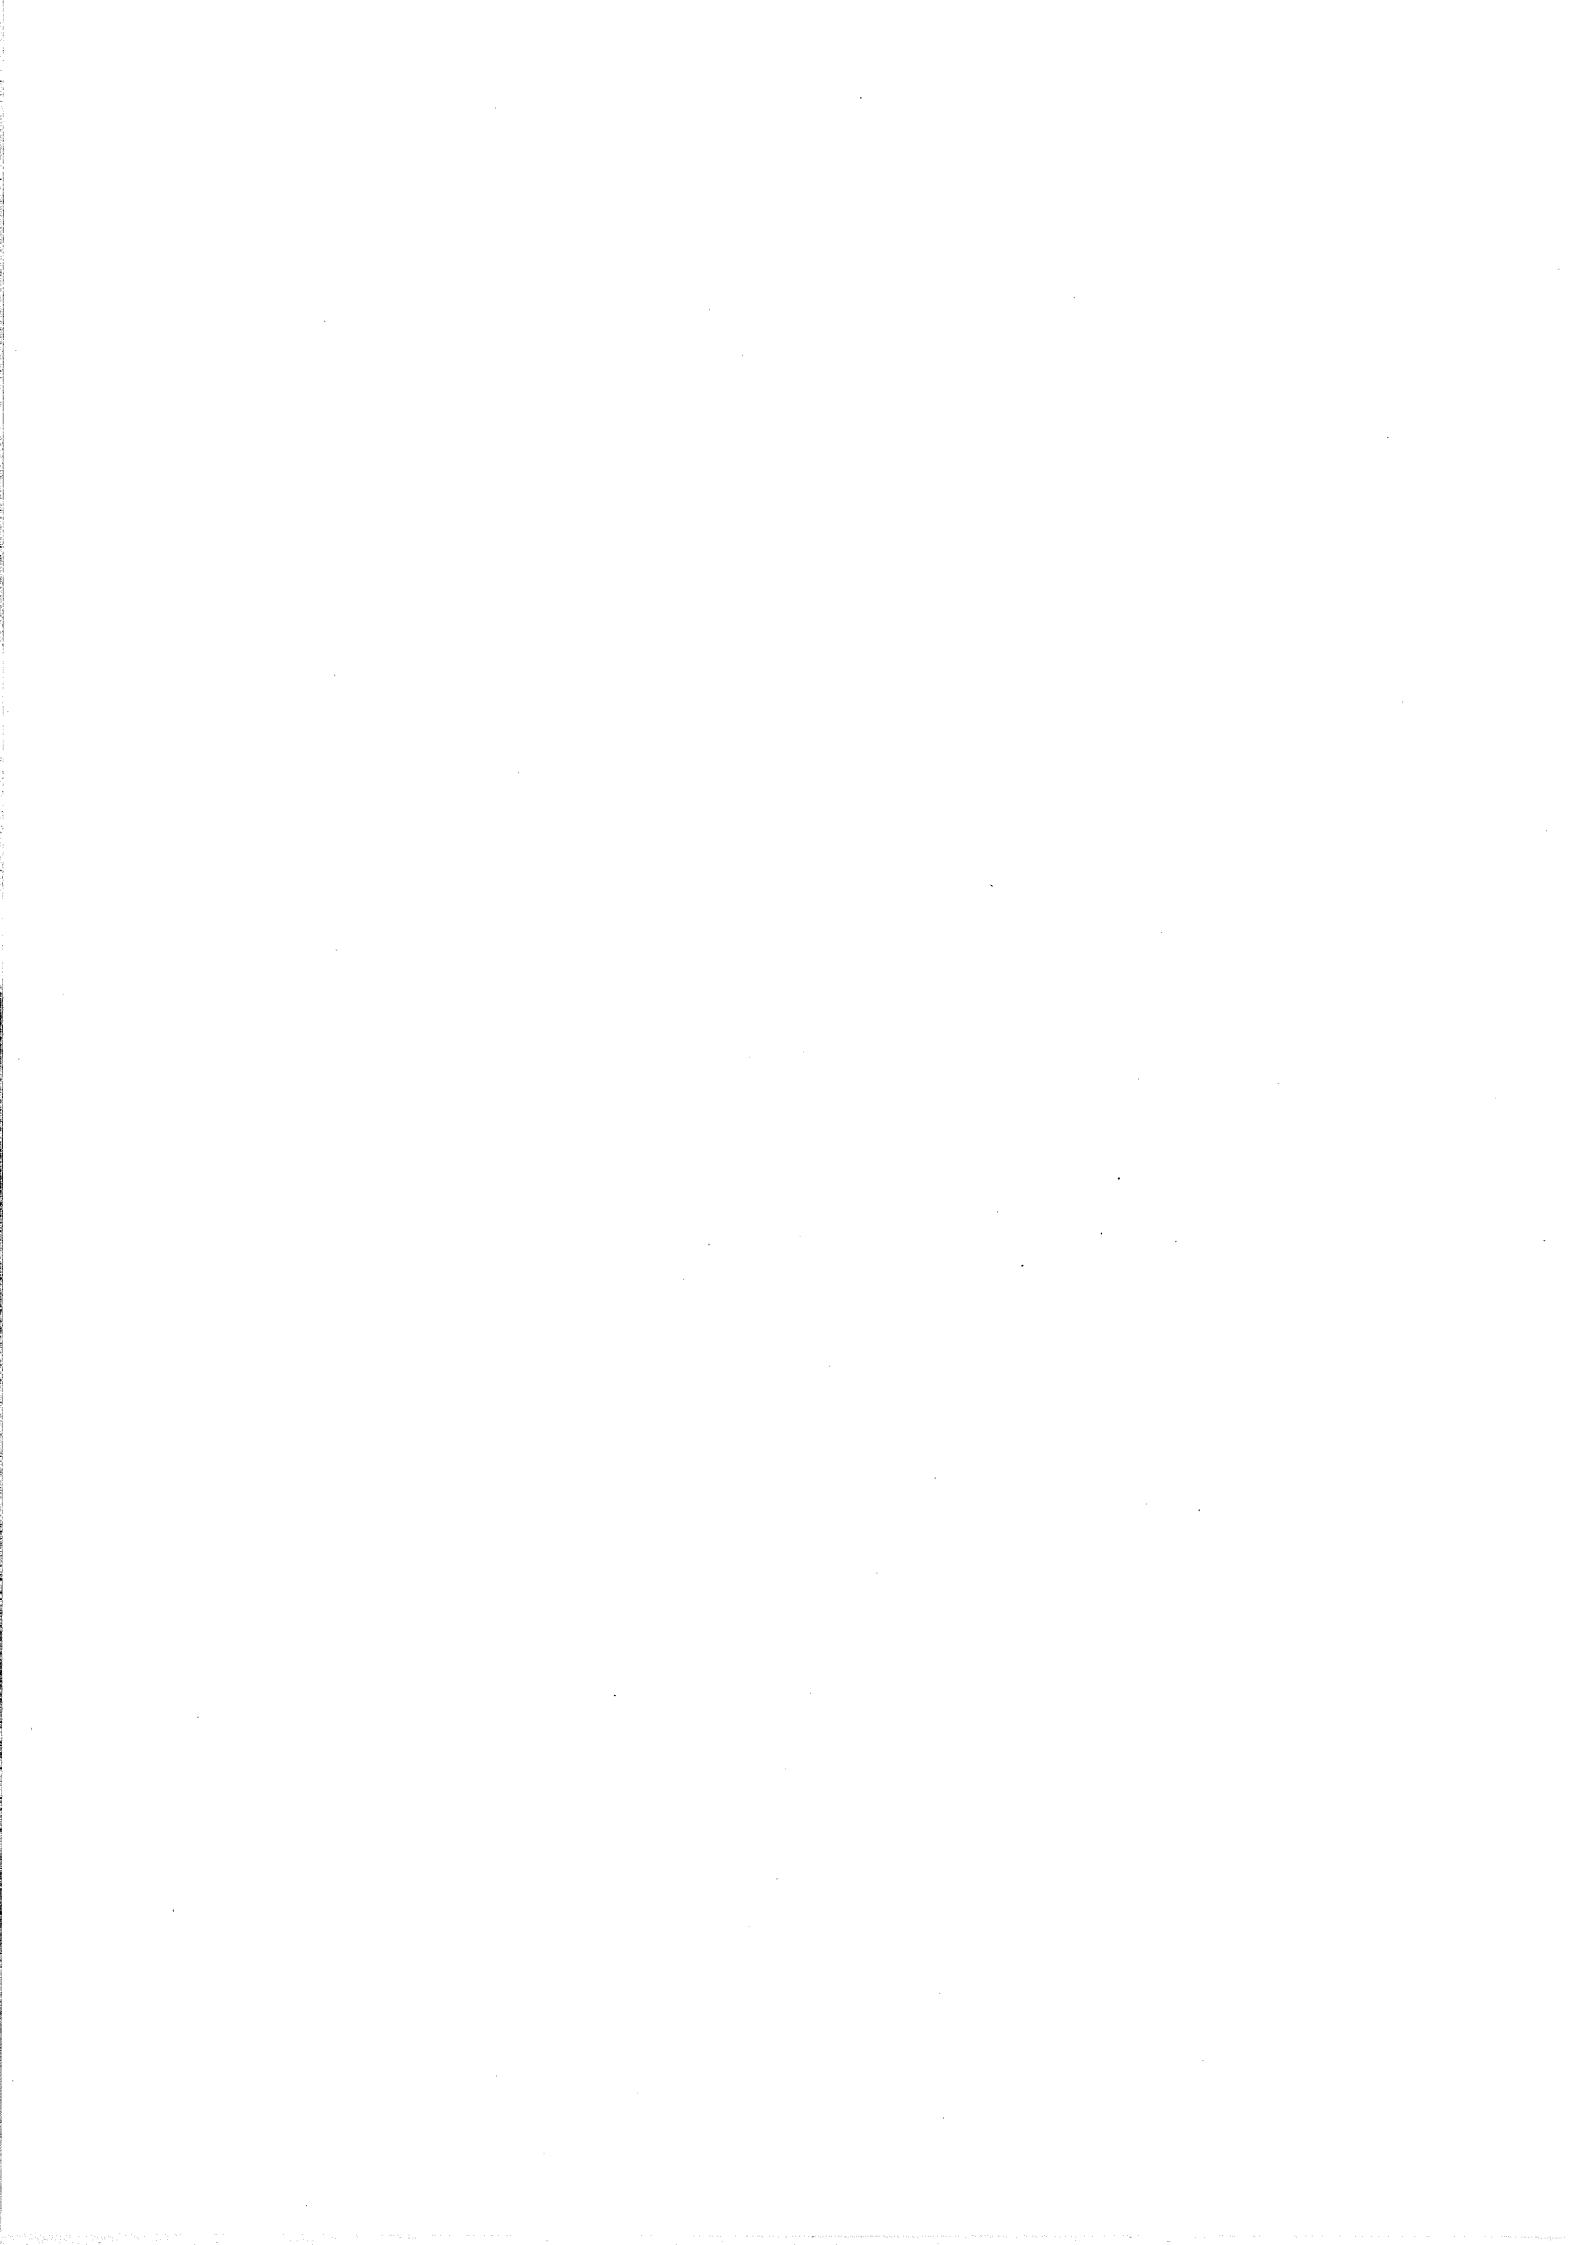
Providing Public Access to Information in Complex and Weakly
Structured Domains: A 3-Layered Model for Hypermedia
Information Systems

M Bearne
S Jones
J Hewitt
J Sapsford Francis

October 1995

Providing Public Access to Information in Complex and Weakly Structured Domains: A 3-Layered Model for Hypermedia Information Systems

Mike Bearne, Sara Jones, Jill Hewitt, John Sapsford-Francis
School of Information Sciences, University of Hertfordshire,
Hatfield, AL10 9AB, UK

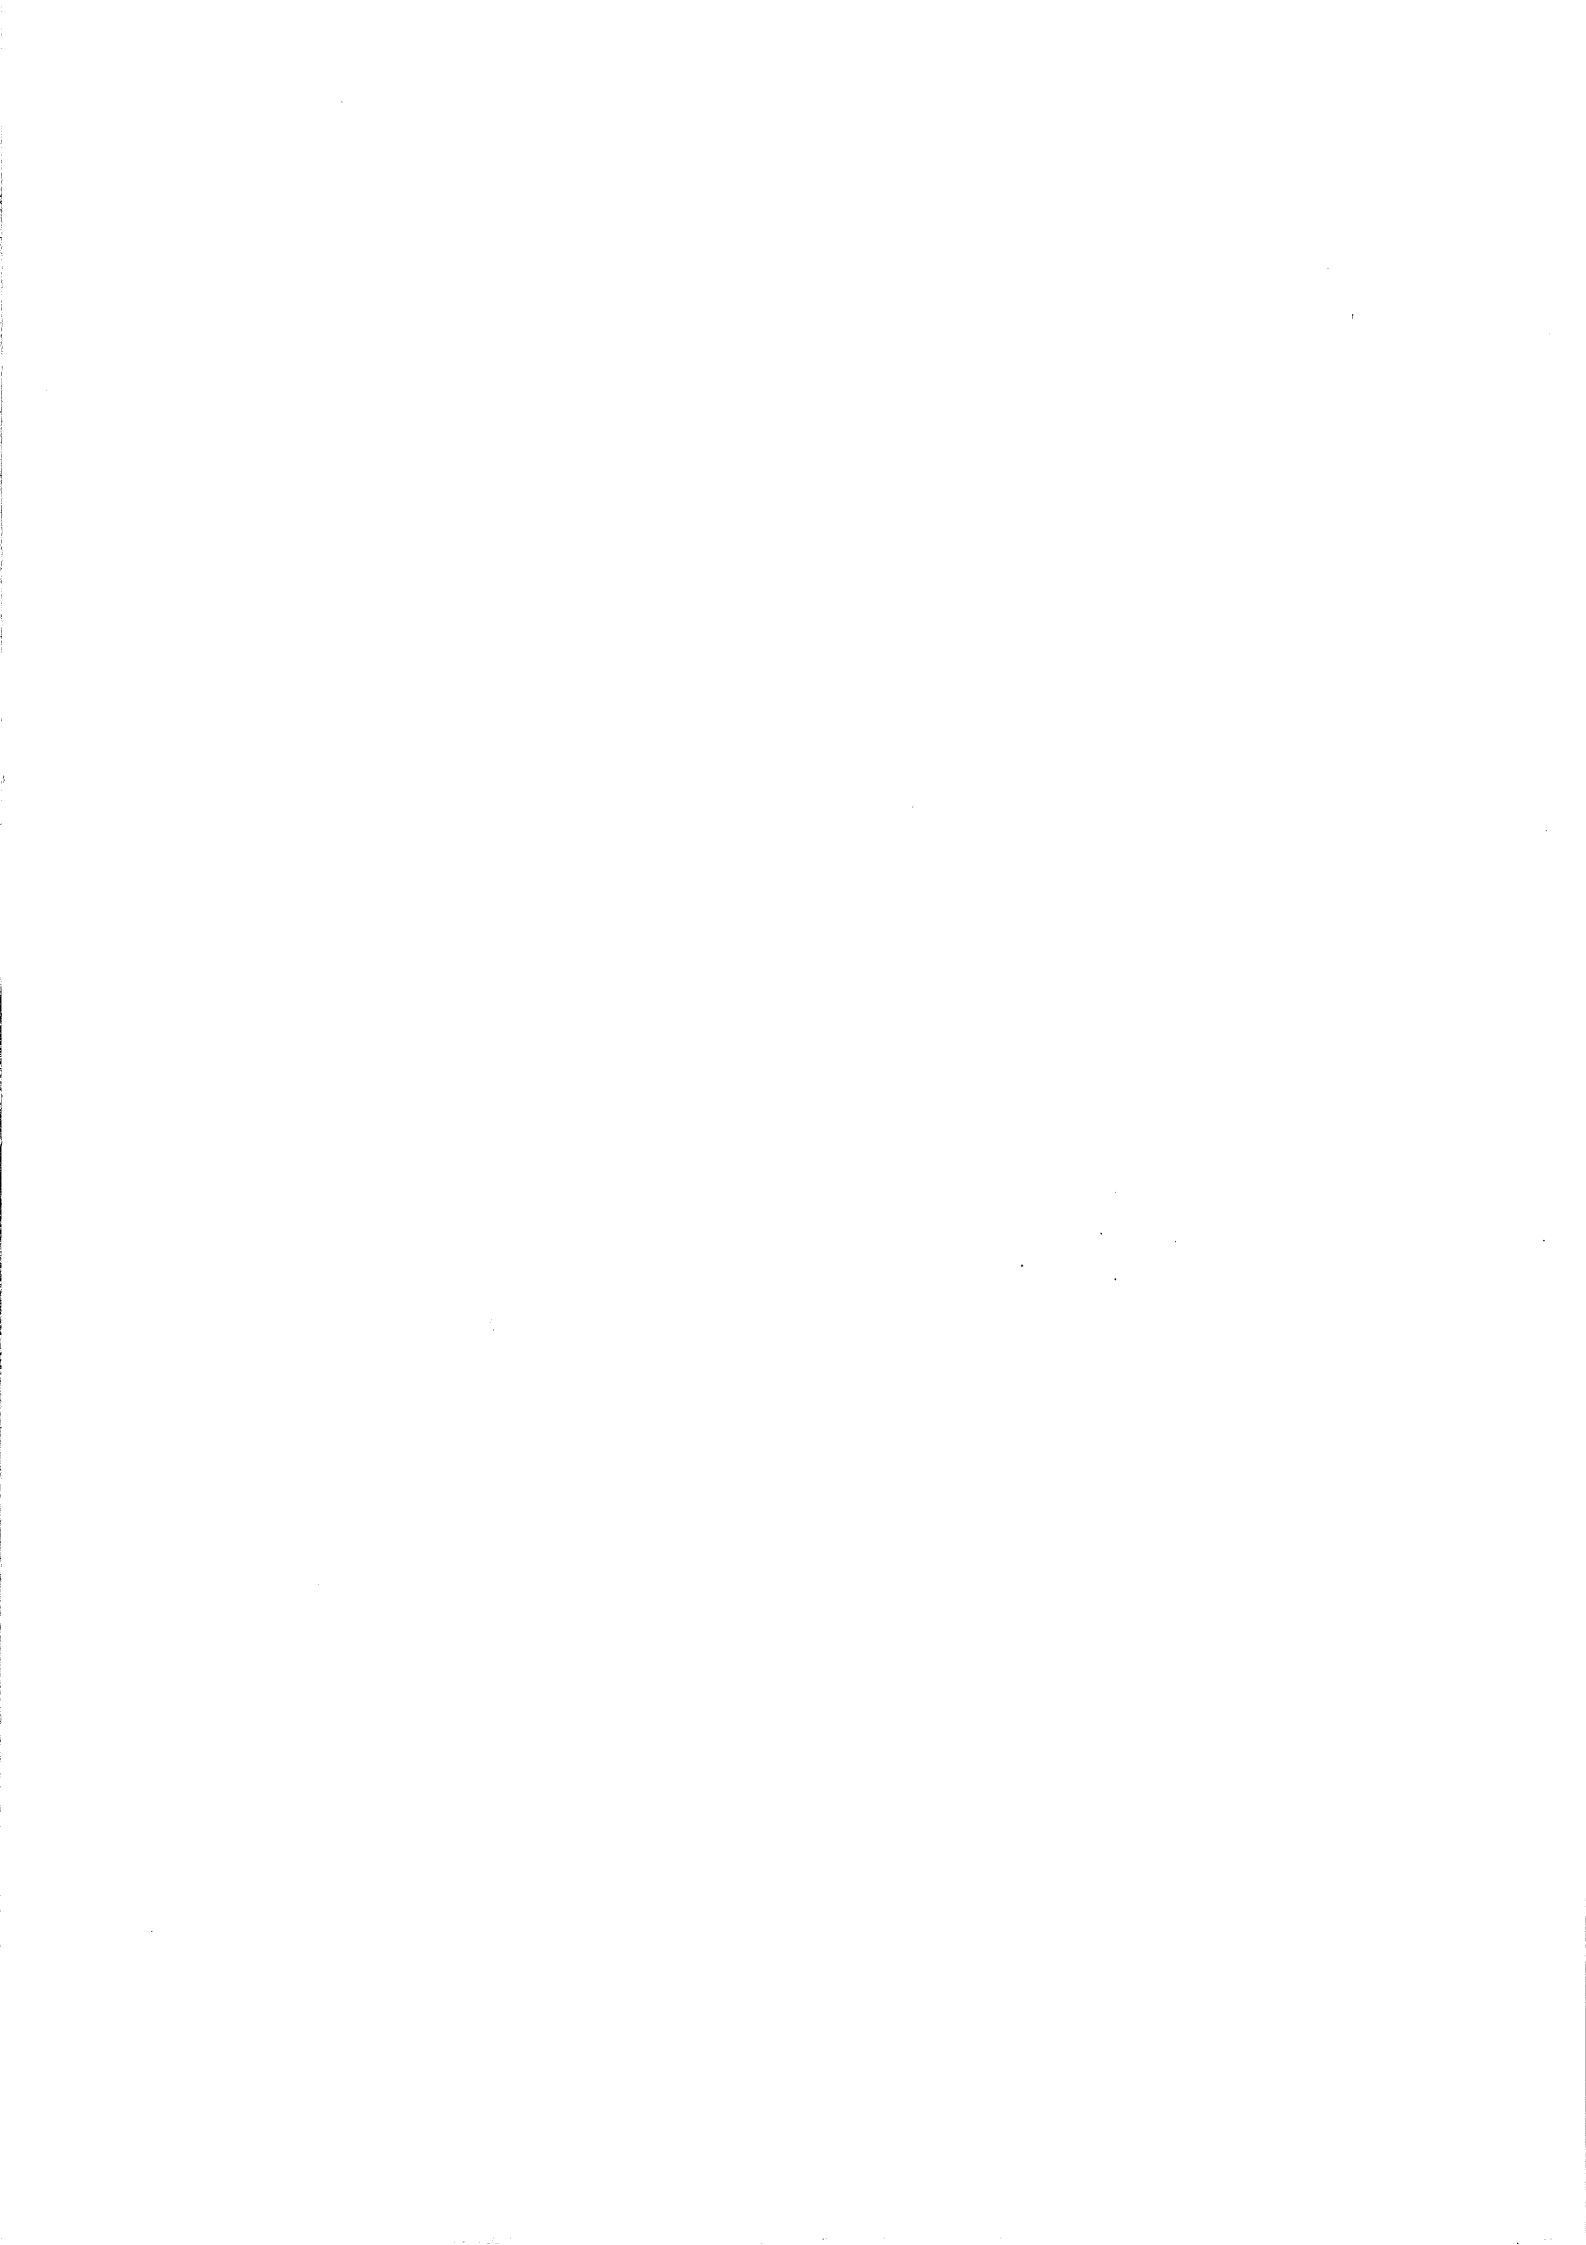Tel: (01707) 284766, 284370, 284327, 284309
Fax: (01707) 284303

comrmb,comrsj,comqjah, comqjs@herts.ac.uk

Abstract

Advice systems are defined as information systems which present users with both information and more subjective expert advice about complex and weakly structured domains. This paper presents a Generic Advice System Architecture (GASA) to assist in developing such systems. It describes how the architecture was used in developing the SPIRE system, whose aim is to provide advice and information which will assist the integration of students with disabilities into higher education. The paper discusses the way in which the GASA addresses key issues in the development of hypermedia advice systems including the need to make such systems available to users with little training and limited knowledge of the domain; the requirement to support diverse information exploration strategies; the provision of purely 'point and select' access; and the minimisation of user disorientation and cognitive overhead.

Providing Public Access to Information in Complex and Weakly Structured Domains: A 3-Layered Model for Hypermedia Information Systems

Mike Bearne, Sara Jones, Jill Hewitt, John Sapsford-Francis
School of Information Sciences, University of Hertfordshire,
Hatfield, AL10 9AB, UK
comrmb,comrsj,comqjah, comqjs@herts.ac.uk

Abstract

Advice systems are defined as information systems which present users with both information and more subjective expert advice about complex and weakly structured domains. This paper presents a Generic Advice System Architecture (GASA) to assist in developing such systems. It describes how the architecture was used in developing the SPIRE system, whose aim is to provide advice and information which will assist the integration of students with disabilities into higher education. The paper discusses the way in which the GASA addresses key issues in the development of hypermedia advice systems including the need to make such systems available to users with little training and limited knowledge of the domain; the requirement to support diverse information exploration strategies; the provision of purely 'point and select' access; and the minimisation of user disorientation and cognitive overhead.

1 Introduction

Computer-based systems are increasingly being used to present general information to people in public places such as shops, museums, parks and libraries. Such systems increase the availability and accessibility of information and complement existing information services by, for example, providing access to more people in more locations and for longer hours. These systems, often termed public access, point of information or 'kiosk' systems (1,2) frequently present information in the form of hypertext.

This paper describes a generic architecture for use in developing a special class of kiosk systems, termed 'advice systems', which present users with both information and more complex expert advice about a particular area. Advice systems are intended to be used in complex and weakly structured domains, in which each piece of information may be related to many other pieces in a number of different ways, and there is no generally agreed view on how the domain is structured. Advice from a number of domain experts can be used to present various perspectives on the relationships between different pieces of information, and to guide users to the answers they are seeking, even though they may begin with only vague ideas about the kinds of questions they want to ask.

The Generic Advice System Architecture (GASA) presented below consists of 3 layers - a hypertext layer, a more traditional hierarchically organised database layer, and a novel data access layer. The architecture provides a template which can be used to structure advice systems in a variety of different domains in such a way as to allow a wide range of users to access the information they require without first needing to formulate a precise question or query.

The architecture was developed as part of the SPIRE project, whose aim was to provide advice and information to assist the integration of disabled students into higher education. In section 2 of this paper, we outline some of the challenges and problems associated with the design and development of hypermedia kiosk and advice systems in general. Section 3 describes the form these challenges took in the context of the SPIRE project, and section 4 discusses the way in which the SPIRE system was designed and implemented using the 3-layered GASA to address some of these challenges. Sections 5 and 6 evaluate the success of the model as implemented in SPIRE and its relation to other work.

2 Challenges in the Design of Hypermedia Advice Systems

Designing public access hypermedia information and advice systems presents the developer with a number of interesting challenges. Some of these challenges are described below.

2.1 Defining Characteristics of Public Access Information and Advice Systems

Public access point of information, or 'kiosk' systems have, in recent years, increasingly been placed in buildings as diverse as museums, libraries, railway stations, and shops. Some of the defining characteristics of such systems which have arisen out of the literature on human-computer interaction (HCI) were set out by Leventhal et al (1) as follows:

- kiosk systems are intended to be used by wide-ranging audiences, of different ages, and different levels of education, experience and computer literacy;

- they serve primarily to provide reference information, and may be seen as providing a service complementary to existing information services or reference library facilities;

- kiosk systems will often be used on a casual or 'walk-up-and-use' basis, so that users may not have the time or motivation to overcome difficulties in using the system or searching for particular pieces of information;

- because of the situations in which they are used (e.g. park, station platform, shopping mall) they often need to be accessible to users using only pointing devices (mouse or touch screen) for input.

Advice systems, as they are classified here, represent a special class of information kiosk. In addition to the above characteristics, advice systems also have the following features:

- advice systems are used to present information relating to complex and weakly structured domains;

- they present the user with both *advice* - information written by domain experts, which is generally text-based, free-form, and highly subjective and contains a large number of cross-referential links - and *data* - 'hard facts' which are objective and can be relatively well-structured.

Because of the wide ranging user group and the variety of information in the system, it is likely that different users of public access information and advice systems will want to employ different strategies for information exploration depending on their differing information needs. Two methods identified (3,4) as strategies that users may employ when using an information system are:

- *browsing* in which users explore the information in the system with fairly open-ended goals in mind. The users' information needs are broad and non-specific. Browsing has been characterised as going from *where* to *what*: users are aware of where they are in the system and they wish to know what is there (3)

and

- *searching* in which users have specific goals, and are able to formulate queries which will help them satisfy those goals (i.e. to retrieve specific pieces of information). The users' information needs are narrow and highly specific. Searching has been characterised as going from *what* to *where:* users are aware of what they want from the system and need to know where in the database it is (3).

While novice users commonly prefer to browse systems, more experienced or expert users are more likely to employ directed search techniques for information retrieval (4,5). Individual users may also

3

employ a combination of searching and browsing in a serendipitous fashion during the course of a single interaction with a system (6).

2.2 The Use of Hypertext in Advice Systems

It has been suggested (1) that hypertext provides an ideal basis for the design and implementation of public access kiosk systems. Hypertext provides a very flexible structure which naturally promotes browsing and navigational access to information, and the travel metaphor it employs supports a very intuitive form of information retrieval. Furthermore, hypertext is naturally suited to representing the sort of free-form, textual and highly inter-related information that makes up a body of domain related advice.

However, in spite of the obvious advantages of a hypertext model in this context, there are also certain disadvantages, or at least challenges, inherent in the use of hypertext. Hypertext systems, whilst naturally supporting browsing (navigation) as an approach to information exploration, offer less support for users wishing to employ a searching strategy. Halasz has identified support for search and query mechanisms to be one of the seven key issues for the next generation of hypertext systems (7). As described above, advice systems must provide support for both browsing and searching strategies if they are to be effective.

It is also well known that navigating through large hypertext networks can cause disorientation in users (the familiar 'lost in space' problem (8)). This problem is also referred to as "the embedded digression problem" (4). Disorientation is obviously a potential problem in advice systems which present large bodies of weakly structured information to a wide range of often untrained users. Conklin's 'cognitive overhead' problem (8), concerning the cognitive resources required by users in order to select which links in a hypertext network to traverse and which to disregard, is also likely to be significant in this context.

Finally, although hypertext systems are very well suited to modelling free-form textual information, they are perhaps less well suited to modelling other, more structured forms of information. While hypertext provides an ideal model for the representation of the advice held in an advice system, it is less well suited to representing the more structured data.

2.3 Summary of Challenges

Summarising the above, designers of public access advice systems must ensure that such systems:

- are readily usable by users with little or no training and limited knowledge of the relevant information domain;

- are usable by a diverse set of users with diverse problem-solving strategies and information exploration which may involve both searching and browsing;

- permit users to implement such strategies using only pointing devices for input.

They must also address some of the potential drawbacks of using a simple hypertext model in the design and implementation of advice systems including:

- the possibility of disorientation and cognitive overload

- the lack of support for presenting structured data

The rest of this paper will describe the Generic Advice System Architecture (GASA), a novel, 3-layered model which augments a basic hypertext structure with a hierarchical database layer and a filtering and guided-route mechanism, thus addressing each of the challenges identified above. Discussion centres on the particular case of SPIRE, a public access system providing advice and information regarding the integration of disabled students into higher education.

3 The SPIRE Project

The aim of the SPIRE system is to provide advice and information to assist with the integration of students with disabilities into higher education. It has all of the important characteristics of advice systems defined in section 2.1:

- SPIRE is intended to be used by a wide range of people including students, lecturers and University support staff;

- it provides reference information about the domain of students with disabilities in higher education, and supplements existing University support services;

- the system is available to all users of the University library on a casual 'walk-up-and-use' basis, and also to support staff working in the University Health Centre;

- the domain about which information is provided (that of students with disabilities in higher education) is a complex domain in which each piece of information may be related to many others in a number of different ways, and different experts may have different views about how particular situations should be approached;

- SPIRE therefore provides users with both factual information about national and local organisations or departments within the University dealing with, for example, accommodation and transport, and advice from a range of different experts about how to tackle particular issues.

3.1 System Users

At the beginning of the SPIRE project, it was envisaged that users of the system would fall mainly into three categories as described below. It was acknowledged that these users would differ considerably in terms of their knowledge of the domain (disability and higher education), their experience with computers, and their information needs.

i) *Students* It was estimated that both potential students and current students might use the system to discover, typically, things like what facilities the University has to offer and what support the University provides for disabled students. In addition, students might require more specific information from the system such as the name and contact details of the person responsible for providing accommodation for disabled students, or the route for wheelchair access to a particular lecture theatre. Students have varying levels of computer literacy, and may have extensive knowledge of some aspects of the domain (for example, those relating to a particular disability) but not others (relating, for example, to University policy and organisation).

ii) *Academic Staff* Judging from their use of existing information services, it was thought that academic staff usage of the system would commonly centre around specific students whose needs are to be supported. Academic staff often require information about the needs of students with particular disabilities, or advice on how to best support a specific student in a lecture, practical or tutorial setting. Levels of computer literacy vary amongst academic staff, as do levels of knowledge about the domain - many will have had no previous experience of teaching students with disabilities, and may, therefore, know little about the University's support facilities.

iii) *Support Staff* University support staff (such as disabled student co-ordinators, or workers in the Health Centre) also typically require information which will help them cater for particular students. Support staff may, for example, need to know particular details about the University's admissions procedure, or the address of an external charity or equipment supplier.

5

They are typically knowledgeable about the domain of disability and higher education, but many have little or no experience of using computers to retrieve information.

3.2 Usage Patterns

Use of the system varies from casual use by some users (for example a prospective student might use the system only once when attending a University open day to find out more about the facilities the University has to offer), to frequent, more goal-directed use by others (for example a member of University support staff might repeatedly use the system to retrieve contact details of either University members of staff or external organisations).

Because of the combination of casual and frequent users, and the variety of their information needs, SPIRE was designed to support both browsing and searching strategies for accessing information. Also, since some users are expert in the domain (and therefore wish to use a searching strategy to directly access the information they require), but are uncomfortable with formulating their search criteria in terms of a specific query, the system supports users in constructing queries by making a number of simple choices.

The following section describes the way in which the GASA was used to respond to the particular challenges of the SPIRE project.

4 The Generic Advice System Architecture

To address the requirements of the SPIRE project regarding the range of expected system users and usage patterns, the design team developed a three-layered Generic Advice System Architecture (GASA). The SPIRE system can, effectively, be seen as a particular instantiation of this architecture. An overview of the model is given in figure 1.

In the GASA, a hypertext 'Advice Layer' is used to model the subjective expert advice in an advice system, whilst a more traditional, hierarchically organised relational database ('Data Layer') is used to model the objective data and factual information. To help support search and query access methods and to capitalise on novice users' existing domain knowledge, a third, top-level layer, was added to the architecture. The structure of this layer was derived from an analysis of domain tasks. This 'Task Layer', in conjunction with a User Profiling (filtering) mechanism, provides a guided route through the system.

The following paragraphs describe the SPIRE instantiation of the GASA. The main features of the architecture are described in the order in which they would appear to a user following the guided route through the SPIRE system. In parallel with the description of the model, a sample user interaction with the system is described in order to provide the reader with a realistic example of the model's use.

4.1 The User Profile

The User Profile is provided for casual, untrained users of the system who may have little experience in using computers for information retrieval. It is the first component of the GASA that a user encounters when following the guided route through the SPIRE system. The User Profile acts as a filtering mechanism for both the Task and Advice Layers, by allowing the user to select certain criteria which are used by the system to narrow the information space that the user will subsequently be able to browse.
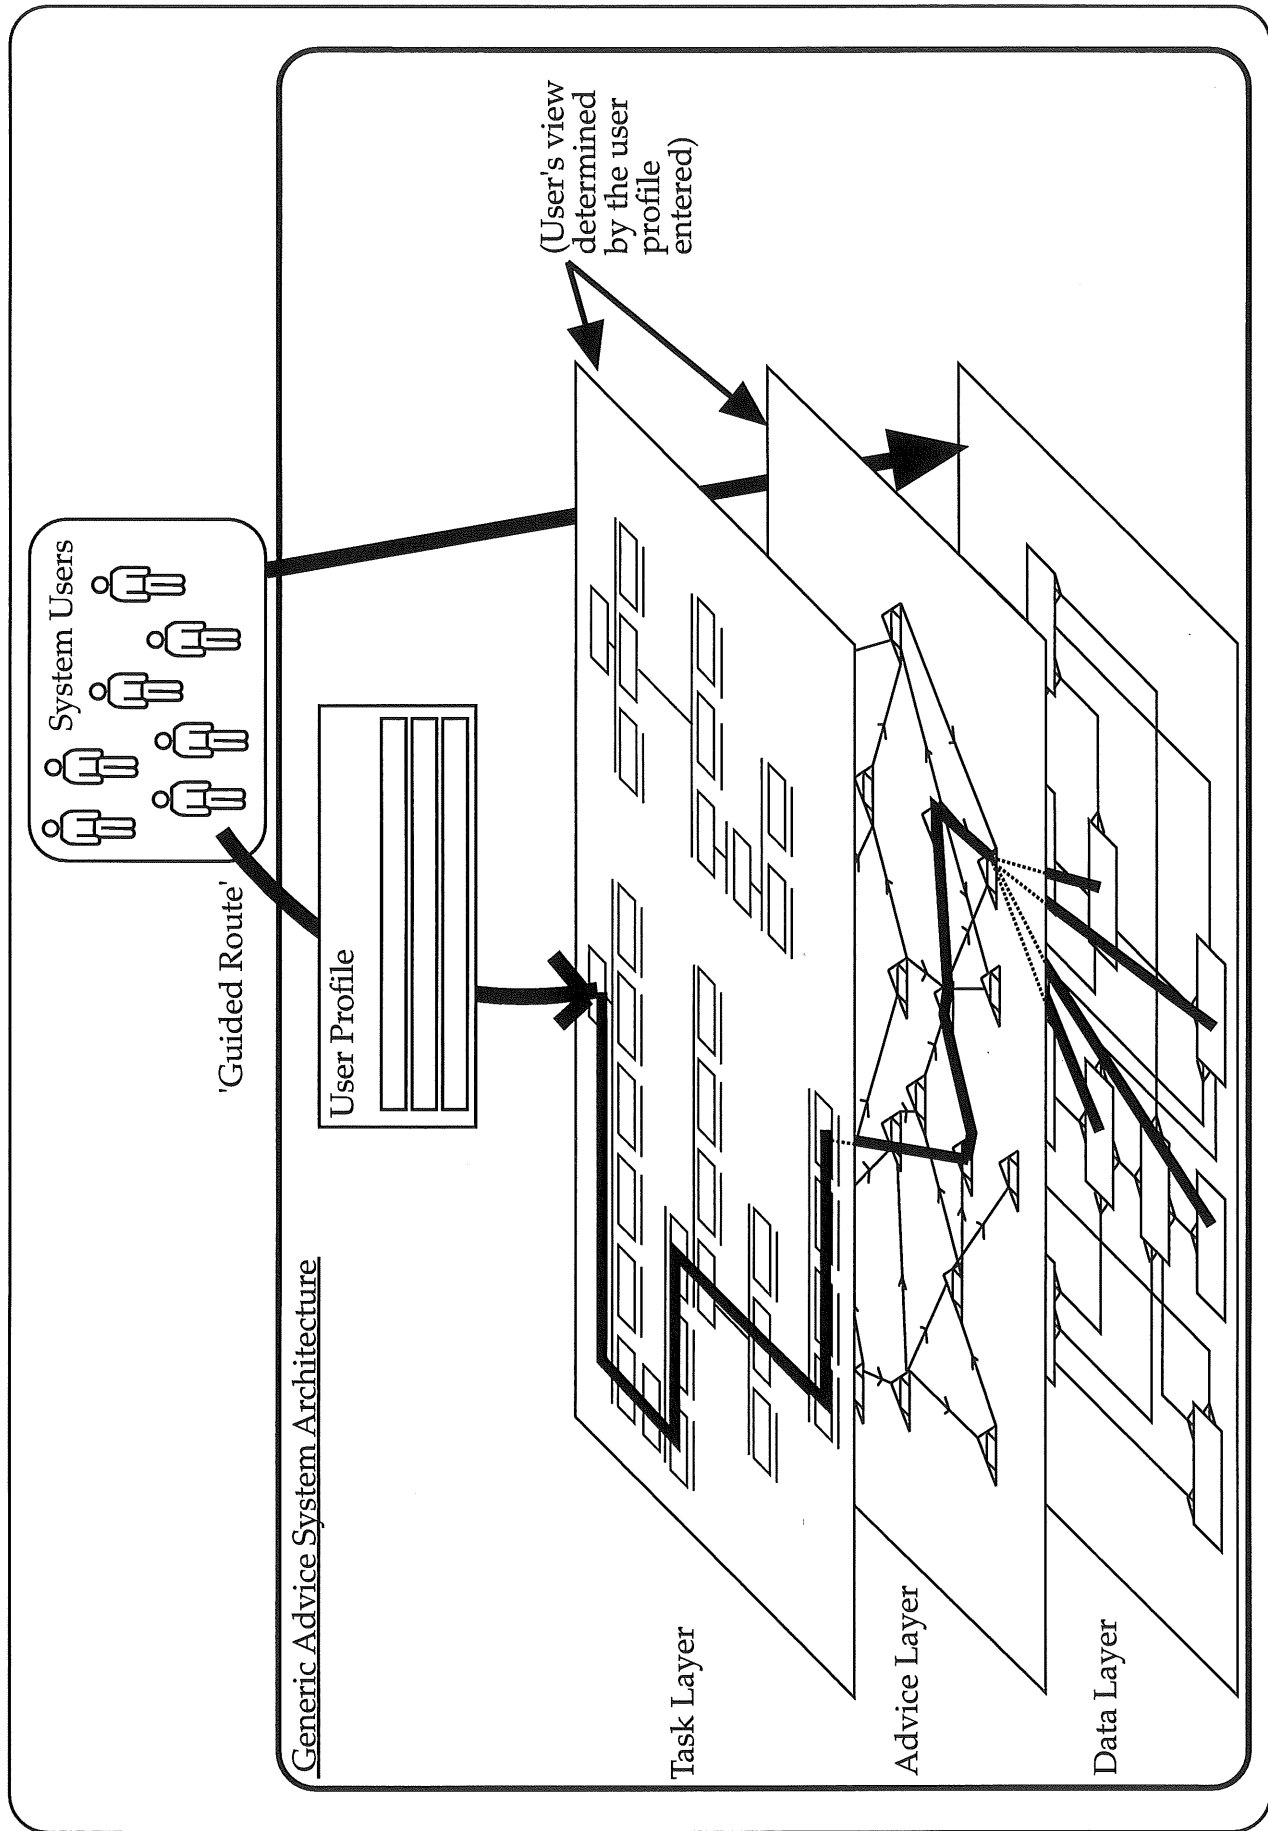
Fig. 1. Generic Advice System Architecture

The elements used in the User Profile should allow users to draw on their existing knowledge of the domain in structuring their use of the system. They should correspond to search criteria which are likely to remain unchanged during a particular user's interaction with the system. In the SPIRE system, the User Profile consists of three elements: User Role, Disability, and Subject. For each element the user may select a value from a list which is then used by the system in selecting relevant views of the Task and Advice Layers. In SPIRE, the possible values for each element are:

User Role: Current Student, Potential Student, Academic Staff, Support Staff
Disability: All, Blind/Visually Impaired, Deaf/Hard of Hearing, Language/Speech Difficulties, Medical Conditions, Mental Health Problems, Physical Disabilities, Specific Learning Difficulties
Subject: All, Accounting, Art and Design, Astronomy, Astrophysics, etc

For example, a visually impaired student who wishes to find out about how s/he might be able to read and answer a question paper in an exam situation might set the following User Profile :

User Role : Current Student.
Disability : Blind/Visually Impaired.
Subject : All.

This Profile would then be used to select appropriate views on the Task and Advice Layers of the system as described below.

Note that in the above example the Subject element is set to the value 'All'. Giving users the option to set one or more of the Profile elements to 'All' allows the user to choose how specific they wish the information filtering mechanism to be.

Had the visually impaired student been specifically interested in sitting exams in the subject of, for example, Physics, s/he could have specified this in the Profile, and thus narrowed the search even further.

4.2 The Task Layer

After selecting a User Profile, the next component of the GASA that a user following a guided route through the system encounters is the Task Layer. This layer is also intended to assist casual and untrained users in structuring their searches, and provides further filtering of information. As with the User Profile, the aim is to capitalise on the user's existing domain knowledge.

In the Task Layer, the user is presented with a hierarchical task analysis of the domain appropriate to the user role specified in the User Profile. Task analysis is a technique commonly used in user-centred or participative design (9) to capture information regarding the structure and sequence of tasks in the domain being studied. In filling out the GASA for a particular system, designers must generate hierarchical domain task analyses for each of the user roles identified for the User Profile. In the SPIRE system, separate task hierarchies were created for students, academic staff, and support staff in higher education.

The relevant task hierarchy is initially presented to the user as a list of top level tasks. Returning to the example of a visually impaired student who wishes to find out about how s/he might be able to read and answer a question paper in an exam situation, a selection of the top level tasks presented after selecting the 'student' user role are as follows :

Get to Campus
Attend Lecture
Attend Tutorial
Use Library
Use Computer Centre
Do Coursework
Take Exam

Use Media Services
                    Do Field Work
                    Use Sports Facilities etc

The user can select a top level task and see a decomposition of that task into it's component sub-tasks. So, for example, if the user selects the task 'Take Exam' the decomposition is as follows (where levels of indentation represent lower levels in the task hierarchy) :

            Take Exam
                    Get to lecture theatre/classroom
                            Get to place on campus
                    Get to seat/position
                    Read Question
                            Read Text
                    Answer Question
                            Write Text
                            Draw Diagrams
                    Observe timing device
                    Listen

The user can progressively disclose lower levels of the task hierarchy by selecting the tasks which best match their search criteria. In the example above, the user might select the following task path :

            Take Exam: Read Question: Read Text

The combination of User Profile and task path is then used by the system to select an appropriate view on the Advice Layer as described below.

4.3 The Advice Layer

The Advice Layer holds the subjective, textual, advice information in the system. It contains 'nodes' of advice relating to the domain of the system (in the case of SPIRE, that of disability in higher education), which are linked together to form a hypertext network.

Each advice node is a frame-like structure consisting of a header, a body, and a footer. The advice node header contains the title of the piece of advice, the author and source of the advice, and a number of slots which correspond to the elements of the User Profile. The advice node body contains the advice text itself, and the footer contains links to other related advice nodes as well as to points in the Data Layer of the GASA. An empty advice node frame from the SPIRE system is shown in Figure 2.

When a user sets values in the User Profile, the system applies a matching algorithm to each prospective advice node in turn, in order to select an appropriate subset of advice nodes to include in that user's view of the Advice Layer. Thus, each time a new User Profile is selected a new, virtual, Advice Layer is created by the system by matching the User Profile against the corresponding slots in each advice node. For example, in the case of the visually impaired student wishing to discover information about sitting exams, selecting the 'Student:Blind/Visually Imapaired:All' Profile causes the system to create an Advice Layer consisting of all advice nodes with matching values in the appropriate header slots.

The use of the key word 'All' for a User Profile element (either in the Profile itself or in a corresponding advice node slot) is important as it provides flexibility both for the browser and for the authors of an Advice Layer. In the case of the browser, specifying 'All' for a User Profile element allows users to widen their searches in the Advice Layer. In the case of the author, specifying 'All' in an advice node slot allows the author of the advice node to widen the applicability of the advice they have written.
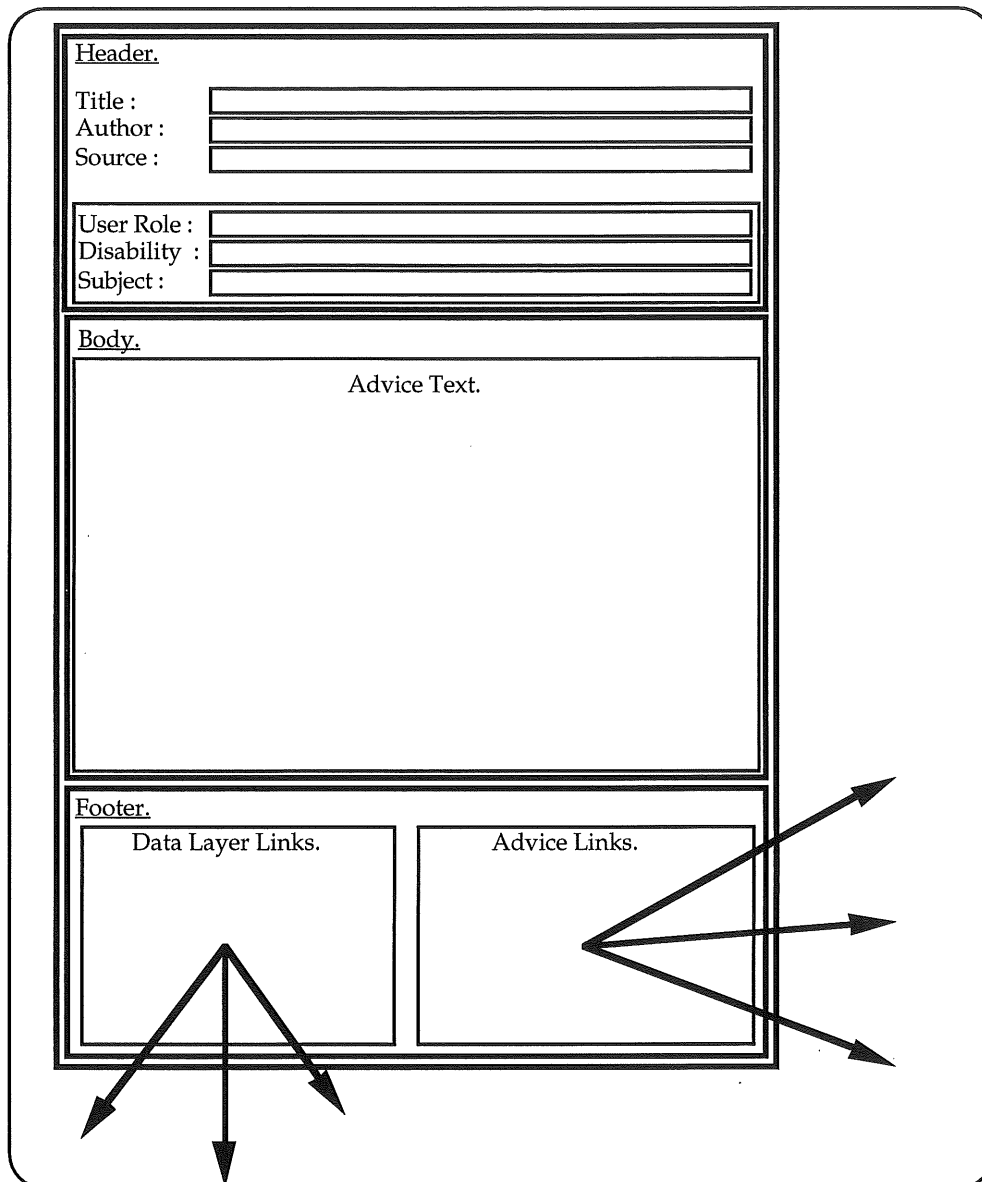
9

Figure 2. SPIRE Advice Node Frame.

For each task in the Task Layer, there is a corresponding set of advice nodes. Thus, by selecting a particular task from the Task Layer, users may further home in on information of interest. Having selected a task path in the Task Layer the user is able to 'drop' into the Advice Layer at an appropriate point. At this point, the user is typically presented with a handful of advice nodes of special relevance to the combination of User Profile and task path selected. S/he is then free to browse around the Advice Layer by following the advice links to other semantically related advice nodes in the network. Alternatively, each advice node also contains links to the bottom (Data) layer of the system which the user may also traverse.

Returning once again to the current example, having selected the 'Take Exam: Read Question: Read Text' task path the user is presented with one of a small selection of advice nodes for this task path. An example of an appropriate node in this case is shown in figure 3. It contains advice written by a domain expert who was, in this case, an ex-student of the University who is blind.

From this advice node the user can either traverse one of the advice links to access other advice nodes ('Take Exam: Observe Timing Device', or 'Finding personal readers'), or traverse a Data Layer link ('Kurzweill Reader', or 'Braille Embosser') in order to access the information in the Data Layer of the

SPIRE system where s/he may view factual information regarding pieces of equipment that a visually impaired person might use to read text in an exam.
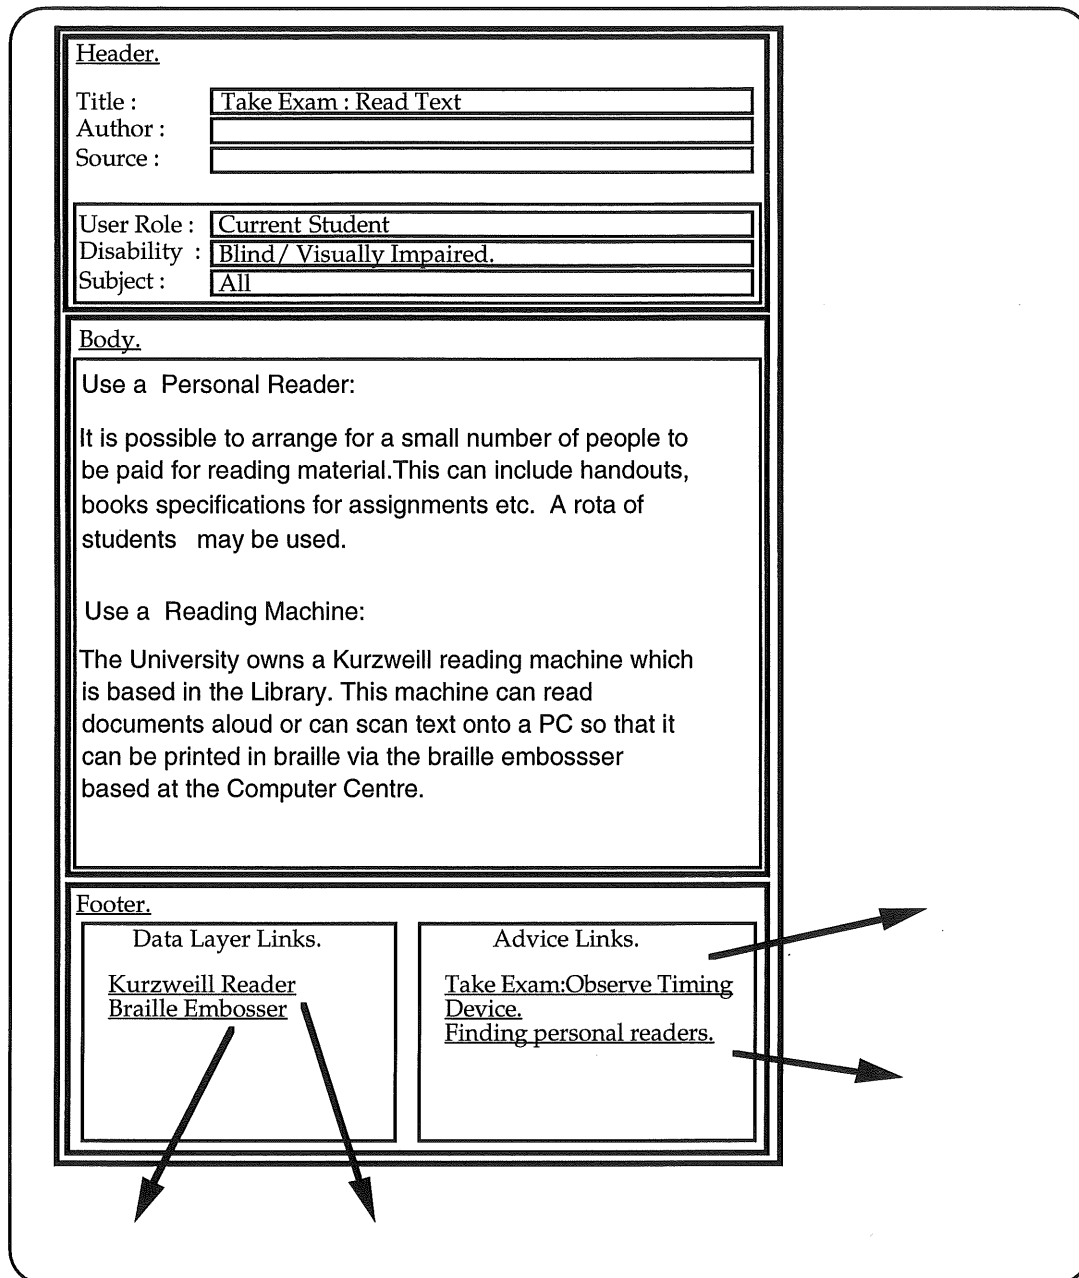

Figure 3. A SPIRE advice node.

Low level tasks such as 'Read Text' which appear at several points in the task hierarchy can refer users to the same pieces of advice regardless of their position in the task hierarchy as a whole. However, it is also possible to ask for advice specific to performing a task in a particular context so that, for example, a user may request information about reading text in an exam, rather than in the library or in a private room.

4.4 The Data Layer

The Data Layer is used to store the objective, hierarchically-structured, data held in an advice system. In this respect the Data Layer resembles a hierarchically presented standard relational database of the

sort commonly found in more traditional information systems. The Data Layer may also contain multimedia information including video, sound and animation.

The Data Layer can be accessed in one of two ways. A user may follow the guided route through the system described above, or alternatively, s/he may access the Data Layer directly from the top level of the model through a standard hierarchical menu-based interface (see Figure 1). This second route is provided for the more confident or experienced user who is familiar both with the domain, and with standard computer-based information retrieval mechanisms, and who has a specific information need. Thus this second route provides the facilities for direct searching of the system rather than browsing.

The Data Layer in SPIRE can be directly accessed through a hierarchical menu containing the following top level items:

> INFORMATION SOURCES
> > People
> > Organisations
>
> FINANCE
> EQUIPMENT
> ADMISSIONS
> STUDENT SERVICES
> ACCESS
> ACCOMMODATION
> SUBJECT SPECIFIC INFORMATION

Each top level item leads to a category selection which allows the users to find a specific record in the database. For example, choosing 'Accommodation', followed by selecting the name of a University campus will allow the user to view information regarding the accommodation available for disabled students at the chosen campus.

As described above, the other method of accessing the Data Layer is by traversing a Data Layer link from a node in the Advice Layer. Returning to the example of the visually impaired student, the advice node shown in Figure 3 contains two direct links into the Data Layer, both of which link to records in the 'Equipment' section of the data base. Selecting a Data Layer link whilst in the Advice Layer allows the user to 'drop' directly to a relevant point in the underlying data layer. Once in the Data Layer s/he can also navigate around the Layer to find other relevant pieces of information. For example if the user had selected 'Kurzweill Reader' from the advice node in Figure 3 s/he would be presented with information about that piece of equipment including a picture and a video clip of the piece of equipment in action. Having seen that information, s/he could then navigate around the Data Layer to find other similar pieces of equipment, or to locate other relevant pieces of information in the database.

4.5 Implementation

The implementation of the SPIRE system was carried out entirely in Microsoft FoxPro 2.5 for Macintosh, a relational database management system (RDBMS) which includes an application development toolkit. FoxPro fully supports the storage and display of all the multimedia elements used in SPIRE's Data Layer. Two further reasons why FoxPro was chosen as the development environment are as follows. Firstly, unlike most RDBMS available for the Apple Macintosh, FoxPro allows developers to create royalty free runtime versions of their systems which can be distributed as standalone applications. Secondly, because FoxPro is a cross-platform application running on both Apple Macintosh and PC platforms, applications can be ported between the two common platforms. SPIRE was initially developed on a Mac but later successfully ported to run under Microsoft Windows Vs. 3.1 using Microsoft FoxPro Vs 2.6 for Windows.

Using a RDBMS to implement the GASA model was an interesting exercise, particularly when it came to implementing the hypertext structure of the Advice Layer within the confines of the relational model.

12

There are certain advantages to implementing a hypertext network within a relational model, some of which were outlined by Shutt and Streitz (10). The referential integrity that can be enforced with a RDBMS can be used to ensure that the integrity of the hypertext network is maintained. This facility may be used to avoid the problem of creating 'dangling' links when a node is deleted. Links within SPIRE are implemented separately from advice nodes and are then associated with those nodes by the use of foreign keys. Storing link information separately from nodes elevates links to 'first class objects' within the system, so that they may have their own attribute sets. This opens up the possibility of extending and enhancing the Advice Layer in future systems: for example, it would be possible to store information regarding creation dates, relevance to different user roles, or the number of times that links are traversed.

Using a RDBMS for the SPIRE system also made the User Profile filtering mechanisms a lot easier to implement. Assigning each advice node a list of attributes corresponding to the elements of the User Profile meant that the process of creating a virtual Advice Layer for each Profile selected by the user was made fairly simple. The routine used to match advice nodes to the chosen User Profile consisted entirely of compound SQL statements. The query optimisation routines used by FoxPro also ensured that the filtering process was fairly rapid, even when several hundred advice nodes had to be matched.

It has been argued (10,11), that Object Oriented Data Base Management Systems (OODBMS) present a more logical model for the implementation of a hypertext network than relational models. Whilst this may be the case, practical considerations relating to the availability of OODBMS can mean that using a RDBMS is a more sensible solution. We found when creating SPIRE that the relational model provided all the means necessary to fully implement the GASA model.

5 Discussion

This section discusses the Generic Advice System Architecture in relation to some of the challenges identified in section 2.

5.1 Users with Little Training

A general requirement for advice systems is that they should be usable by users with little or no training and limited knowledge of the relevant information domain. The User Profile and Task Layer of the GASA address this requirement by providing the user with the ability to construct queries implicitly through selection of appropriate values from structured lists of key terms. This ability to construct queries by making simple choices is particularly important in advice systems where users may be expert in the domain of the system but unable to formulate their search criteria in terms of a specific query. It is also important in cases where users have only a limited knowledge of the information domain as it relies on recognition and comprehension of key terms, rather than recall and complex query construction.

5.2 Diverse Information Exploration Strategies

Advice systems must also be usable by a diverse set of users with diverse information exploration and problem-solving strategies which may involve the use of both browsing and searching. As described in section 4, the GASA model combines hypertext, database, and data access layers in such a way as to support both browsing and searching. Novice users who are unsure as to their information needs can browse a GASA-based system by following a guided route through each layer. More expert users with more specific information needs can search the system by going directly into the Data Layer, or by selecting more specific search criteria using the User Profile.

The User Profile mechanism can also be seen as a way of creating virtual hypertext networks (Advice Layers) in the manner described by Halasz (7). The advantage of creating virtual hypertext structures in this way is that they are dynamic and are therefore able to re-configure themselves in response to changing information needs or search criteria.

5.3 'Point and Select' Access to Information

As described in section 2, advice systems often need to be installed as touch screen systems and must therefore be operated solely through the use of pointing and selection. The aim of the GASA is to present the user with pre-defined lists of key terms to enable them to search a system simply by selecting appropriate choices from the User and Task Layers. The Advice and Data Layers may also be navigated solely by selection of appropriate key terms. No keyboard input is necessary.

5.4 Disorientation and Cognitive Overhead

Conklin has suggested that using the use of a filtering mechanism to reduce the number of nodes in a hypertext network can a have a positive effect on cognitive overhead and disorientation problems (8). The GASA addresses this issue by providing a User Profile as a filter so that the user views only those nodes in the Advice Layer that are relevant to his or her current query, thus reducing the number of navigation options available at any one time, and cutting down on cognitive overhead.

It has been suggested (12) that using more hierarchical structures in an information system can be an aid to navigation and can reduce user disorientation. The GASA model addresses the problem of disorientation by 'sandwiching' the hypertext (Advice) layer of the model between two more structured layers, the Task and Data Layers. The hierarchical structuring of the Task and Data Layers provides the user with a place to retreat to should they get lost in the Advice Layer.

Finally, it is well known that disorientation can also be reduced by 'land marking', or providing the user with an indication of where s/he is in a system by the use of visual cues. In the SPIRE version of the GASA model the user is made aware of which Layer of the model s/he is currently traversing by the use of different coloured screen backgrounds in the Task, Advice, and Data Layers. Colours are also used to indicate links between Layers, so, for example, a link from an Advice Layer node to a point in the Data Layer is shown in the colour of the Data Layer.

5.5 Support for Structured Data

While complex and weakly structured advice information is held in the GASA in the form of a hypertext network (the Advice Layer), the objective data component of an advice system is held in the Data Layer, which provides all the structuring of a relational database.

6 Related Work.

Other attempts to combine hypertext and database models have mainly concentrated on using a hypertext layer as a front end to an existing database. For example, the HIFI project (13) uses a hypertext layer to add navigational access to the data in a relational database, as does the Hypermuse system (5). While these approaches do have something in common with the GASA model, a major difference is that in the GASA there is a distinction between the different types information that are held in the hypertext and database components of the system. Unlike HIFI or Hypermuse, the GASA actually holds a separate class of information (advice) in the hypertext layer, rather than using the hypertext layer merely as a front end to an underlying database.

A number of other layered hypertext models also share some of the GASA's features. Such models typically separate hypertexts into an index layer which is used to index, search, and navigate through the hypertext, and a document layer, which stores the data. Examples of other models which use this type of classification are the Hyperindices model (6), which is divided into 'Hyperindex' and 'HyperBase' layers, the model used by the Dynamic medical handbook project (14), which separates the 'Document Space' from the 'Index Space', and Lucarella's model (3) which incorporates a 'Concept Network' and 'Document Network'. These models have many similarities to the GASA in that they all separate the navigation and search semantics of a hypertext network from the actual information in the system. In addition, each model recognises the importance of augmenting the basic navigational

method of accessing a hypertext with search and query based mechanisms. Lucarella in particular stresses the benefits that can be gained from the interaction between standard information retrieval and hypertext models and methods (3). In common with the GASA, certain models (3,6) also allow the user to apply a filter to the hypertext network as a way of narrowing the information space they are browsing, though the GASA is unique in structuring the filtering mechanism around a model of the user's knowledge about the domain in question.

7 Conclusions

The structure of the Generic Advice System Architecture was found to be useful in addressing a number of challenges presented by the design and implementation of SPIRE, a hypermedia advice system which presents users with both expert advice and factual information about a complex domain and weakly structured domain. Following the model provided by the GASA, advice in SPIRE is presented in the form of a hypertext network which provides the user with numerous links into a multimedia database of factual information. Access to this information is further facilitated by the filtering mechanisms provided by the GASA's User Profile and Task Layers.

Preliminary evaluations of the system suggest that users of SPIRE find it easy to learn how to use the system, and feel confident that they can get at the information they want. The system is now in use in a number of institutions in the UK. On the basis of our experiences on the SPIRE project, we believe that the GASA may also be useful in developing hypermedia advice systems in a number of other complex domains in which factual information can be usefully supplemented by expert advice in order to assist members of the general public in finding the answers they are seeking, even when they are not initially clear about what questions they want to ask.

Acknowledgements

References

(1) LEVENTHAL, L., INSTONE, K. and TEASLEY, B. Hypertext-based kiosk systems: Seven challenges and an empirical study. *In:* Proceedings of EWHCI94, East-West International Conference on Human-Computer Interaction, volume II, August 2 - 6th, St. Petersburg, 1994, 109 - 114.

(2) HOLFELDER, W. and HEHMANN, H. A networked multimedia retrieval management system for distributed kiosk applications. *In:* Proceedings of the International Conference on Multimedia Computing and Systems: IEEE Computer Society Press, 1994, 342 - 351.

(3) LUCARELLA, D. A model for hypertext based information retrieval, *In:* Proceedings of ECHT'90, the European Conference on Hypertext, 1990, 81-94.

(4) CARMEL, E., CRAWFORD, S., and HSINCHUN, C. Browsing in hypertext: A cognitive study, *IEEE Transactions on Systems Man and Cybernetics*, 22(5),1992, 865-883.

(5) POLTER, A., SARGENT, G. and FAHY, A. Hypermuse : A prototype hypermedia front-end for museum information systems, *Hypermedia*, 5(3), 1993, 165-186.

(6) BRUZA, P.D. Hyperindices : A novel aid for searching in hypermedia, *In:* Proceedings of ECHT'90, the European Conference on Hypertext, 1990, 109-122.

(7) HALASZ, F. Reflections on NoteCards : Seven issues for the next generation of hypertext systems, *Communications of the ACM,* 31(7), July 1988, 836-852.

(8) CONKLIN, J. Hypertext : An introduction and survey, *IEEE Computer,* September 1987, 17-41.

(9) JOHNSON, P. *Human-Computer Interaction,* Mcgraw-Hill, 1992.

(10) SHUTT, H.A. and STREITZ, N.A. HyperBase : A hypermedia engine based on a relational data base management system, *In:* Proceedings of ECHT'90, the European Conference on Hypertext, 1990, 95-108.

(11) SMITH, K. and ZDONIK, S. Intermedia : A case study of the differences between relational and object-oriented database systems", *In:* Proceedings of OOPSLA 87, ACM, 1987.

(12) AKSCYN, R. YODER, E. and McCRACKEN, The data model is the heart of interface design, *In:* Proceedings of CHI'88, ACM, 1988.

(13) CAVALLARO, U., GARZOTTO, F., PAOLINI, P. and TOTARO, D. HIFI : Hypertext interface for information systems, *IEEE Software,* November 1993, 48-51.

(14) FRISSE, M.E. and COUSINS, S.B. Information retrieval from hypertext : Update on the dynamic medical imaging project, *In:* Proceedings of Hypertext '89, ACM Press, 199-212.